

THE FLUTED FRAGMENT REVISITED

IAN PRATT-HARTMANN, WIESŁAW SZWAST, AND LIDIA TENDERA

Abstract. We study the fluted fragment, a decidable fragment of first-order logic with an unbounded number of variables, motivated by the work of W. V. Quine. We show that the satisfiability problem for this fragment has nonelementary complexity, thus refuting an earlier published claim by W. C. Purdy that it is in NEXPTIME. More precisely, we consider \mathcal{FL}^m , the intersection of the fluted fragment and the m -variable fragment of first-order logic, for all $m \geq 1$. We show that, for $m \geq 2$, this subfragment forces $\lfloor m/2 \rfloor$ -tuply exponentially large models, and that its satisfiability problem is $\lfloor m/2 \rfloor$ -NEXPTIME-hard. We further establish that, for $m \geq 3$, any satisfiable \mathcal{FL}^m -formula has a model of at most $(m - 2)$ -tuply exponential size, whence the satisfiability (= finite satisfiability) problem for this fragment is in $(m - 2)$ -NEXPTIME. Together with other, known, complexity results, this provides tight complexity bounds for \mathcal{FL}^m for all $m \leq 4$.

§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. Fluted formulas arise naturally as first-order translations of quantified English sentences in which no quantifier-rescoping occurs, thus:

No student admires every professor
 $\forall x_1(\text{student}(x_1) \rightarrow \neg \forall x_2(\text{prof}(x_2) \rightarrow \text{admires}(x_1, x_2)))$; (1)

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

The origins of the fluted fragment can be traced to a paper given by W. V. Quine to the 1968 *International Congress of Philosophy* [16], in which the author defined what he called the *homogeneous m -adic formulas*. In these formulas, all predicates have the same arity m and all atomic formulas have the same argument sequence x_1, \dots, x_m . Boolean operators and quantifiers may be freely applied, except that the order of quantification must follow the order of arguments: a quantifier binding an occurrence of x_i may only be applied to a subformula in which all occurrences of x_{i+1}, \dots, x_m are already bound. Quine explained how Herbrand's decision procedure for monadic first-order logic extends to cover all homogeneous m -adic formulas.

Received July 31, 2018.

2010 *Mathematics Subject Classification.* 68Q17.

Key words and phrases. first-order logic, decidability, satisfiability, transitivity, complexity.

© 2019, Association for Symbolic Logic
0022-4812/19/8403-0008
DOI:10.1017/jsl.2019.33

The term *fluted logic* first appears (to the present authors' knowledge) in Quine [18], where the restriction that all predicates have the same arity is abandoned, a relaxation which, according to Quine, does not affect the proof of decidability of satisfiability. The allusion is presumably architectural: we are invited to think of arguments of predicates as being 'lined up' in columns. Quine's motivation for defining the fluted formulas was to locate the boundary of decidability in the context of his reconstruction of first-order logic in terms of *predicate-functors*, which Quine himself described as a 'modification of Bernays' modification of Tarski's cylindrical algebra' [17, p. 299]. Specifically, the fluted fragment can be identified by dropping from full predicate functor logic those functors associated with the permutation and identification of variables, while retaining those concerned with cylindrification and Boolean combination.

Notwithstanding its predicate-functorial lineage, the fluted fragment has, as we shall see, a completely natural characterization within the standard régime of bound variable quantification, and thus constitutes an interesting fragment of first-order logic in its own right. In fact, \mathcal{FL} overlaps in expressive power with various other such fragments. For example, *Boolean modal logic* (Lutz and Sattler [7]) maps, under the standard first-order translation, to \mathcal{FL} —in fact, to \mathcal{FL}^2 , the fluted fragment restricted to just two variables. On the other hand, even \mathcal{FL}^2 is not contained within the so-called *guarded fragment* of first-order logic (Andréka, van Benthem, and Némethi [1]): the formula (2), for example, is not equivalent to any guarded formula. A more detailed comparison of the fluted fragment to other familiar decidable fragments can be found in Hustadt, Schmidt, and Georgieva [6].

Quine never published a proof of his later claims regarding the fluted formulas; indeed, Noah [9] later claimed that, on the contrary, Herbrand's technique does not obviously extend from homogeneous m -adic logic as Quine suggested, and that consequently, the decidability of the satisfiability problem for the latter should be regarded as open. This problem—together with the corresponding problems for various related systems—was considered in a series of articles in the 1990s by W. C. Purdy [12–15]. The decidability of Purdy's \mathcal{FL} is proved in [13], while in [15, Corollary 10] it is claimed that this fragment 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 [15, Theorem 13] that the satisfiability problem for \mathcal{FL} is NEXPTIME-complete.

These latter claims are false. In the sequel, we show that, for $m \geq 2$, the fluted fragment restricted to just m variables, denoted \mathcal{FL}^m , can force models of $(\lfloor m/2 \rfloor)$ -tuply exponential size, and that its satisfiability problem is $(\lfloor m/2 \rfloor)$ -NEXPTIME-hard. It follows that there is no elementary bound on the size of models of satisfiable fluted formulas, and that the satisfiability problem for \mathcal{FL} is nonelementary.¹ On the other hand, we also show that, for $m \geq 3$, any satisfiable formula of the m -variable fluted fragment has a model of $(m - 2)$ -tuply exponential size, so that the satisfiability problem for this subfragment is contained in $(m - 2)$ -NEXPTIME. Thus, \mathcal{FL} has the finite model property, and its satisfiability (= finite satisfiability) problem is decidable, but not elementary. In the case $m = 2$, \mathcal{FL}^2 is contained within the 2-variable

¹An extended abstract with this result was published in [11].

fragment of first-order logic, whence its satisfiability problem is in NEXPTIME by the well-known result of Grädel, Kolaitis, and Vardi [4], which matches the lower bound reported above. The fragment \mathcal{FL}^0 is evidently the same as propositional logic, and so its satisfiability problem is NPTIME-complete; the fragment \mathcal{FL}^1 likewise coincides with the 1-variable fragment of first-order logic (over a signature of unary predicates), and hence also has an NPTIME-complete satisfiability problem, since all satisfiable formulas clearly have models of polynomially bounded size. Counting “0-tuply exponential” as a synonym for “polynomial,” we see that for $0 \leq m \leq 4$, \mathcal{FL}^m is $(\lfloor m/2 \rfloor)$ -NEXPTIME-complete. For $m > 4$, the above complexity bounds for \mathcal{FL}^m leave a gap between $(\lfloor m/2 \rfloor)$ -NEXPTIME and $(m - 2)$ -NEXPTIME.

We mention at this point another incorrect claim by Purdy concerning an extension of the fluted fragment. In Purdy [14], the author considers what he calls *extended fluted logic* (EFL), in which, in addition to the usual predicate functors of fluted logic, we have an *identity functor* (essentially: the equality predicate), *binary conversion* (the ability to exchange arguments in binary atomic formulas) and *functions* (the requirement that certain specified binary predicates be interpreted as the graph of a function). Purdy claims (Corollary 19, p. 1460) that EFL has the *finite model property*: if a formula of this fragment is satisfiable, then it is satisfiable over a finite domain. But EFL evidently contains the formula

$$\forall x_1 \forall x_2 (r(x_1, x_2) \rightarrow f(x_1, x_2)) \wedge \exists x_1 \forall x_2 \neg r(x_1, x_2) \wedge \forall x_1 \exists x_2. r(x_2, x_1),$$

where f is required to be interpreted as the graph of a binary function; and this is an axiom of infinity. In view of these observations, it seems only prudent to treat Purdy’s series of articles with caution.

An independent, resolution-based decision procedure for the fluted fragment was presented by Schmidt and Hustadt [19]. No complexity bounds are given. Moreover, that article omits detailed proofs, and these have, to the present authors’ knowledge, never been published. The logic \mathcal{FL} is rather similar in spirit to the little-known *ordered fragment* of first-order logic, whose satisfiability problem was shown to be decidable in [5]. The decision procedure given there is based on a polynomial time translation to the modal logic D of serial frames, whence the satisfiability problem is in PSPACE. In the ordered fragment, predicate arguments are lined up ‘the other way’: thus, for example, $\forall x_1 \exists x_2 (p(x_1) \wedge q(x_1, x_2))$ is ordered but not fluted, while $\forall x_1 \exists x_2 (p(x_2) \wedge q(x_1, x_2))$ is fluted but not ordered. It is, however, straightforward to translate any ordered formula into an equivalent fluted formula.

In the sequel, we show that, for $m \geq 3$, the satisfiability problem for \mathcal{FL}^m is in $(m - 2)$ -NEXPTIME. Specifically, we use a model-construction-based technique to show that any satisfiable formula of \mathcal{FL}^3 has a model of size bounded by an exponential function of the size of φ ; and we use resolution theorem proving to reduce the satisfiability problem for \mathcal{FL}^m ($m \geq 4$) to the corresponding problem for \mathcal{FL}^{m-1} , at the cost of an exponential increase in the signature and the size of the formula. Our proof is shorter and more perspicuous than the arguments for the decidability of the satisfiability problem for \mathcal{FL} given in either Purdy [13] or Schmidt and Hustadt [19], and yields better complexity bounds than could—in the absence of nontrivial refinements—be derived from those approaches.

The high complexity of \mathcal{FL} is perhaps remarkable when compared to other well-known decidable first-order fragments. For example, the classical decidable

quantifier-prefix fragments [2], the guarded fragment [3] and the two-variable fragment (mentioned above) all have elementary complexity. One first-order fragment that comprises a similar hierarchy of hard problems to \mathcal{FL} , however, is the recently discovered *separated fragment* [23].

The structure of this article is as follows. Section 2 gives some basic definitions. In Section 3, we show that formulas of \mathcal{FL}^{2m} can force models of m -tuply exponential size, and indeed that the satisfiability problem for \mathcal{FL}^{2m} is m -NEXPTIME-hard, thus disproving the results claimed in Purdy [15]. In Section 4, we show that, for $m \geq 3$, any satisfiable formula φ of \mathcal{FL}^m has a model of size at most $(m - 2)$ -tuply exponential in the size of φ , and hence that the satisfiability (= finite satisfiability) problem for \mathcal{FL}^m is in $(m - 2)$ -NEXPTIME.

§2. Preliminaries. Fix a sequence of variables $\bar{x}_\omega = x_1, x_2, \dots$. Let σ be a purely relational signature—i.e., a signature containing predicates of any arity (including 0), but no function-symbols or individual constants. A *fluted atomic formula* (or: *fluted atom*) of $\mathcal{FL}_\sigma^{[k]}$ ($k \geq 0$) is an expression $p(x_\ell, \dots, x_k)$, where $1 \leq \ell \leq k + 1$, $p \in \sigma$ has arity $(k - \ell + 1)$, and x_ℓ, \dots, x_k is a contiguous subsequence of \bar{x}_ω . If $\ell = k + 1$, then p has arity 0—in other words, is a propositional letter. A *fluted literal* of $\mathcal{FL}_\sigma^{[k]}$ is either a fluted atom of $\mathcal{FL}_\sigma^{[k]}$ or the negation of such. We define the sets of formulas $\mathcal{FL}_\sigma^{[k]}$ (for $k \geq 0$) over σ by structural induction as follows: (i) any fluted atom of $\mathcal{FL}_\sigma^{[k]}$ is a formula of $\mathcal{FL}_\sigma^{[k]}$; (ii) $\mathcal{FL}_\sigma^{[k]}$ is closed under Boolean combinations; (iii) if φ is in $\mathcal{FL}_\sigma^{[k+1]}$, then $\exists x_{k+1}.\varphi$ and $\forall x_{k+1}.\varphi$ are in $\mathcal{FL}_\sigma^{[k]}$. We normally suppress reference to σ , writing $\mathcal{FL}^{[k]}$ for $\mathcal{FL}_\sigma^{[k]}$. In this context, a fluted atom of $\mathcal{FL}_\sigma^{[k]}$ will simply be called a *fluted k -atom*, and similarly for literals. Observe that any proposition letter is a fluted k -atom for all $k \geq 0$ (and hence is a fluted k -literal and indeed a formula of $\mathcal{FL}^{[k]}$). The set of *fluted formulas* is defined as $\mathcal{FL} = \bigcup_{k \geq 0} \mathcal{FL}^{[k]}$. A *fluted sentence* is a fluted formula over an empty set of variables, i.e., an element of $\mathcal{FL}^{[0]}$. Thus, when forming Boolean combinations in the fluted fragment, all the combined formulas must have as their free variables some contiguous subsequence of \bar{x}_ω ; and when quantifying, only the free variable with highest index may be bound. Note however that proposition letters may occur freely in fluted formulas. (See [13], p. 609.)

Denote by \mathcal{FL}^m the subfragment 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 m variables, free or bound) with $\mathcal{FL}^{[m]}$. For example, formulas (1) and (2) are in \mathcal{FL}^2 and \mathcal{FL}^3 , respectively, but both are in $\mathcal{FL}^{[0]}$.

To avoid tedious repetition of formulas, we write $\pm\varphi$ to stand ambiguously for the formulas φ and $\neg\varphi$. However, we adopt the convention that multiple occurrences of the symbol \pm in a displayed formula are all resolved in the same way. Thus, for example, the expression

$$\bigwedge_{i=0}^{n-1} \bigwedge_{\ell=0}^L \forall x_1 (\text{int}_1(x_1) \wedge \pm p_i(x_1) \rightarrow \forall x_2 \cdots \forall x_{\ell+1} \pm p_i^\ell(x_1, \dots, x_{\ell+1})),$$

stands for a *pair* of \mathcal{FL}^{L+1} -formulas: one formula in which all $2n(L + 1)$ occurrences of the symbol \pm are deleted, and another in which they are all replaced by the symbol \neg .

We make extensive use of the tetration function $t(k, n)$, defined, for $n, k \geq 0$, by induction as follows:

$$t(0, n) = n;$$

$$t(k + 1, n) = 2^{t(k, n)}.$$

Thus, $t(1, n) = 2^n$, $t(2, n) = 2^{2^n}$, and so on.

§3. Lower bound. In this section, we establish lower complexity bounds for the fluted fragment. Theorem 3.1 shows that an \mathcal{FL}^{2^m} -formula of size $O(n^2)$ can force models of size at least $t(m, n)$, thus contradicting Corollary 10 of Purdy [15]. Theorem 3.2 shows that the satisfiability problem for \mathcal{FL}^{2^m} is m -NEXPTIME-hard, thus contradicting Theorem 11 of Purdy [15].

As a preliminary, for any $z \geq 0$, we take the (canonical) representation of any integer n in the range $(0 \leq n < 2^z)$ to be the bit-string $\bar{s} = s_{z-1}, \dots, s_0$ of length z , where $n = \sum_{i=0}^{z-1} s_i \cdot 2^i$. (Thus, s_0 is the least significant bit.) Where z is clear from context, this representation is unique. Observe that, if, in addition, an integer n' in the same range is represented by s'_{z-1}, \dots, s'_0 , then $n' = n - 1 \pmod{2^z}$ if and only if, for all i ($0 \leq i < z$):

$$s'_i = \begin{cases} 1 - s_i & \text{if, for all } j \ (0 \leq j < i), s_j = 0; \\ s_i & \text{otherwise.} \end{cases}$$

This simple observation—effectively, the algorithm for decrementing an integer represented in binary—will feature at various points in the proof of the following theorem.

THEOREM 3.1. *For all $m \geq 1$, there exists a sequence of satisfiable sentences $\{\Phi_{m,n}\}_{n \in \mathbb{N}} \in \mathcal{FL}^{2^m}$ such that $\|\Phi_{m,n}\|$ grows polynomially with m and n (and indeed quadratically in n for fixed m), but the smallest satisfying model of $\Phi_{m,n}$ has at least $t(m, n)$ elements. Hence, there is no elementary bound on the size of models of satisfiable sentences in \mathcal{FL} .*

PROOF. The proof employs a well-known construction due to Stockmeyer [21, p. 112] for showing nonelementary complexity lower bounds, sometimes called the ‘yardstick/ruler construction’. For each $k \geq 0$, we write formulas that encode counters with values ranging from 0 to a tower of exponentials of height k . We construct these formulas iteratively and call the counters they define k -integers. The challenge we face is to carry out this construction within the confined syntax of \mathcal{FL} .

Fix positive integers m and n . Consider a signature $\sigma_{m,n}$ featuring:

- unary predicates p_0, \dots, p_{n-1} ;
- for all k in the range $1 \leq k \leq m$, a unary predicate int_k ;
- for all k in the range $1 \leq k < m$, binary predicates $\text{in}_k, \text{out}_k$.

(We shall add further predicates to $\sigma_{m,n}$ in the course of the proof.) When working within a particular structure, we call any element satisfying the unary predicate int_k in that structure a k -integer. Each k -integer, b , will be associated with an integer value, $\text{val}_k(b)$, between 0 and $t(k, n) - 1$. For $k = 1$, this value will be encoded by b ’s satisfaction of the unary predicates p_0, \dots, p_{n-1} . Specifically, for any 1-integer b ,

define $\text{val}_1(b)$ to be the integer canonically represented by the n -element bit-string s_{n-1}, \dots, s_0 , where, for all i ($0 \leq i < n$),

$$s_i = \begin{cases} 1 & \text{if } \mathfrak{A} \models p_i[b]; \\ 0 & \text{otherwise.} \end{cases}$$

On the other hand, if b is a $(k + 1)$ -integer ($k \geq 1$), then $\text{val}_{k+1}(b)$ will be encoded by the way in which the various k -integers are related to b via the predicate in_k . Specifically, for any k ($1 \leq k < m$) and any $(k + 1)$ -integer b , define $\text{val}_{k+1}(b)$ to be the integer canonically represented by the bit-string s_{N-1}, \dots, s_0 of length $N = t(k, n)$ where, for all i ($0 \leq i < N$),

$$s_i = \begin{cases} 1 & \text{if } \mathfrak{A} \models \text{in}_k[a, b] \text{ for some } k\text{-integer } a \text{ s.t. } \text{val}_k(a) = i; \\ 0 & \text{otherwise.} \end{cases}$$

We shall be interested in the case where \mathfrak{A} satisfies the following property, for all k ($1 \leq k \leq m$).

k-covering: The function $\text{val}_k : \text{int}_k^{\mathfrak{A}} \rightarrow [0, t(k, n) - 1]$ is surjective.

For technical reasons, we create a duplicate (mirror image) encoding of $\text{val}_{k+1}(b)$ in terms of the way in which b is related to the various k -integers via the predicate out_k . Specifically, we shall be interested in the case where \mathfrak{A} satisfies the following property, for all k ($1 \leq k \leq m$).

k-harmony: If $k > 1$, then, for all k -integers b and all $(k - 1)$ -integers a, a' in \mathfrak{A} such that $\text{val}_{k-1}(a) = \text{val}_{k-1}(a')$, $\mathfrak{A} \models \text{in}_{k-1}[a, b] \Leftrightarrow \mathfrak{A} \models \text{out}_{k-1}[b, a']$.

Let us pause to consider the properties of k -covering and k -harmony for various values of k . If $k < m$, k -covering ensures that, when we want to know what the i th bit in the canonical binary representation of a $(k + 1)$ -integer b is (where $0 \leq i < t(k, n)$), then *there exists* a k -integer a such that $\text{val}_k(a) = i$, and for which we can ask whether $\mathfrak{A} \models \text{in}_k[a, b]$. Conversely, $(k + 1)$ -harmony ensures that, if there are *many* k -integers a satisfying $\text{val}_k(a) = i$, then it does not matter which one we consult. For if $\text{val}_k(a) = \text{val}_k(a')$, then by two applications of $(k + 1)$ -harmony, $\mathfrak{A} \models \text{in}_k[a, b] \Leftrightarrow \mathfrak{A} \models \text{out}_k[b, a] \Leftrightarrow \mathfrak{A} \models \text{in}_k[a', b]$.

Our strategy will be to construct a satisfiable \mathcal{FL}^{2m} -formula $\Phi_{m,n}$ in the signature $\sigma_{m,n}$ such that any model $\mathfrak{A} \models \Phi_{m,n}$ satisfies k -covering and k -harmony for all k ($1 \leq k \leq m$). It then follows from m -covering that $|A| \geq t(m, n)$, proving the theorem. The signature $\sigma_{m,n}$ will feature several auxiliary predicates. In particular, we take $\sigma_{m,n}$ to contain:

- (i) the unary predicates $\text{zero}_1, \dots, \text{zero}_m$;
- (ii) the binary predicates $\text{pred}_{1,0}, \dots, \text{pred}_{m,0}$;
- (iii) the ternary predicates $\text{pred}_{1,1}, \dots, \text{pred}_{m-1,1}$;
- (iv) for all k ($1 \leq k \leq m$) and all ℓ ($0 \leq \ell \leq 2(m - k)$), the $(\ell + 2)$ -ary predicate $\text{eq}_{k,\ell}$.

Further predicates in $\sigma_{m,n}$ will be introduced later, as and when they are needed. Observe from (iii) that, setting $m = 1$, the list of ternary predicates in $\sigma_{1,n}$ is empty. This is as it should be: if $m = 1$, the formula $\Phi_{m,n}$ we want to construct must lie in \mathcal{FL}^2 , and thus may not use any ternary predicates. Observe also in this regard that, in (iv), as k increases from 1 to m , the maximal value of the index ℓ in the predicates

$\text{eq}_{k,\ell}$ decreases, in steps of 2, from $2m - 2$ down to 0; hence the maximal arity of these predicates decreases from $2m$ to 2, whence these predicates may all be used in \mathcal{FL}^{2m} -formulas.

We show that any model $\mathfrak{A} \models \Phi_{m,n}$ satisfies—in addition to k -covering and k -harmony—the following properties for all k ($1 \leq k \leq m$) concerning the interpretation of these predicates.

k-zero: For all k -integers b , $\mathfrak{A} \models \text{zero}_k[b] \Leftrightarrow \text{val}_k(b) = 0$.

k-equality: For all ℓ ($0 \leq \ell \leq 2(m - k)$), all k -integers b, b' and all ℓ -tuples of elements \bar{c} , $\mathfrak{A} \models \text{eq}_{k,\ell}[b, \bar{c}, b'] \Leftrightarrow \text{val}_k(b) = \text{val}_k(b')$.

k-predecessor: For all ℓ ($0 \leq \ell \leq \min(m - k, 1)$), all k -integers b, b' and all ℓ -tuples of elements \bar{c} , $\mathfrak{A} \models \text{pred}_{k,\ell}[b, \bar{c}, b'] \Leftrightarrow \text{val}_k(b') = \text{val}_k(b) - 1$, modulo $t(k, n)$.

Note that, in the definition of the property k -predecessor, \bar{c} is either the empty sequence or a singleton. Indeed, the bounds on ℓ amount to saying that ℓ takes values 0 or 1, except when $k = m$, in which case it takes only the value 0. (Recall that $\sigma_{m,n}$ does not feature the predicate $\text{pred}_{m,1}$.)

Thus, in a structure satisfying k -zero, $\text{zero}_k(x_1)$ can be read as “ x_1 is zero”; and in a structure satisfying k -predecessor, $\text{pred}_{k,0}(x_1, x_2)$, as “ x_2 is the predecessor of x_1 ”, and $\text{pred}_{k,1}(x_1, x_2, x_3)$ as “ x_3 is the predecessor of x_1 ”. Notice that, in the latter case, the argument x_2 is semantically inert. Similarly, in a structure satisfying k -equality, $\text{eq}_{k,\ell}(x_1, \dots, x_{\ell+2})$ can be read as “ x_1 is equal to $x_{\ell+2}$ ”, with the ℓ arguments $x_2, \dots, x_{\ell+1}$ again semantically inert. When naming predicates, we employ the convention that the first subscript, k , serves as a reminder that its primary arguments are typically assumed to be k -integers; the second subscript, ℓ , indicates that ℓ (possibly 0) semantically inert arguments have been inserted between the primary arguments.

To prove that any model $\mathfrak{A} \models \Phi_{m,n}$ satisfies the properties of k -covering, k -harmony, k -zero, k -equality and k -predecessor for all k ($1 \leq k \leq m$), we proceed by induction on k . For ease of reading, we introduce the various conjuncts of $\Phi_{m,n}$ as they are required in the proof. Appeals to the inductive hypothesis are indicated by the initials IH.

Base case ($k = 1$). Let b be a 1-integer, and recall that $\text{val}_1(b)$ is defined by b 's satisfaction of the predicates p_0, \dots, p_{n-1} . We proceed to secure the properties required for the base case of the induction. The property 1-harmony is trivially satisfied. We secure 1-zero by adding to $\Phi_{m,n}$ the conjunct

$$\forall x_1 (\text{int}_1(x_1) \rightarrow (\text{zero}_1(x_1) \leftrightarrow \bigwedge_{i=0}^{n-1} \neg p_i(x_1))). \tag{\Phi_1}$$

Thus, if b is a 1-integer, $\mathfrak{A} \models \text{zero}_1[b] \Leftrightarrow \text{val}_1(b) = 0$. To do the same for 1-predecessor and 1-equality, we proceed as follows. Letting $L = 2m - 1$, we add to $\sigma_{m,n}$ an $(\ell + 1)$ -ary predicate, p_i^ℓ , for all i ($0 \leq i < n$) and all ℓ ($0 \leq \ell \leq L$), and we add to $\Phi_{m,n}$ the corresponding pair of conjuncts

$$\bigwedge_{i=0}^{n-1} \bigwedge_{\ell=0}^L \forall x_1 (\text{int}_1(x_1) \wedge \pm p_i(x_1) \rightarrow \forall x_2 \cdots \forall x_{\ell+1} \pm p_i^\ell(x_1, \dots, x_{\ell+1})). \tag{\Phi_2}$$

Note that this really is a *pair* of formulas: all occurrences of the \pm sign must be resolved in the same way.

Then, for any 1-integer b and any ℓ -tuple \bar{c} from A ,

$$\mathfrak{A} \models p_i^\ell[b, \bar{c}] \Leftrightarrow \mathfrak{A} \models p_i[b]. \tag{3}$$

In effect, the conjuncts (Φ_2) append semantically inert arguments to each of the predicates p_i . This technique will be helpful at several points in the sequel, and we employ the convention that a superscript ℓ on a predicate letter indicates that the corresponding undecorated predicate has ℓ semantically inert arguments *appended* to its primary arguments. Note that p_i^0 is simply equivalent to p_i .

Now we can secure the property 1-equality. For all ℓ ($0 \leq \ell \leq 2m - 2$), let $\varepsilon_{1,\ell}(x_1, \dots, x_{\ell+2})$ abbreviate the formula:

$$\bigwedge_{i=0}^{n-1} (p_i^{\ell+1}(x_1, \dots, x_{\ell+2}) \leftrightarrow p_i(x_{\ell+2})).$$

We see from (3) that $\varepsilon_{1,\ell}(x_1, \dots, x_{\ell+2})$ in effect states that (for x_1 and $x_{\ell+2}$ 1-integers) the values of x_1 and $x_{\ell+2}$ are identical. We therefore add to $\Phi_{m,n}$ the conjuncts

$$\bigwedge_{\ell=0}^{2m-2} \forall x_1 (\text{int}_1(x_1) \rightarrow \forall x_2 \dots \forall x_{\ell+2} (\text{int}_1(x_{\ell+2}) \rightarrow \text{eq}_{1,\ell}(x_1, \dots, x_{\ell+2}) \leftrightarrow \varepsilon_{1,\ell}(x_1, \dots, x_{\ell+2}))). \tag{\Phi_3}$$

Thus, for any 1-integers b, b' in \mathfrak{A} and any ℓ -tuple \bar{c} from A ($0 \leq \ell \leq 2m - 2$), $\mathfrak{A} \models \text{eq}_{1,\ell}[b, \bar{c}, b'] \Leftrightarrow \text{val}_1(b) = \text{val}_1(b')$.

Turning to the property 1-predecessor, assume for the moment that $m > 1$, so that the predicates $\text{pred}_{1,0}$ and $\text{pred}_{1,1}$ are both in $\sigma_{m,n}$. For $0 \leq \ell \leq 1$, let $\pi_{1,\ell}(x_1, \dots, x_{\ell+2})$ abbreviate the formula

$$\bigwedge_{i=0}^{n-1} \left(\left[\bigvee_{j=0}^{i-1} p_j^{\ell+1}(x_1, \dots, x_{\ell+2}) \right] \leftrightarrow [p_i^{\ell+1}(x_1, \dots, x_{\ell+2}) \leftrightarrow p_i(x_{\ell+2})] \right).$$

From our preliminary remarks on the canonical representations of numbers by bit-strings, we see that $\pi_{1,\ell}(x_1, \dots, x_{\ell+2})$ codes the statement that (for x_1 and $x_{\ell+2}$ 1-integers) the value of $x_{\ell+2}$ is one less than that of $x_1 \pmod{2^n}$ (empty disjunction is defined as false). We then add to $\Phi_{m,n}$ the conjuncts

$$\bigwedge_{\ell=0}^1 \forall x_1 (\text{int}_1(x_1) \rightarrow \forall x_2 \dots \forall x_{\ell+2} (\text{int}_1(x_{\ell+2}) \rightarrow (\text{pred}_{1,\ell}(x_1, \dots, x_{\ell+2}) \leftrightarrow \pi_{1,\ell}(x_1, \dots, x_{\ell+2}))))), \tag{\Phi_4}$$

securing the property 1-predecessor, as required.

If, on the other hand, $m = 1$, we proceed in the same way, except that we add only the conjunct of (Φ_4) with index $\ell = 0$, i.e., the formula

$$\forall x_1 (\text{int}_1(x_1) \rightarrow \forall x_2 (\text{int}_1(x_2) \rightarrow (\text{pred}_{1,0}(x_1, x_2) \leftrightarrow \pi_{1,0}(x_1, x_2)))). \tag{\Phi'_4}$$

This suffices to satisfy the property 1-predecessor without resorting to any predicates outside $\sigma_{1,n}$.

Finally, to secure 1-covering, we add to $\Phi_{m,n}$ the conjuncts

$$\exists x_1(\text{int}_1(x_1) \wedge \text{zero}_1(x_1)); \tag{\Phi_5}$$

$$\forall x_1(\text{int}_1(x_1) \rightarrow \exists x_2(\text{int}_1(x_2) \wedge \text{pred}_{1,0}(x_1, x_2))). \tag{\Phi_6}$$

Observe that (Φ_6) features only $\text{pred}_{1,0}$, and not $\text{pred}_{1,1}$, and so does not stray outside $\sigma_{m,n}$, even when $m = 1$.

Inductive case. This case arises only if $m \geq 2$. Assume that, for some $k < m$, $\text{val}_k : \text{int}_k^{\mathfrak{A}} \rightarrow [0, \mathfrak{t}(k, n) - 1]$ satisfies the properties of k -harmony, k -zero, k -predecessor, k -covering and k -equality. We show, by adding appropriate conjuncts to $\Phi_{m,n}$, that these properties hold with k replaced by $k + 1$.

For $(k + 1)$ -harmony, we add to $\Phi_{m,n}$ the following pair of conjuncts:

$$\begin{aligned} \forall x_1(\text{int}_k(x_1) \rightarrow \\ \forall x_2(\text{int}_{k+1}(x_2) \wedge \pm \text{in}_k(x_1, x_2) \rightarrow \\ \forall x_3(\text{int}_k(x_3) \wedge \text{eq}_{k,1}(x_1, x_2, x_3) \rightarrow \pm \text{out}_k(x_2, x_3))))). \end{aligned} \tag{\Phi_7}$$

If a, a' are k -integers such that $\text{val}_k(a) = \text{val}_k(a')$, and b is any $(k + 1)$ -integer, then, by k -equality (IH), $\mathfrak{A} \models \text{eq}_{k,1}[a, b, a']$, whence (Φ_7) evidently secures $(k + 1)$ -harmony.

We remind ourselves at this point of the role of $(k + 1)$ -harmony in the subsequent argument, and, in particular, on its relationship to k -covering. Let b be a $(k + 1)$ -integer, and recall that $\text{val}_{k+1}(b)$ is defined by b 's satisfaction of the predicates in_k in relation to the various k -integers in \mathfrak{A} . By k -covering (IH), for all i ($0 \leq i < \mathfrak{t}(k, n)$), there is a k -integer a with $\text{val}_k(a) = i$; and by $(k + 1)$ -harmony (just established), all such k -integers a agree on what the i th bit in $\text{val}_{k+1}(b)$ should be.

To secure $(k + 1)$ -zero, we add to $\Phi_{m,n}$ the conjunct

$$\forall x_1(\text{int}_{k+1}(x_1) \rightarrow (\text{zero}_{k+1}(x_1) \leftrightarrow \forall x_2(\text{int}_k(x_2) \rightarrow \neg \text{out}_k(x_1, x_2)))). \tag{\Phi_8}$$

From $(k + 1)$ -harmony and (Φ_8) we see that, for all $(k + 1)$ -integers b , $\mathfrak{A} \models \text{zero}_{k+1}[b] \Leftrightarrow (\text{val}_{k+1}(b) = 0)$. For if there were any k -integer a , such that $\mathfrak{A} \models \text{in}_k[a, b]$, then we would have $\mathfrak{A} \models \text{out}_k[b, a]$.

Establishing the property $(k + 1)$ -predecessor is more involved. We add to $\sigma_{m,n}$ binary predicates $\text{in}_k^{\triangleleft}, \text{out}_k^{\triangleleft}$. The idea is that, for any k -integer a and any $(k + 1)$ -integer b :

$$\mathfrak{A} \models \text{in}_k^{\triangleleft}[a, b] \Leftrightarrow \tag{4}$$

$$(\text{for any } k\text{-integer } a', \text{val}_k(a') < \text{val}_k(a) \Rightarrow \mathfrak{A} \not\models \text{in}_k[a', b]);$$

$$\mathfrak{A} \models \text{in}_k^{\triangleleft}[a, b] \Leftrightarrow \mathfrak{A} \models \text{out}_k^{\triangleleft}[b, a]. \tag{5}$$

Condition (4) allows us to read $\text{in}_k^{\triangleleft}(x_1, x_2)$ as ‘‘all the bits in the value of the $(k + 1)$ -integer x_2 whose index is less than the value of the k -integer x_1 are zero.’’ Condition (5) is somewhat analogous to $(k + 1)$ -harmony.

Securing Condition (5) is easy. We add to $\Phi_{m,n}$ the pair of conjuncts

$$\begin{aligned} \forall x_1(\text{int}_k(x_1) \rightarrow \\ \forall x_2(\text{int}_{k+1}(x_2) \wedge \pm \text{in}_k^{\triangleleft}(x_1, x_2) \rightarrow \\ \forall x_3(\text{int}_k(x_3) \wedge \text{eq}_{k,1}(x_1, x_2, x_3) \rightarrow \pm \text{out}_k^{\triangleleft}(x_2, x_3))))). \end{aligned} \tag{\Phi_9}$$

For let a be a k -integer and b a $(k + 1)$ -integer. By the property k -equality (IH), $\mathfrak{A} \models \text{eq}_{k,1}[a, b, a]$, whence (5) follows.

Securing Condition (4) is harder. We first add to $\sigma_{m,n}$ a binary predicate zero_k^1 , which appends one semantically inert argument to the unary predicate zero_k . That is, we add to $\Phi_{m,n}$ the pair of conjuncts

$$\forall x_1(\text{int}_k(x_1) \wedge \pm \text{zero}_k(x_1) \rightarrow \forall x_2 \pm \text{zero}_k^1(x_1, x_2)). \tag{\Phi_{10}}$$

We can then secure (4) by adding to $\Phi_{m,n}$ the conjunct

$$\begin{aligned} \forall x_1(\text{int}_k(x_1) \rightarrow \\ \forall x_2(\text{int}_{k+1}(x_2) \rightarrow (\text{in}_k^\triangleleft(x_1, x_2) \leftrightarrow (\text{zero}_k^1(x_1, x_2) \vee \\ \forall x_3(\text{int}_k(x_3) \wedge \text{pred}_{k,1}(x_1, x_2, x_3) \rightarrow \\ (\text{out}_k^\triangleleft(x_2, x_3) \wedge \neg \text{out}_k(x_2, x_3))))))). \end{aligned} \tag{\Phi_{11}}$$

To see this, we perform a subsidiary induction on the quantity $\text{val}_k(a)$. Let a be any k -integer and b any $(k + 1)$ -integer. For the base case, suppose $\text{val}_k(a) = 0$. Then $\mathfrak{A} \models \text{zero}_k[a]$ by the property k -zero (IH), whence $\mathfrak{A} \models \text{zero}_k^1[a, b]$ by (Φ_{10}) , whence $\mathfrak{A} \models \text{in}_k^\triangleleft[a, b]$ by (Φ_{11}) . For the inductive step, suppose that $\text{val}_k(a) > 0$; thus, by k -zero again, $\mathfrak{A} \not\models \text{zero}_k[a]$. Assume first that $\mathfrak{A} \models \text{in}_k^\triangleleft[a, b]$. By k -covering (IH), we may pick some k -integer a^* with $\text{val}_k(a^*) = \text{val}_k(a) - 1$. By k -predecessor (IH), setting $\ell = 1$, $\mathfrak{A} \models \text{pred}_{k,1}[a, b, a^*]$, whence, taking x_1, x_2 and x_3 in (Φ_{11}) to be a, b and a^* , respectively, $\mathfrak{A} \models \text{out}_k^\triangleleft[b, a^*]$ and $\mathfrak{A} \not\models \text{out}_k[b, a^*]$. The situation is illustrated in Figure 1. Applying the subsidiary inductive hypothesis, it follows from (4) and (5), with a replaced by a^* , that for any k -integer a' with $\text{val}_k(a') < \text{val}_k(a^*)$, $\mathfrak{A} \not\models \text{in}_k[a', b]$. Moreover, by $(k + 1)$ -harmony (just established), $\mathfrak{A} \not\models \text{out}_k[b, a^*]$ implies that, for any k -integer a' with $\text{val}_k(a') = \text{val}_k(a^*)$, $\mathfrak{A} \not\models \text{in}_k[a', b]$. Thus, for any k -integer a' , $\text{val}_k(a') < \text{val}_k(a) \Rightarrow \mathfrak{A} \not\models \text{in}_k[a', b]$. Conversely, suppose that $\mathfrak{A} \not\models \text{in}_k^\triangleleft[a, b]$. Then, from (Φ_{11}) , there exists some k -integer a^* such that $\text{pred}_{k,1}[a, b, a^*]$, but either $\mathfrak{A} \not\models \text{out}_k^\triangleleft[b, a^*]$ or $\mathfrak{A} \models \text{out}_k[b, a^*]$. By k -predecessor (IH), again setting $\ell = 1$, $\text{val}_k(a^*) = \text{val}_k(a) - 1$, and hence, applying the subsidiary inductive hypothesis, (4) and (5) ensure that, if $\mathfrak{A} \not\models \text{out}_k^\triangleleft[b, a^*]$, then, for some k -integer a' with $\text{val}_k(a') < \text{val}_k(a^*) < \text{val}_k(a)$, $\mathfrak{A} \models \text{in}_k[a', b]$. On the other hand, by k -equality and $(k + 1)$ -harmony, $\mathfrak{A} \models \text{out}_k[b, a^*]$ implies $\mathfrak{A} \models \text{in}_k[a^*, b]$. Either




FIGURE 1. Fixing the interpretation of $\text{in}_k^\triangleleft$.

way, there exists a k -integer a' such that $\text{val}_k(a') < \text{val}_k(a)$, but $\mathfrak{A} \models \text{in}_k[a', b]$. This completes the (subsidiary) induction, and establishes (4).

Having fixed the interpretation of $\text{in}_k^{\triangleleft}$, we proceed to secure the property $(k + 1)$ -predecessor. Assume first that $k + 1 < m$. We add to $\sigma_{m,n}$ the predicates in_k^2 , in_k^3 , $\text{in}_k^{\triangleleft 2}$, $\text{in}_k^{\triangleleft 3}$ and we add to $\Phi_{m,n}$ the conjuncts

$$\bigwedge_{\ell=0}^1 \forall x_1 (\text{int}_k(x_1) \rightarrow \forall x_2 (\text{int}_{k+1}(x_2) \wedge \pm \text{in}_k(x_1, x_2) \rightarrow \forall x_3 \cdots \forall x_{\ell+4} \pm \text{in}_k^{\ell+2}(x_1, \dots, x_{\ell+4}))), \quad (\Phi_{12})$$

$$\bigwedge_{\ell=0}^1 \forall x_1 (\text{int}_k(x_1) \rightarrow \forall x_2 (\text{int}_{k+1}(x_2) \wedge \pm \text{in}_k^{\triangleleft}(x_1, x_2) \rightarrow \forall x_3 \cdots \forall x_{\ell+4} \pm \text{in}_k^{\triangleleft \ell+2}(x_1, \dots, x_{\ell+4}))), \quad (\Phi_{13})$$

fixing these predicates to be the result of adding either 2 or 3 semantically inert arguments to in_k and $\text{in}_k^{\triangleleft}$, as indicated by the superscripts.

Now add to $\sigma_{m,n}$ the ternary predicate $\text{predDig}_{k+1,0}$ and quaternary predicate $\text{predDig}_{k+1,1}$ and let $\varrho_{k,\ell}(x_1, \dots, x_{\ell+4})$ abbreviate the formula

$$\text{in}_k^{\triangleleft \ell+2}(x_1, \dots, x_{\ell+4}) \leftrightarrow (\text{in}_k^{\ell+2}(x_1, \dots, x_{\ell+4}) \leftrightarrow \neg \text{out}_k(x_{\ell+3}, x_{\ell+4})),$$

where $0 \leq \ell \leq 1$. To understand the motivation for this formula, one needs to confine attention to the case where x_1 and $x_{\ell+4}$ are k -integers taking the *same value*, and x_2 and $x_{\ell+3}$ are $(k + 1)$ -integers. In that case, $\varrho_{k,\ell}(x_1, \dots, x_{\ell+4})$ states that the x_1 th digits of x_2 and $x_{\ell+3}$ are opposite if all previous digits of x_2 are zero, and identical otherwise. More formally, suppose that a, a' are k -integers, b, b' $(k + 1)$ -integers and \bar{c} an ℓ -tuple of elements. From (Φ_{12}) , $\mathfrak{A} \models \text{in}_k^{\ell+4}[a, b, \bar{c}, b', a'] \Leftrightarrow \mathfrak{A} \models \text{in}_k[a, b]$, and from (Φ_{13}) , $\mathfrak{A} \models \text{in}_k^{\triangleleft \ell+2}[a, b, \bar{c}, b', a'] \Leftrightarrow \mathfrak{A} \models \text{in}_k^{\triangleleft}[a, b]$. Furthermore, by $(k + 1)$ -harmony, $\mathfrak{A} \models \text{out}_k[b', a'] \Leftrightarrow \mathfrak{A} \models \text{in}_k[a', b']$. Hence, from our preliminary remarks on the canonical representations of numbers by bit-strings, $\mathfrak{A} \models \varrho_{k,\ell}[a, b, \bar{c}, b', a']$ just in case the $\text{val}_k(a')$ th digit in the encoding of $\text{val}_{k+1}(b')$ is the same as the $\text{val}_k(a)$ th digit in the encoding of $\text{val}_{k+1}(b) - 1$, modulo $t(k + 1, n)$.

We then add to $\Phi_{m,n}$ the conjunct

$$\bigwedge_{\ell=0}^1 \forall x_1 (\text{int}_k(x_1) \rightarrow \forall x_2 (\text{int}_{k+1}(x_2) \rightarrow \forall x_3 \cdots \forall x_{\ell+3} (\text{int}_{k+1}(x_{\ell+3}) \rightarrow \forall x_{\ell+4} (\text{int}_k(x_{\ell+4}) \wedge \text{eq}_{k,\ell+2}(x_1, \dots, x_{\ell+4}) \rightarrow (\text{predDig}_{k+1,\ell}(x_2, \dots, x_{\ell+4}) \leftrightarrow \varrho_{k,\ell}(x_1, \dots, x_{\ell+4})))))). \quad (\Phi_{14})$$

(Remember: we are assuming that $k < m - 1$, so that $\text{eq}_{k,3}$ exists in $\sigma_{m,n}$.) This formula is illustrated in the diagram of Figure 2a in the case $\ell = 1$: here, $\varrho_{k,1}$ holds of the tuple a, b, c, b', a , and $\text{predDig}_{k+1,1}$ of the tuple b, c, b', a , just in case the $\text{val}_k(a)$ th digit of $\text{val}_{k+1}(b')$ agrees with the $\text{val}_k(a)$ th digit of $\text{val}_{k+1}(b) - 1$. (Note that the single element a is depicted twice in this diagram.) Suppose a is a

k -integer, b, b' are $(k + 1)$ -integers in \mathfrak{A} , and \bar{c} is any ℓ -tuple from A with $0 \leq \ell \leq 1$. By k -equality (IH), $\mathfrak{A} \models \text{eq}_{k,\ell+2}[a, b, \bar{c}, b', a]$, and from the properties of $\varrho_{k,\ell}$ just established (setting $a' = a$), $\mathfrak{A} \models \text{predDig}_{k+1,\ell}[b, \bar{c}, b', a]$ just in case the $\text{val}_k(a)$ th digit of $\text{val}_{k+1}(b')$ is equal to the $\text{val}_k(a)$ th digit of $\text{val}_{k+1}(b) - 1$, modulo $t(k + 1, n)$.

To establish $(k + 1)$ -predecessor, therefore, we add to $\Phi_{m,n}$ the conjuncts

$$\bigwedge_{\ell=0}^1 \forall x_1 (\text{int}_{k+1}(x_1) \rightarrow \forall x_2 \cdots \forall x_{\ell+2} (\text{int}_{k+1}(x_{\ell+2}) \rightarrow (\text{pred}_{k+1,\ell}(x_1, \dots, x_{\ell+2}) \leftrightarrow \forall x_{\ell+3} (\text{int}_k(x_{\ell+3}) \rightarrow \text{predDig}_{k+1,\ell}(x_1, \dots, x_{\ell+3})))))). \quad (\Phi_{15})$$

From (Φ_{15}) , $\mathfrak{A} \models \text{pred}_{k+1,\ell}[b, \bar{c}, b']$ just in case each digit of $\text{val}_{k+1}(b')$ is equal to the corresponding digit of $\text{val}_{k+1}(b) - 1$, modulo $t(k + 1, n)$.

If, on the other hand, $k + 1 = m$, we proceed as above, but we add to $\sigma_{m,n}$ only the predicates in_k^2 , $\text{in}_k^{<2}$, $\text{predDig}_{k+1,0}$ (not in_k^3 , $\text{in}_k^{<3}$ or $\text{predDig}_{k+1,1}$), and we add to $\Phi_{m,n}$ only those conjuncts of (Φ_{12}) – (Φ_{15}) with $\ell = 0$ (not with $\ell = 1$). This suffices for $(k + 1)$ -predecessor in the case $k + 1 = m$, and does not require the use of any predicates outside $\sigma_{m,n}$. It is important to be careful with indices here: the conjunct $\ell = 1$ in (Φ_{12}) – (Φ_{14}) requires predicates of arity 5, which are not available in \mathcal{FL}^4 , as required by the theorem when $m = 2$.

To establish the property $(k + 1)$ -covering, we add to $\Phi_{m,n}$ the conjuncts

$$\exists x_1 (\text{int}_{k+1}(x_1) \wedge \text{zero}_{k+1}(x_1)); \quad (\Phi_{16})$$

$$\forall x_1 (\text{int}_{k+1}(x_1) \rightarrow \exists x_2 (\text{int}_{k+1}(x_2) \wedge \text{pred}_{k+1,0}(x_1, x_2))). \quad (\Phi_{17})$$

Note that (Φ_{17}) features only $\text{pred}_{k+1,0}$, and not $\text{pred}_{k+1,1}$, so it is defined even when $k + 1 = m$

It remains only to establish $(k + 1)$ -equality. Conceptually, this is rather easier than $(k + 1)$ -predecessor; however, we do need to consider larger numbers of semantically inert variables. Let $L = 2(m - k - 1)$. The property $(k + 1)$ -equality concerns the interpretation of the $(\ell + 2)$ -ary predicate $\text{eq}_{k+1,\ell}$ for all ℓ ($0 \leq \ell \leq L$). Observe that, if $k = 1$ (first inductive step), then $L = 2m - 4$, and if $k = m - 1$ (last inductive step), then $L = 0$. Thus, in the sequel, we always have $L \leq 2m - 4$. (Remember that the inductive case is encountered only if $m \geq 2$.) To ease the pain of reading, we split the task into three stages.

For the first stage, for all ℓ ($0 \leq \ell \leq L$), add to $\sigma_{m,n}$ an $(\ell + 2)$ -ary predicate in_k^ℓ , and add to $\Phi_{m,n}$ the conjuncts

$$\bigwedge_{\ell=0}^L \forall x_1 (\text{int}_k(x_1) \rightarrow \forall x_2 (\text{int}_{k+1}(x_2) \wedge \pm \text{in}_k(x_1, x_2) \rightarrow \forall x_3 \cdots \forall x_{\ell+2} (\pm \text{in}_k^\ell(x_1, x_2, \dots, x_{\ell+1}, x_{\ell+2}))))), \quad (\Phi_{18})$$

thus fixing in_k^ℓ to be the result of adding ℓ semantically inert arguments to in_k . (For $\ell \leq 3$, this repeats the work of (Φ_{12}) , but no matter.)

In the second stage, for all ℓ ($0 \leq \ell \leq L$), add to $\sigma_{m,n}$ an $(\ell + 3)$ -ary predicate $\text{eqDig}_{k+1,\ell}$, and add to $\Phi_{m,n}$ the conjuncts

$$\bigwedge_{\ell=0}^L \forall x_1 (\text{int}_k(x_1) \rightarrow \forall x_2 (\text{int}_{k+1}(x_2) \rightarrow \forall x_3 \cdots \forall x_{\ell+3} (\text{int}_{k+1}(x_{\ell+3}) \rightarrow \forall x_{\ell+4} (\text{int}_k(x_{\ell+4}) \wedge \text{eq}_{k,\ell+2}(x_1, \dots, x_{\ell+4}) \rightarrow (\text{eqDig}_{k+1,\ell}(x_2, \dots, x_{\ell+4}) \leftrightarrow \eta_{k,\ell+2}(x_1, \dots, x_{\ell+4}))))))))), \quad (\Phi_{19})$$

where $\eta_{k,\ell+2}(x_1, \dots, x_{\ell+4})$ is the formula: $\text{in}_k^{\ell+2}(x_1, \dots, x_{\ell+4}) \leftrightarrow \text{out}_k(x_{\ell+3}, x_{\ell+4})$.
 Let b, b' be $(k + 1)$ -integers in \mathfrak{A} , a a k -integer in \mathfrak{A} , and \bar{c} any ℓ -tuple from A . We claim that $\mathfrak{A} \models \text{eqDig}_{k+1,\ell}[b, \bar{c}, b', a]$ just in case $\text{val}_{k+1}(b)$ and $\text{val}_{k+1}(b')$ agree on their $\text{val}_k(a)$ th bit. For, by k -equality (IH), $\mathfrak{A} \models \text{eq}_{k,\ell+2}[a, b, \bar{c}, b', a]$. Hence, by (Φ_{19}) $\mathfrak{A} \models \text{eqDig}_{k+1,\ell}[b, \bar{c}, b', a]$ holds just in case $\mathfrak{A} \models \eta_{k,\ell+2}[a, b, \bar{c}, b', a]$. But, by (Φ_{18}) , $\mathfrak{A} \models \text{in}_k^{\ell+2}[a, b, \bar{c}, b', a]$ if and only if $\mathfrak{A} \models \text{in}_k[a, b]$, i.e., if and only if the $\text{val}_k(a)$ th bit of $\text{val}_{k+1}(b)$ is 1. That is, $\mathfrak{A} \models \text{eqDig}_{k+1,\ell}[b, \bar{c}, b', a]$ is equivalent to the statement that $\mathfrak{A} \models \text{in}_k[a, b]$ if and only if $\mathfrak{A} \models \text{out}_k[b', a]$. But, by $(k + 1)$ -harmony, $\mathfrak{A} \models \text{out}_k[b', a]$ if and only if $\mathfrak{A} \models \text{in}_k[a, b']$. This establishes the claim. The situation is illustrated (for the case where $\mathfrak{A} \models \text{eqDig}_{k+1,\ell}[b, \bar{c}, b', a]$ holds) in Figure 2b, where all polarity alternatives \pm are assumed to be resolved in the same way.

In the third stage, we add to $\Phi_{m,n}$ the conjunct

$$\bigwedge_{\ell=0}^L \forall x_1 (\text{int}_{k+1}(x_1) \rightarrow \forall x_2 \cdots \forall x_{\ell+2} (\text{int}_{k+1}(x_{\ell+2}) \rightarrow (\text{eq}_{k+1,\ell}(x_1, \dots, x_{\ell+2}) \leftrightarrow \forall x_{\ell+3} (\text{int}_k(x_{\ell+3}) \rightarrow \text{eqDig}_{k+1,\ell}(x_1, \dots, x_{\ell+3})))))). \quad (\Phi_{20})$$

Given the properties of $\text{eqDig}_{k+1,\ell}$ just established, this evidently secures $(k + 1)$ -equality, completing the induction.

We have remarked that, by m -covering, any model of $\Phi_{m,n}$ has cardinality at least $t(m, n)$. We claim that $\Phi_{m,n}$ is satisfiable. Let $A = A_1 \dot{\cup} \cdots \dot{\cup} A_m$, where $A_k = \{\langle k, i \rangle \mid 0 \leq i < t(k, n)\}$. (That is, A is the disjoint union of the various sets of integers $[0, t(k, n) - 1]$.) Let $\text{int}_k^{\mathfrak{A}} = A_k$ for all k ($1 \leq k \leq m$), and interpret the other predicates of $\sigma_{m,n}$ as described above. It is easily verified that $\mathfrak{A} \models \Phi_{m,n}$.

It remains only to check the number of variables featured in $\Phi_{m,n}$. Consider first the conjuncts introduced in the base case. By inspection, (Φ_1) – (Φ_3) and (Φ_5) – (Φ_6) are in \mathcal{FL}^{2m} . For $m > 1$, (Φ_4) is in \mathcal{FL}^3 ; but if $m = 1$, only the conjunct with index




FIGURE 2. Fixing the interpretations of predicates to ensure $(k + 1)$ -predecessor and $(k + 1)$ -covering.

$\ell = 0$ is present, which is in \mathcal{FL}^2 . Either way, (Φ_1) – (Φ_6) are in \mathcal{FL}^{2m} . Consider now the conjuncts introduced in the inductive case. By inspection, these feature only $\max(5, 2m)$ variables. If $m > 2$, $\max(5, 2m) = 2m$. If, on the other hand, $m = 2$, then the inductive step only runs once, with $k + 1 = m$, in which case only those conjuncts of (Φ_{12}) – (Φ_{15}) occur for which $\ell = 0$, and these feature only 4 variables. Either way, (Φ_7) – (Φ_{20}) are in \mathcal{FL}^{2m} . \dashv

In the following theorem, we assume polynomial-time reductions for simplicity.

THEOREM 3.2. *Let $m \geq 1$. The satisfiability problem for \mathcal{FL}^{2m} is m -NEXPTIME-hard. Hence, the satisfiability problem for \mathcal{FL} is nonelementary.*

PROOF. For $m = 1$, we use the fact that Boolean modal logic—which, as we mentioned earlier, can be embedded in \mathcal{FL}^2 —is NEXPTIME-complete [7]. Thus we henceforth assume that $m \geq 2$. We employ the well-known technique of reduction from tiling problems. A *tiling system* is a triple (C, H, V) , where C is a nonempty, finite set, and H, V are subsets of $C \times C$. For any $N \geq 2$, an $N \times N$ *tiling* for (C, V, H) is a function $f : [0, N - 1] \times [0, N - 1] \rightarrow C$ such that, for all i, j ($0 \leq i, j < N$), $\langle f(i, j), f(i + 1, j) \rangle \in H$ and $\langle f(i, j), f(i, j + 1) \rangle \in V$, where arithmetic in arguments is interpreted modulo N . Intuitively, we are to imagine an $N \times N$ grid (with toroidal wrap-around) whose squares are covered with tiles having the colours in C : the relation H specifies which colours are allowed to go immediately ‘to the right of’ which other colours, and the relation V specifies which colours are allowed to go immediately ‘above’ which other colours. Let $w \in C^*$ with $|w| = n$ (i.e., w is a word of length n over the alphabet C). We say that the tiling f *satisfies the initial condition* w if $w = f(0, 0)f(1, 0) \cdots f(n - 1, 0)$. Intuitively, we are to imagine w written by f in the ‘bottom left’ corner of the grid. For any nondeterministic Turing machine M over an alphabet Σ , running in time g , there exists a tiling system (C, H, V) and a linear-time computable, length-preserving transduction $\tau : \Sigma^* \rightarrow C^*$, such that, for any $x \in \Sigma^*$, M has an accepting run on input x if and only if there is a $g(|w|) \times g(|w|)$ -tiling for (C, H, V) satisfying initial condition $w = \tau(x)$ (see, e.g., Börger, Grädel, and Gurevich [2, pp. 243, Theorem 6.1.2]). Intuitively, the tiling in question is a ‘movie’ of the run of M , with the tiles in the i th row representing the first $g(|w|)$ squares of M ’s tape at time i .

Let $m \geq 2$ be fixed. We wish to reduce an arbitrary problem P in $\text{NTIME}(t(m, p(n)))$, for p a polynomial, to the satisfiability problem for \mathcal{FL}^{2m} . Using a standard padding argument, we may assume that P is in fact in $\text{NTIME}(t(m, n))$. It suffices to show that, if (C, H, V) is a tiling system, then we can map, in polynomial time, any word $w \in C^*$ to an \mathcal{FL}^{2m} -formula Ψ_w such that Ψ_w is satisfiable if and only if there exists a $t(m, n) \times t(m, n)$ tiling for (C, H, V) satisfying initial condition w , where $|w| = n$. We shall proceed by constructing an \mathcal{FL}^{2m} -formula Ψ_n —depending only on the *length* of w —any of whose models defines a (C, H, V) -tiling of the $t(m, n) \times t(m, n)$ -grid. It will then be a simple matter to define Ψ_w .

To construct Ψ_n , we first carry out the construction of $\Phi_{m,n}$ in the proof of Theorem 3.1, and start by setting Ψ_n to be $\Phi_{m,n}$. (Remember, m is fixed here.) The various conjuncts of this formula thus establish the existence of k -integers, with a valuation function val_k satisfying the properties k -harmony, k -zero, k -predecessor,

k -covering, k -equality, for all values of k from 1 to m . In the sequel, we shall concern ourselves mainly with $(m - 1)$ -integers. In fact, we do not require m -integers directly, but instead employ objects we refer to as *vertices*, which are, in effect, pairs of m -integers. We establish the requisite properties of vertices by adding to Ψ_n further conjuncts as described below.

Let vtx be a unary predicate, and $\text{in}_X, \text{in}_Y, \text{out}_X, \text{out}_Y$, binary predicates. If a structure \mathfrak{A} is clear from context, we call any element of A satisfying vtx a *vertex*. Define the function $\text{val}_X : \text{vtx}^{\mathfrak{A}} \rightarrow [0, t(m, n) - 1]$ by setting $\text{val}_X(b)$, for any vertex b , to be the integer coded by the bit-string s_{N-1}, \dots, s_0 of length $N = t(m - 1, n)$ where, for all i ($0 \leq i < N$),

$$s_i = \begin{cases} 1 & \text{if } \mathfrak{A} \models \text{in}_X[a, b] \text{ for some } (m - 1)\text{-integer } a \text{ such that } \text{val}_{m-1}(a) = i; \\ 0 & \text{otherwise.} \end{cases}$$

Thus, we use the $(m - 1)$ -integers to represent positions in the horizontal coordinate of any vertex, employing the binary predicate in_X to encode the bits at the positions in question. Similarly, we may use the $(m - 1)$ -integers to represent positions in the vertical coordinate of any vertex, employing the binary predicate in_Y to encode the bits at the positions in question. Here again, we implicitly rely on $(m - 1)$ -covering to show that there *is* an $(m - 1)$ -integer having any value i in the range $[0, t(m - 1, n) - 1]$. In the sequel, we establish harmony-like properties showing that it does not matter *which* such $(m - 1)$ -integer we choose.

We write $\bar{X} = Y$ and $\bar{Y} = X$, and we use the symbol D to stand for either of the letters X or Y . We shall ensure that any model of Ψ_n has the following properties, for $D \in \{X, Y\}$.

- D-harmony*: For all vertices b and all $(m - 1)$ -integers a, a' in \mathfrak{A} such that $\text{val}_{m-1}(a) = \text{val}_{m-1}(a')$, $\mathfrak{A} \models \text{in}_D[a, b] \Leftrightarrow \mathfrak{A} \models \text{in}_D[a', b]$.
- D-zero*: For all vertices b in \mathfrak{A} , $\mathfrak{A} \models \text{zero}_D[b] \Leftrightarrow \text{val}_D(b) = 0$.
- D-equality*: For all vertices b, b' in \mathfrak{A} , $\mathfrak{A} \models \text{eq}_D[b, b'] \Leftrightarrow \text{val}_D(b) = \text{val}_D(b')$.
- D-predecessor*: For all vertices b, b' in \mathfrak{A} , $\mathfrak{A} \models \text{pred}_D[b, b'] \Leftrightarrow \text{val}_D(b') = \text{val}_D(b) - 1 \text{ modulo } t(m, n)$.

To do so, we add to Ψ_n the following conjuncts. The argumentation is in each case virtually identical to that given in the proof of Theorem 3.1. For X - and Y -harmony, we add to our signature binary predicates $\text{out}_X, \text{out}_Y$, and add to Ψ_n the conjuncts

$$\bigwedge_{D \in \{X, Y\}} \forall x_1 (\text{int}_{m-1}(x_1) \rightarrow \forall x_2 (\text{vtx}(x_2) \wedge \pm \text{in}_D(x_1, x_2) \rightarrow \forall x_3 (\text{int}_{m-1}(x_3) \rightarrow \text{eq}_{m-1,1}(x_1, x_2, x_3) \rightarrow \pm \text{out}_D(x_2, x_3))))).$$

We observe in passing that, if b is a vertex and c, c' are $(m - 1)$ -integers with $\text{val}_k(c) = \text{val}_k(c')$, then $\mathfrak{A} \models \text{out}_D[b, c] \Leftrightarrow \mathfrak{A} \models \text{out}_D[b, c']$. At this point, we have established that, in any model \mathfrak{A} of Ψ_n , any vertex b has well-defined coordinates $(\text{val}_X(b), \text{val}_Y(b))$.

For X - and Y -zero, add to Ψ_n the conjuncts

$$\bigwedge_{D \in \{X, Y\}} \forall x_1 (\text{vtx}(x_1) \rightarrow (\text{zero}_D(x_1) \leftrightarrow \forall x_2 (\text{int}_{m-1}(x_2) \rightarrow \neg \text{out}_D(x_1, x_2))))).$$

To secure the remaining properties, we need more predicates. We first introduce the binary predicates $\text{in}_D^\triangleleft$, $\text{out}_D^\triangleleft$. The idea is that, for any $(m - 1)$ -integer a and any vertex b , we should have

$$\mathfrak{A} \models \text{in}_D^\triangleleft[a, b] \Leftrightarrow (\text{for any } (m - 1)\text{-integer } a', \tag{6}$$

$$\text{val}_{m-1}(a') < \text{val}_{m-1}(a) \Rightarrow \mathfrak{A} \not\models \text{in}_D[a', b]);$$

$$\mathfrak{A} \models \text{in}_D^\triangleleft[a, b] \Leftrightarrow \mathfrak{A} \models \text{out}_D^\triangleleft[b, a]. \tag{7}$$

To secure these conditions, we proceed exactly as for (4) and (5), adding to Ψ_n the conjuncts

$$\bigwedge_{D \in \{X, Y\}} \forall x_1 (\text{int}_{m-1}(x_1) \rightarrow$$

$$\forall x_2 (\text{vtx}(x_2) \wedge \pm \text{in}_D^\triangleleft(x_1, x_2) \rightarrow$$

$$\forall x_3 (\text{int}_{m-1}(x_3) \wedge \text{eq}_{m-1,1}(x_1, x_2, x_3) \rightarrow \pm \text{out}_D^\triangleleft(x_2, x_3)))));$$

$$\bigwedge_{D \in \{X, Y\}} \forall x_1 (\text{int}_{m-1}(x_1) \rightarrow$$

$$\forall x_2 (\text{vtx}(x_2) \rightarrow (\text{in}_D^\triangleleft(x_1, x_2) \leftrightarrow (\text{zero}_{m-1}^1(x_1, x_2) \vee$$

$$\forall x_3 (\text{int}_{m-1}(x_3) \wedge \text{pred}_{m-1,1}(x_1, x_2, x_3) \rightarrow$$

$$(\text{out}_D^\triangleleft(x_2, x_3) \wedge \neg \text{out}_{m-1}(x_2, x_3))))))))).$$

(The interpretation of zero_{m-1}^1 was fixed during the construction of $\Phi_{m,n}$.) That the truth of these formulas in a structure \mathfrak{A} ensures (6) and (7) follows by reasoning analogous to the earlier argument that formulas (Φ_9) – (Φ_{11}) secure (4) and (5).

Having fixed the interpretation of $\text{in}_D^\triangleleft$, we introduce the predicates in_D^2 and $\text{in}_D^{\triangleleft 2}$, and we add to Ψ_n conjuncts fixing these predicates to be the result of adding 2 semantically inert arguments to in_D and $\text{in}_D^\triangleleft$, as indicated by the superscripts. The following evidently suffice:

$$\forall x_1 (\text{int}_{m-1}(x_1) \rightarrow \forall x_2 (\text{vtx}(x_2) \wedge \pm \text{in}_D(x_1, x_2) \rightarrow \forall x_3 \forall x_4 \pm \text{in}_D^2(x_1, \dots, x_4)));$$

$$\forall x_1 (\text{int}_{m-1}(x_1) \rightarrow \forall x_2 (\text{vtx}(x_2) \wedge \pm \text{in}_D^\triangleleft(x_1, x_2) \rightarrow \forall x_3 \forall x_4 \pm \text{in}_D^{\triangleleft 2}(x_1, \dots, x_4))).$$

We may now secure X - and Y -equality by adding to Ψ_n the conjuncts

$$\bigwedge_{D \in \{X, Y\}} \forall x_1 (\text{int}_{m-1}(x_1) \rightarrow$$

$$\forall x_2 (\text{vtx}(x_2) \rightarrow$$

$$\forall x_3 (\text{vtx}(x_3) \rightarrow$$

$$\forall x_4 (\text{int}_{m-1}(x_4) \wedge \text{eq}_{m-1,2}(x_1, x_2, x_3, x_4) \rightarrow$$

$$(\text{eqDig}_D(x_2, x_3, x_4) \leftrightarrow$$

$$(\text{in}_D^2(x_1, x_2, x_3, x_4) \leftrightarrow \text{out}_D(x_3, x_4))))))))),$$

$$\bigwedge_{D \in \{X, Y\}} \forall x_1 (\text{vtx}(x_1) \rightarrow$$

$$\forall x_2 (\text{vtx}(x_2) \rightarrow$$

$$(\text{eq}_D(x_1, x_2) \leftrightarrow$$

$$\forall x_3 (\text{int}_{m-1}(x_3) \rightarrow \text{eqDig}_D(x_1, x_2, x_3))))),$$

analogous to (Φ_{19}) and (Φ_{20}) . Notice that $\text{eq}_D(x_1, x_2)$ has no semantically inert variables; we no longer need them now that the recursive construction is over.

Turning finally to X - and Y -predecessor, let $\varrho_D(x_1, \dots, x_4)$ be the formula

$$\text{in}_D^{\leq 2}(x_1, \dots, x_4) \leftrightarrow (\text{in}_D^2(x_1, \dots, x_4) \leftrightarrow \neg \text{out}_D(x_3, x_4))$$

for $D \in \{X, Y\}$. Using reasoning entirely analogous to that for (Φ_{14}) and (Φ_{15}) , we see that adding to Ψ_n the conjuncts

$$\bigwedge_{D \in \{X, Y\}} \forall x_1(\text{int}_{m-1}(x_1) \rightarrow \forall x_2(\text{vtx}(x_2) \rightarrow \forall x_3(\text{vtx}(x_3) \rightarrow \forall x_4(\text{int}_{m-1}(x_4) \wedge \text{eq}_{m-1,2}(x_1, \dots, x_4) \rightarrow (\text{predDig}_D(x_2, x_3, x_4) \leftrightarrow \varrho_D(x_1, \dots, x_4))))))$$

$$\bigwedge_{D \in \{X, Y\}} \forall x_1 \forall x_2 (\text{pred}_D(x_1, x_2) \leftrightarrow \forall x_3 (\text{int}_{m-1}(x_3) \rightarrow \text{predDig}_D(x_1, x_2, x_3)))$$

secures D -predecessor. At this point, we have established that, in any model \mathfrak{A} of Ψ_n , and for $D \in \{X, Y\}$, the properties D -harmony, D -zero, D -predecessor and D -equality obtain.

The following conjuncts of Ψ_n now establish that, for all pairs of integers i, j in the range $[0, t(m, n) - 1]$, there exists a vertex with coordinates (x, y) :

$$\begin{aligned} &\exists x_1(\text{vtx}(x_1) \wedge \text{zero}_X(x_1) \wedge \text{zero}_Y(x_1)); \\ &\forall x_1(\text{vtx}(x_1) \rightarrow \exists x_2(\text{vtx}(x_2) \wedge \text{pred}_Y(x_1, x_2) \wedge \text{eq}_X(x_1, x_2))); \\ &\forall x_1(\text{vtx}(x_1) \rightarrow \exists x_2(\text{vtx}(x_2) \wedge \text{pred}_X(x_1, x_2) \wedge \text{eq}_Y(x_1, x_2))). \end{aligned}$$

Treating the colours in C as unary predicates, the following conjuncts assign, to each vertex, a , a unique colour, namely, the colour which a satisfies:

$$\begin{aligned} &\forall x_1(\text{vtx}(x_1) \rightarrow \bigvee_{c \in C} c(x_1)); \\ &\bigwedge_{\substack{c \neq d \\ c, d \in C}} \forall x_1(\text{vtx}(x_1) \rightarrow \neg(c(x_1) \wedge d(x_1))). \end{aligned}$$

Note that there is no requirement that vertices be uniquely defined by their X - and Y -coordinates. Nevertheless, we obtain a well-defined encoding of a grid-colouring by securing the property

chromatic harmony: For all vertices b and b' such that $\text{val}_X(b) = \text{val}_X(b')$ and $\text{val}_Y(b) = \text{val}_Y(b')$, and for all colours $c \in C$, $\mathfrak{A} \models c[b] \Leftrightarrow \mathfrak{A} \models c[b']$.

To do so, we add to Ψ_n the conjunct

$$\bigwedge_{c \in C} \forall x_1(\text{vtx}(x_1) \wedge c(x_1) \rightarrow \forall x_2(\text{vtx}(x_2) \wedge \text{eq}_X(x_1, x_2) \wedge \text{eq}_Y(x_1, x_2) \rightarrow c(x_2))).$$

Thus, any model \mathfrak{A} of Ψ_n defines a function $f : [0, t(m, n) - 1]^2 \rightarrow C$. To ensure that f is a tiling for the system (C, H, V) , we simply add to Ψ_n the conjuncts

$$\bigwedge_{(c,d) \notin H} \forall x_1 (\text{vtx}(x_1) \wedge d(x_1) \rightarrow \forall x_2 (\text{vtx}(x_2) \wedge \text{pred}_X(x_1, x_2) \wedge \text{eq}_Y(x_1, x_2) \rightarrow \neg c(x_2)));$$

$$\bigwedge_{(c,d) \notin V} \forall x_1 (\text{vtx}(x_1) \wedge d(x_1) \rightarrow \forall x_2 (\text{vtx}(x_2) \wedge \text{pred}_Y(x_1, x_2) \wedge \text{eq}_X(x_1, x_2) \rightarrow \neg c(x_2))).$$

This completes the construction of Ψ_n . We have shown that any model of Ψ_n defines a (C, H, V) -tiling of the $t(m, n) \times t(m, n)$ -grid. To obtain the formula Ψ_w , we simply need to state that w is written in the bottom left corner of the grid. Let $w = c_0 \dots c_{n-1}$, and let q_0, \dots, q_{n-1} be new unary predicates. (Notice that the various c_i may contain repeats, but the q_i are by assumption all different.) We define Ψ_w to be the conjunction of Ψ_n and

$$\exists x_1. q_{n-1}(x_1) \wedge \forall x_1 (q_0(x_1) \rightarrow \text{zero}_X(x_1));$$

$$\bigwedge_{i=0}^{n-1} \forall x_1 (q_i(x_1) \rightarrow \text{vtx}(x_1) \wedge \text{zero}_Y(x_1) \wedge c_i(x_1));$$

$$\bigwedge_{i=0}^{n-2} \forall x_1 (q_{i+1}(x_1) \rightarrow \exists x_2 (q_i(x_2) \wedge \text{pred}_X(x_1, x_2))).$$

Thus, if Ψ_w is satisfiable, then there is a (C, H, V) -tiling of the $t(m, n) \times t(m, n)$ -grid with initial condition w . Conversely, a simple check shows that, if there is a (C, H, V) -tiling of the $t(m, n) \times t(m, n)$ -grid with initial condition w , then by interpreting the predicates involved in Ψ_w as suggested above over a two-dimensional toroidal grid of m -integers, we obtain a model of Ψ_w . This completes the reduction. ⊣

§4. Upper bound. In this section we first show that \mathcal{FL}^3 has the exponential-sized model property, and then prove a bounded model property for every \mathcal{FL}^m with $m > 3$. These results give the upper complexity bounds for the satisfiability problem for every \mathcal{FL}^m with $m \geq 3$.

It will be convenient to work with fluted formulas having a special form. Fix some purely relational signature σ and some positive integer k . A *fluted k -clause* is a disjunction of fluted k -literals. We allow the absurd formula \perp (i.e., the empty disjunction) to count as a fluted k -clause. Thus, any literal of a fluted k -clause is either 0-ary (i.e., a proposition letter or its negation) or has arguments x_h, \dots, x_k , in that order, for some h ($1 \leq h \leq k$). When writing fluted k -clauses, we silently remove bracketing, reorder literals and delete duplicated literals as necessary. It is easy to see that the number of fluted k -clauses, modulo these operations, is precisely $2^{2^{|\sigma|}}$. To reduce notational clutter, we identify finite sets of fluted k -clauses with their conjunctions where convenient, writing Γ instead of the more correct $\bigwedge \Gamma$.

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

$$\forall x_1 \cdots \forall x_m. \Omega \wedge \bigwedge_{i=1}^s \forall x_1 \cdots \forall x_{m-1} (\alpha_i \rightarrow \exists x_m. \Gamma_i) \wedge \bigwedge_{j=1}^t \forall x_1 \cdots \forall x_{m-1} (\beta_j \rightarrow \forall x_m. \Delta_j), \quad (8)$$

where $\Omega, \Gamma_1, \dots, \Gamma_s, \Delta_1, \dots, \Delta_t$ are finite sets of fluted m -clauses and $\alpha_1, \dots, \alpha_s, \beta_1, \dots, \beta_t$ fluted $(m - 1)$ -atoms. (We allow empty conjunctions, interpreted as \top .) We refer to $\forall x_1 \cdots \forall x_m. \Omega$ as the *static conjunct* of φ , to conjuncts of the form $\forall x_1 \cdots \forall x_{m-1} (\alpha_i \rightarrow \exists x_m. \Gamma_i)$ as the *existential conjuncts* of φ , and to conjuncts of the form $\forall x_1 \cdots \forall x_{m-1} (\beta_j \rightarrow \forall x_m. \Delta_j)$ as the *universal conjuncts* of φ . Normal-form \mathcal{FL}^m -formulas are by definition sentences, i.e., have no free variables.

LEMMA 4.1. *Let φ be an \mathcal{FL}^m -sentence, $m \geq 1$. We can compute, in time bounded by a polynomial function of $\|\varphi\|$, an \mathcal{FL}^m -formula ψ in normal form such that: (i) $\models \psi \rightarrow \varphi$; and (ii) any model of φ can be expanded to a model of ψ .*

PROOF. We take Q to stand for either of the quantifiers \forall or \exists , writing $\overline{\forall} := \exists$ and $\overline{\exists} := \forall$. If α is a fluted $(k + 1)$ -atom ($0 \leq k < m$), write $\overline{\alpha} := \neg\alpha$ and $\overline{\overline{\alpha}} := \alpha$; thus, if μ is a fluted $(k + 1)$ -literal, so is $\overline{\mu}$. If θ is a formula of the form $Qx_{k+1}. \mu$, where μ is a fluted $(k + 1)$ -literal, write $\overline{\theta} := \overline{Q}x_{k+1}. \overline{\mu}$; thus, $\overline{\theta}$ is (also of this form and) logically equivalent to $\neg\theta$. We prove the lemma using standard ‘re-writing’ techniques (see, e.g., [10]). By moving negations inward in the usual way, we may assume without loss of generality that negation symbols in φ apply only to atoms. Let $\varphi_0 := \varphi$.

Suppose the \mathcal{FL}^m -sentence φ_0 has a subformula $\theta = \mu \circ v$, where μ, v are fluted k -literals ($0 \leq k \leq m$) and \circ is a Boolean operator. Let p be a fresh predicate of arity k , let $\varphi_1 := \varphi_0[\theta/p(x_1, \dots, x_k)]$ (i.e., φ_0 with θ replaced by the new fluted atom), and let $\psi_1 := \forall x_1 \cdots \forall x_k (p(x_1, \dots, x_k) \leftrightarrow \theta)$. It is obvious that $\models \varphi_1 \wedge \psi_1 \rightarrow \varphi_0$, and any model of φ_0 can be expanded to a model of $\varphi_1 \wedge \psi_1$. Moreover ψ_1 can be equivalently written in the form $\forall x_1 \cdots \forall x_k. \Omega$, where Ω is a set of fluted k -clauses, with at most a linear increase in size.

Suppose, alternatively, φ_0 has a subformula $\theta = Qx_{k+1}. \mu$, where μ is a fluted $(k + 1)$ -literal ($0 \leq k < m$). Let p and p' be fresh predicates of arity k , let $\varphi_1 := \varphi_0[\theta/p(x_1, \dots, x_k)]$, and let ψ_1 be the formula

$$\forall x_1 \cdots \forall x_{k+1} (p(x_2, \dots, x_{k+1}) \vee p'(x_2, \dots, x_{k+1})) \wedge \forall x_1 \cdots \forall x_k (p(x_1, \dots, x_k) \rightarrow \theta) \wedge \forall x_1 \cdots \forall x_k (p'(x_1, \dots, x_k) \rightarrow \overline{\theta}).$$

Hence, $\models \varphi_1 \wedge \psi_1 \rightarrow \varphi_0$, and any model of φ_0 can be expanded to a model of $\varphi_1 \wedge \psi_1$. We see that ψ_1 is a normal-form formula of \mathcal{FL}^{k+1} with $s = t = 1$, and Ω, Γ_1 and Δ_1 singletons. (In fact the clauses in Γ_1 and Δ_1 consist of just one $(k + 1)$ -literal.)

Proceeding similarly with φ_1 in place of φ_0 , we obtain φ_2 and ψ_2 , and so on, until we reach some point where the \mathcal{FL}^m -sentence φ_s is a fluted 0-literal, and hence a proposition letter or a negated proposition letter, at which point φ_s may be equivalently written $\forall x_1. \varphi_s$. Defining ψ' to be $\psi_1 \wedge \cdots \wedge \psi_s \wedge \varphi_s$, we see that $\models \psi' \rightarrow \varphi$, and any model of φ can be expanded to a model of ψ' . It should be clear

that the size of ψ' is at most linear in the size of φ , and that the above computation runs in polynomial time. Each of the ψ_i and indeed φ_s is a normal-form formula of \mathcal{FL}^k for some $k \leq m$. To obtain ψ in the appropriate form, simply increase all the variable indices in the various conjuncts of ψ' and prepend vacuous universal quantifiers as required, and finally gather all the static conjuncts into a single static conjunct. \dashv

We begin by showing that \mathcal{FL}^3 has the exponential-sized model property. For the proof we define the notion of a *connector-type*, which resembles Mortimer’s notion of a *star* from his proof of the double exponential model property for FO^2 [8]. In Lemma 4.4 below we show that in contrast to the (doubly exponential) number of stars, it suffices to consider sets of connector-types of singly exponential size.

We repeat the normal-form (8) in the case $m = 3$ for convenience as

$$\forall x_1 \forall x_2 \forall x_3. \Omega \wedge \bigwedge_{i=1}^s \forall x_1 \forall x_2 (\alpha_i \rightarrow \exists x_3. \Gamma_i) \wedge \bigwedge_{j=1}^t \forall x_1 \forall x_2 (\beta_j \rightarrow \forall x_3. \Delta_j). \tag{9}$$

In the sequel, we shall additionally suppose that formulas do not contain any proposition letters. After all, when testing for satisfiability, the relevant truth-values can be simply guessed and eliminated accordingly. Since the complexities involved are all super-exponential, such a guessing processes can be carried out without additional cost.

For $k \geq 1$, and finite relational signature σ , a *fluted k -type over σ* is a maximal, consistent set of fluted k -literals over σ . We identify fluted k -types with their conjunctions when convenient. If \mathfrak{A} is a structure and \bar{a} a tuple from A , then there exists a unique fluted k -type τ satisfied by \bar{a} in \mathfrak{A} : we call τ the fluted k -type of \bar{a} , and denote it by $\text{ftp}^{\mathfrak{A}}[\bar{a}]$. If τ is a fluted k -type, we denote by $\tau^{[1]}$ the result of incrementing the indices of all the variables in τ , and if, in addition, $k \geq 2$, we denote by τ_{\uparrow} the result of removing all literals featuring the variable x_1 (i.e., all literals of arity k) and decrementing the indices of all variables. Notice that τ_{\uparrow} will be a fluted $(k - 1)$ -type; however, $\tau^{[1]}$ will not be a fluted $(k + 1)$ -type over σ if σ features any predicates of arity $(k + 1)$.

A *connector-type (over σ)* is a triple $\langle \pi, I, O \rangle$, where π is a fluted 1-type over σ , I is a set of fluted 2-types over σ such that, for all $\tau \in I$, $\tau_{\uparrow} = \pi$, and O is a set of fluted 2-types over σ . We refer to I as the connector-type’s *inputs*, and to O as its *outputs*. If \mathfrak{A} is any structure interpreting σ , and $b \in A$, we define

$$\text{Con}^{\mathfrak{A}}[b] = \langle \text{ftp}^{\mathfrak{A}}[b], \{ \text{ftp}^{\mathfrak{A}}[a, b] \mid a \in A \}, \{ \text{ftp}^{\mathfrak{A}}[b, c] \mid c \in A \} \rangle.$$

It is obvious that $\text{Con}^{\mathfrak{A}}[b]$ is a connector-type; we call it *the connector-type of b in \mathfrak{A}* .

Let φ be a formula of the form (9) over some signature σ . A connector-type $c = \langle \pi, I, O \rangle$ over σ is said to be *locally compatible with φ* if the following conditions hold:

- $\text{LC}\exists$. for every i ($1 \leq i \leq s$) and every $\tau \in I$ such that $\models \tau \rightarrow \alpha_i$, there exists $\tau' \in O$ such that the formula $\psi_i(x_1, x_2, x_3)$ is consistent, where

$$\psi_i(x_1, x_2, x_3) := \tau \wedge \tau'^{[1]} \wedge \Gamma_i \wedge \Omega \wedge \bigwedge_{j=1}^t (\beta_j \rightarrow \Delta_j);$$

LC \forall . for every $\tau \in I$ and every $\tau' \in O$, the formula $\psi(x_1, x_2, x_3)$ is consistent, where

$$\psi(x_1, x_2, x_3) := \tau \wedge \tau'^{[1]} \wedge \Omega \wedge \bigwedge_{j=1}^t (\beta_j \rightarrow \Delta_j).$$

The formulas ψ_i and ψ are of course not fluted.

LEMMA 4.2. *If $\mathfrak{A} \models \varphi$ and $b \in A$, then $\text{Con}^{\mathfrak{A}}[b]$ is locally compatible with φ .*

PROOF. Let $\text{Con}^{\mathfrak{A}}[b] = \langle \pi, I, O \rangle$. For LC \exists , pick $\tau \in I$ and i ($1 \leq i \leq s$). From the construction of I , select $a \in A$ such that $\mathfrak{A} \models \tau[a, b]$. If $\models \tau \rightarrow \alpha_i$, then, since $\mathfrak{A} \models \varphi$, select $c \in A$ such that the triple a, b, c satisfies $\Gamma_i \wedge \Omega \wedge \bigwedge_{j=1}^t (\beta_j \rightarrow \Delta_j)$ in \mathfrak{A} , and set $\tau' = \text{ftp}^{\mathfrak{A}}[b, c]$. Condition LC \forall is treated similarly. \dashv

A set C of connector-types is said to be *globally coherent* if the following conditions hold:

GC \exists . for all $\mathfrak{c} = \langle \pi, I, O \rangle \in C$ and all $\tau \in O$, there exists $\mathfrak{d} = \langle \pi', I', O' \rangle \in C$ such that $\tau \in I'$;

GC \forall . for all $\mathfrak{c} = \langle \pi, I, O \rangle \in C$ and all $\mathfrak{d} = \langle \pi', I', O' \rangle \in C$, $O \cap I' \neq \emptyset$.

LEMMA 4.3. *If \mathfrak{A} is a structure interpreting σ , then $\{\text{Con}^{\mathfrak{A}}[b] \mid b \in A\}$ is globally coherent.*

PROOF. Let \mathfrak{A} be a structure interpreting σ , $C = \{\text{Con}^{\mathfrak{A}}[b] \mid b \in A\}$ and $\mathfrak{c} = \langle \pi, I, O \rangle \in C$. Let $c \in A$ be such that $\mathfrak{c} = \text{Con}^{\mathfrak{A}}[c]$. Then by definition of $\text{Con}^{\mathfrak{A}}[c]$ we have $O = \{\text{ftp}^{\mathfrak{A}}[c, d] \mid d \in A\}$. For GC \exists , pick $\tau \in O$, say $\tau = \text{ftp}^{\mathfrak{A}}[c, d]$, for some $d \in A$, and choose $\mathfrak{d} = \text{Con}^{\mathfrak{A}}[d] = \langle \pi', I', O' \rangle$. Again by definition of $\text{Con}^{\mathfrak{A}}[d]$ we have $I' = \{\text{ftp}^{\mathfrak{A}}[a, d] \mid a \in A\}$. So, $\tau \in I'$. Condition GC \forall is treated similarly. \dashv

The relations between connector-types required by global coherence involve only fluted 2-types. This fact underlies the following lemma.

LEMMA 4.4. *If C is a globally coherent, nonempty set of connector-types, there exists a globally coherent, nonempty subset $D \subseteq C$ of cardinality at most $2^{|\sigma|}$.*

PROOF. Pick any connector-type in C , initialize D to be the singleton containing this connector-type, and initialize I^* to be the set of its inputs, I . We shall add connector-types to the set D and fluted 2-types to the set I^* , maintaining the invariant that I^* is the union of all the inputs of the connector-types in D . We call a connector-type in D *satisfied* if its outputs are also included in the set I^* .

Now execute the following procedure until D contains no unsatisfied connector-types. Pick some unsatisfied $\langle \pi, I, O \rangle \in D$ and some $\tau \in O \setminus I^*$. By GC \exists , there exists a connector-type $\langle \pi', I', O' \rangle \in C$ such that $\tau \in I'$. Set $D := D \cup \{\langle \pi', I', O' \rangle\}$ and $I^* := I^* \cup I'$. These assignments maintain the invariant on D and I^* . This process terminates after, say, $h < 2^{|\sigma|}$ steps, since there are only $2^{|\sigma|}$ fluted 2-types over σ , and $|I^*|$ increases by at least one at each step. When it does so, D is globally coherent, and indeed $|D| \leq h + 1$. \dashv

LEMMA 4.5. *Let $\varphi \in \mathcal{FL}^3$ be in normal form over a signature σ , and let φ have s existential conjuncts. The following are equivalent: (i) φ is satisfiable; (ii) there exists a nonempty, globally coherent set C of connector-types locally compatible with φ such that $|C| \leq 2^{|\sigma|}$; (iii) φ is satisfiable over a domain of size at most $s \cdot 2^{2^{|\sigma|}}$.*

PROOF. The implication (i) \Rightarrow (ii) follows from Lemmas 4.2, 4.3, and 4.4. The implication (iii) \Rightarrow (i) is trivial. It remains only to show (ii) \Rightarrow (iii).

Suppose that C is a nonempty, globally coherent set of connector-types locally compatible with φ and that $|C| \leq 2^{|\sigma|}$. We will define a model \mathfrak{A} in which the set of connector types of the elements coincides with C . The universe, A , of the intended model is partitioned into segments A_c of elements whose connector-types are precisely c , for every $c \in C$. We proceed in three stages. First we define fluted 1-type of every element of A_c as π , where $c = \langle \pi, I, O \rangle \in C$. Next, using $\text{GC}\exists$ and $\text{GC}\forall$, we assign fluted 2-types to every ordered pair of elements of A . At the third step, to ensure that $\mathfrak{A} \models \varphi$, we exploit $\text{LC}\exists$ and $\text{LC}\forall$ to assign fluted 3-types.

Denote the set of fluted 2-types over σ by $\text{ftp}^2(\sigma)$. For all $c \in C$ and all $\tau \in \text{ftp}^2(\sigma)$, let $A_{c,\tau}$ be a fresh set of s elements, $A_{c,\tau} = \{a_{c,\tau,1}, \dots, a_{c,\tau,s}\}$. For each $c \in C$, let $A_c = \bigcup_{\tau \in \text{ftp}^2(\sigma)} A_{c,\tau}$; and let $A = \bigcup_{c \in C} A_c$. Hence $|A| \leq s \cdot 2^{|\sigma|} \cdot |C| \leq s \cdot 2^{2|\sigma|}$.

First of all, we fix the fluted 1-type of every element $a \in A$. For every $c = \langle \pi, I, O \rangle \in C$ and every $a \in A_c$, set $\text{ftp}^1[a] = \pi$. Next, we fix the fluted 2-type of every ordered pair a, b of elements of A (possibly $a = b$). For every $c = \langle \pi, I, O \rangle \in C$ and every $\tau \in O$, pick some $\mathfrak{d} = \langle \pi', I', O' \rangle \in C$ such that $\tau \in I'$, by $\text{GC}\exists$, and then, for every $a \in A_c$ and every i ($1 \leq i \leq s$), set $\text{ftp}^2[a, a_{\mathfrak{d},\tau,i}] = \tau$. The conditions on connector-types ensure that this assignment of fluted 2-types is consistent with the fluted 1-type π' assigned to $a_{\mathfrak{d},\tau,i}$; moreover, no clashes can arise by the disjointness of the sets $A_{\mathfrak{d},\tau}$. These assignments having been made, we complete the assignment of fluted 2-types in \mathfrak{A} as follows. As long as there exist $a \in A_c$ and $b \in A_{\mathfrak{d}}$, with $c = \langle \pi, I, O \rangle \in C$ and $\mathfrak{d} = \langle \pi', I', O' \rangle \in C$, such that $\text{ftp}^2[a, b]$ is not yet defined, pick some $\tau \in O \cap I'$, by $\text{GC}\forall$, and set $\text{ftp}^2[a, b] = \tau$. Again, the conditions on connector-types ensure that this assignment is consistent with the fluted 1-type π' assigned to b .

At the end of this process, the interpretation of every unary and binary predicate has been defined in such a way that, if $b \in A_c$, where $c = \langle \pi, I, O \rangle \in C$, then:

- P1. for all $a \in A$, $\text{ftp}^1[a, b] \in I$;
- P2. for all $\tau \in O$, there exists a $\mathfrak{d} \in C$ such that, for each of the s elements $c \in A_{\mathfrak{d},\tau}$, $\text{ftp}^2[b, c] = \tau$;
- P3. for all $c \in A$, $\text{ftp}^2[b, c] \in O$.

We now assign fluted 3-types in \mathfrak{A} so as to ensure that $\mathfrak{A} \models \varphi$. We deal first with the existential conjuncts. Let a, b be any elements in A (not necessarily distinct). Let $b \in A_c$, where $c = \langle \pi, I, O \rangle$, and let $\text{ftp}^2[a, b] = \tau$. By P1, $\tau \in I$, so that, by $\text{LC}\exists$, for all i ($1 \leq i \leq s$), if $\models \tau \rightarrow \alpha_i$, then there exists $\tau' \in O$ such that

$$\tau \wedge \tau'^{[1]} \wedge \Gamma_i \wedge \Omega \wedge \bigwedge_{j=1}^i (\beta_j \rightarrow \Delta_j) \tag{10}$$

is consistent. By P2, pick some $\mathfrak{d} \in C$ such that, for each of the s elements $c \in A_{\mathfrak{d},\tau'}$, $\text{tp}^2[b, c] = \tau'$. Choose some fluted 3-type ξ consistent with (10), and set $\text{ftp}^3[a, b, a_{\mathfrak{d},i}] = \xi$. Notice that the only new predicate assignments involve ternary predicates, all binary predicates having been fixed previously. Clearly, then, we may carry out this process for each i ($1 \leq i \leq s$), and indeed for each pair of elements

$a, b \in A$, without clashes. At the end of this process, all existential conjuncts of φ will be satisfied (irrespective of how the model is completed): moreover, none of the fluted 3-types that have been set violates any of the static or universal conjuncts of φ .

Finally, suppose a, b and c are elements of A (not necessarily distinct) such that $\text{ftp}^{\mathfrak{A}}[a, b, c]$ has not been defined. Let $\text{ftp}^{\mathfrak{A}}[a, b] = \tau$ and $\text{ftp}^{\mathfrak{A}}[b, c] = \tau'$. Suppose $b \in A_c$, where $c = \langle \pi, I, O \rangle$, so that, by P1 and P3, $\tau \in I$ and $\tau' \in O$. By LC \forall ,

$$\tau \wedge \tau'^{[1]} \wedge \Omega \wedge \bigwedge_{j=1}^t (\beta_j \rightarrow \Delta_j) \tag{11}$$

is consistent. Choose some fluted 3-type ξ consistent with (11), and set $\text{ftp}^{\mathfrak{A}}[a, b, c] = \xi$. At the end of this process, \mathfrak{A} will be fully defined: moreover, none of the fluted 3-types that have been set violates any of the static or universal conjuncts of φ . Thus, $\mathfrak{A} \models \varphi$ as required. \dashv

We now turn our attention to obtaining a small model property for \mathcal{FL}^m for all $m \geq 3$. Our strategy is to show that, given an \mathcal{FL}^m formula φ with s existential conjuncts, we can compute an \mathcal{FL}^{m-1} formula ψ such that: if φ is satisfiable, then ψ is satisfiable; and if ψ has a model of size M , then φ has a model of size sM . However, the size of ψ and indeed the size of its signature, both increase by an exponential. We employ the standard apparatus of resolution theorem proving. Fixing some relational signature, let p be a predicate of arity k and let γ' and δ' be fluted k -clauses. Then, $\gamma = p(x_1, \dots, x_k) \vee \gamma'$ and $\delta = \neg p(x_1, \dots, x_k) \vee \delta'$ are also fluted k -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 $p(x_1, \dots, x_k)$. 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, k . It may be helpful to note the following at this point: (i) if γ and δ resolve to form ϵ , then $\models \forall x_1 \dots \forall x_k (\gamma \wedge \delta \rightarrow \epsilon)$; (ii) the fluted resolvent of two fluted k -clauses may or may not involve predicates of arity k ; (iii) in fluted resolution, the arguments of the literals in the fluted k -clauses undergo no change when forming the resolvent; (iv) if the fluted k -clause γ involves no predicates of arity k , then it cannot undergo fluted resolution at all.

If Γ is a set of fluted k -clauses, denote by Γ^* the smallest set of fluted k -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 fluted k -clause involving a predicate of arity k . Clearly, if Γ is finite, so are Γ^* and Γ° . Observe also that Γ° does not feature the variable x_1 .

The following lemma is, in effect, nothing more than the familiar completeness theorem for (ordered) propositional resolution.

LEMMA 4.6. *Let Γ be a set of fluted k -clauses, and τ a fluted $(k - 1)$ -type. If $\Gamma^\circ \cup \{\tau(x_2, \dots, x_k)\}$ is consistent, then there exists a fluted k -type τ^+ such that $\tau^+ \supseteq \tau(x_2, \dots, x_k)$ and $\Gamma \cup \{\tau^+\}$ is consistent.*

PROOF. Enumerate the k -ary predicates occurring in Γ as p_1, \dots, p_n . Note that none of these predicates occurs in τ . Define a *level- i extension* of τ inductively as

follows: (i) $\tau(x_2, \dots, x_k)$ is a level-0 extension of τ ; (ii) if τ' is a level- i extension of τ ($0 \leq i < n$), then $\tau' \cup \{p_{i+1}(x_1, \dots, x_k)\}$ and $\tau' \cup \{\neg p_{i+1}(x_1, \dots, x_k)\}$ are level- $(i + 1)$ extensions of τ . Thus, the level- n extensions of τ are exactly the fluted k -types over σ extending $\tau(x_2, \dots, x_k)$. If τ' is a level- i extension of τ ($0 \leq i \leq n$), we say that τ' *violates* a fluted k -clause δ if, for every literal in δ , the opposite literal is in τ' ; we say that τ' *violates* a set of fluted k -clauses Δ if τ' violates some $\delta \in \Delta$. Suppose now that τ' is a level- i extension of τ ($0 \leq i < n$). We claim that, if both $\tau' \cup \{p_{i+1}(x_1, \dots, x_k)\}$ and $\tau' \cup \{\neg p_{i+1}(x_1, \dots, x_k)\}$ violate Γ^* , then so does τ' . For otherwise, there must be a clause $\neg p_{i+1} \vee \gamma' \in \Gamma^*$ violated by $\tau' \cup \{p_{i+1}(x_1, \dots, x_k)\}$ and a clause $p_{i+1} \vee \gamma' \in \Gamma^*$ violated by $\tau' \cup \{\neg p_{i+1}(x_1, \dots, x_k)\}$. But in that case τ' violates the fluted resolvent $\gamma' \vee \delta'$, contradicting the supposition that τ' does not violate Γ^* . This proves the claim. Now, since $\tau(x_2, \dots, x_k)$ is by hypothesis consistent with Γ° , it certainly does not violate Γ° . Moreover, since it involves no predicates of arity k , $\tau(x_2, \dots, x_k)$ does not violate Γ^* either. By the above claim, then, there must be at least one level- n extension τ^+ of τ which does not violate $\Gamma^* \supseteq \Gamma$. Since τ^+ is a fluted k -type, this proves the lemma. \dashv

We require one further technical lemma.

LEMMA 4.7. *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 $\psi(x_1, \dots, x_{k-1}) = \exists x_k. \chi(x_1, \dots, x_k)$ is a first-order formula without equality, and $\mathfrak{B} \models \psi[b_1, \dots, b_{k-1}]$, then there exist at least z distinct elements b of B such that $\mathfrak{B} \models \chi[b_1, \dots, b_{k-1}, b]$.*

PROOF. Let $B = \{1, \dots, z\} \times A$. If p is any predicate of arity k , set $\langle (i_j, a_j), \dots, (i_k, a_k) \rangle \in p^{\mathfrak{B}}$ if and only if $\langle a_1, \dots, a_k \rangle \in p^{\mathfrak{A}}$. It is simple to verify that \mathfrak{B} has the desired properties. \dashv

These preliminaries having been dealt with, we are now ready obtain the promised small-model property for \mathcal{FL}^m .

LEMMA 4.8. *Let φ be a normal-form \mathcal{FL}^m -formula ($m \geq 3$) over some relational signature σ , and suppose that φ has $s \geq 1$ existential conjuncts and t universal conjuncts. Then there exists a normal-form \mathcal{FL}^{m-1} -formula φ' over a relational signature σ' 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 .*

PROOF. We repeat the form of φ given in (8) for convenience:

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

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 predicates of arity $m - 2$. The intended interpretation of $p_{i,J}(x_2, \dots, x_{m-1})$ is “for some x_1 , the tuple x_1, \dots, x_{m-1} satisfies α_i and also satisfies β_j for every $j \in J$,” and the intended interpretation of $q_J(x_2, \dots, x_{m-1})$ is “for some x_1 , the

tuple x_1, \dots, x_{m-1} satisfies β_j for every $j \in J$." Let φ' be the conjunction of the sentences

$$\bigwedge_{i=1}^s \bigwedge_{J \subseteq T} \forall x_1 \cdots \forall x_{m-1} ((\alpha_i \wedge \bigwedge_{j \in J} \beta_j) \rightarrow p_{i,J}(x_2, \dots, x_{m-1})); \tag{12}$$

$$\bigwedge_{J \subseteq T} \forall x_1 \cdots \forall x_{m-1} ((\bigwedge_{j \in J} \beta_j) \rightarrow q_J(x_2, \dots, x_{m-1})); \tag{13}$$

$$\bigwedge_{i=1}^s \bigwedge_{J \subseteq T} \forall x_2 \cdots \forall x_{m-1} (p_{i,J}(x_2, \dots, x_{m-1}) \rightarrow \exists x_m (\Gamma_i \cup \Omega \cup \bigcup_{j \in J} \Delta_j)^\circ); \tag{14}$$

$$\bigwedge_{J \subseteq T} \forall x_2 \cdots \forall x_{m-1} (q_J(x_2, \dots, x_{m-1}) \rightarrow \forall x_m (\Omega \cup \bigcup_{j \in J} \Delta_j)^\circ). \tag{15}$$

We claim that, if φ is satisfiable, then so is φ' . For suppose $\mathfrak{A} \models \varphi$. We expand \mathfrak{A} to a model $\mathfrak{A}' \models \varphi'$ by setting, for all i ($1 \leq i \leq s$) and all $J \subseteq T$,

$$\begin{aligned} p_{i,J}^{\mathfrak{A}'} &= \{ \langle a_2, \dots, a_{m-1} \rangle \mid \text{for some } a_1 \in A: \\ &\quad \mathfrak{A} \models \alpha_i[a_1, \dots, a_{m-1}] \text{ and } \mathfrak{A} \models \beta_j[a_1, \dots, a_{m-1}] \text{ for all } j \in J \}; \\ q_J^{\mathfrak{A}'} &= \{ \langle a_2, \dots, a_{m-1} \rangle \mid \text{for some } a_1 \in A: \\ &\quad \mathfrak{A} \models \beta_j[a_1, \dots, a_{m-1}] \text{ for all } j \in J \}. \end{aligned}$$

To see that $\mathfrak{A}' \models \varphi'$, we simply check the truth of conjuncts (12)–(15) in \mathfrak{A}' in turn. Sentences (12) and (13) are immediate. For (14), fix i and J , and suppose $\mathfrak{A}' \models p_{i,J}[a_2, \dots, a_{m-1}]$. By the definition of \mathfrak{A}' , let $a_1 \in A$ be such that $\mathfrak{A} \models \alpha_i[a_1, \dots, a_{m-1}]$ and $\mathfrak{A} \models \beta_j[a_1, \dots, a_{m-1}]$ for all $j \in J$. Since $\mathfrak{A} \models \varphi$, there exists b such that $\mathfrak{A} \models \Gamma_i[a_1, \dots, a_{m-1}, b]$, $\mathfrak{A} \models \Omega[a_1, \dots, a_{m-1}, b]$ and $\mathfrak{A} \models \Delta_j[a_1, \dots, a_{m-1}, b]$ for all $j \in J$. Since resolution is a valid inference step, $\mathfrak{A}' \models (\Gamma_i \cup \Omega \cup \bigcup_{j \in J} \Delta_j)^\circ[a_2, \dots, a_{m-1}, b]$. This establishes the truth of (14) in \mathfrak{A}' . Sentence (15) is handled similarly.

Conversely, we claim that, if φ' is satisfiable over a domain A , then φ is satisfiable over a domain of size $s \cdot |A|$. For suppose $\mathfrak{A} \models \varphi'$. Let \mathfrak{B} be the model of φ' guaranteed by Lemma 4.7, where $z = s$. We may assume that \mathfrak{A} and hence \mathfrak{B} interpret no predicates of arity m . We proceed to expand \mathfrak{B} to a model $\mathfrak{B}' \models \varphi$ by interpreting the predicates of arity m occurring in φ . Pick any tuple $\langle a_1, a_2, \dots, a_{m-1} \rangle$ from B (elements not necessarily distinct), and let J be the set of all j ($1 \leq j \leq t$) such that $\mathfrak{B} \models \beta_j[a_1, a_2, \dots, a_{m-1}]$. Suppose also that, for some i ($1 \leq i \leq s$) $\mathfrak{A} \models \alpha_i[a_1, a_2, \dots, a_{m-1}]$. From (12), $\mathfrak{B} \models p_{i,J}[a_2, \dots, a_{m-1}]$; and from (14), we may pick $b_i \in B$ such that $\mathfrak{B} \models (\Gamma_i \cup \Omega \cup \bigcup_{j \in J} \Delta_j)^\circ[a_2, \dots, a_{m-1}, b_i]$. From the properties of \mathfrak{B} due to Lemma 4.7, we know that if, for fixed a_1, \dots, a_{m-1} , we have $\mathfrak{A} \models \alpha_i[a_1, a_2, \dots, a_{m-1}]$ for more than one value of i , then we may choose the corresponding elements b_i so that they are all distinct. For each such b_i , then, let $\tau_i = \text{ftp}^{\mathfrak{B}}[a_2, \dots, a_{m-1}, b_i]$. Thus, $\tau_i(x_2, \dots, x_m)$ is consistent with $(\Gamma_i \cup \Omega \cup \bigcup_{j \in J} \Delta_j)^\circ$. By Lemma 4.6, there exists a fluted m -type $\tau_i^+ \supseteq \tau_i(x_2, \dots, x_m)$ such that τ_i^+ is consistent with $\Gamma_i \cup \Omega \cup \bigcup_{j \in J} \Delta_j$. Set $\text{ftp}^{\mathfrak{B}'}[a_1, a_2, \dots, a_{m-1}, b_i] = \tau_i^+$. Since $\tau_i^+ \supseteq \tau_i(x_2, \dots, x_m)$, only predicates

of arity m are being assigned, so that there is no clash with \mathfrak{B} . Moreover, since the b_i are all distinct, for a given tuple a_1, \dots, a_{m-1} , these assignments do not clash with each other. In this way, every existential conjunct of φ is witnessed in \mathfrak{B}' for the tuple a_1, a_2, \dots, a_{m-1} , and no static or universal conjunct of φ is violated for the tuples from B for which the m -ary predicates of σ have been defined. Now let $\langle a_1, a_2, \dots, a_{m-1}, a_m \rangle$ be any tuple from B for which the m -ary predicates of σ have not been defined, and let J be the set of all j ($1 \leq j \leq m$) such that $\mathfrak{B} \models \beta_j[a_1, a_2, \dots, a_{m-1}]$. From (13), $\mathfrak{B} \models q_J[a_2, \dots, a_{m-1}]$; and from (15), $\mathfrak{B} \models (\Omega \cup \bigcup_{j \in J} \Delta_j)^\circ [a_2, \dots, a_{m-1}, a_m]$. Let $\tau = \text{ftp}^{\mathfrak{B}}[a_2, \dots, a_{m-1}, a_m]$. Hence $\tau(x_2, \dots, x_m)$ is consistent with the set of clauses $(\Omega \cup \bigcup_{j \in J} \Delta_j)^\circ$. By Lemma 4.6, there exists a fluted m -type $\tau^+ \supseteq \tau(x_2, \dots, x_m)$ such that τ^+ is consistent with $\Omega \cup \bigcup_{j \in J} \Delta_j$. Now set $\text{ftp}^{\mathfrak{B}'}[a_1, a_2, \dots, a_{m-1}, a_m] = \tau^+$. Since $\tau^+ \supseteq \tau(x_2, \dots, x_m)$, only predicates of arity m are being assigned, so that there is no clash with \mathfrak{B} . Evidently, no static or universal conjunct of φ is violated in this process. Thus, $\mathfrak{B}' \models \varphi$, as required.

Conjuncts (12) and (13) do not involve x_m , while (14) and (15) do not involve x_1 . Thus, by decrementing variable indices where necessary, φ' becomes a formula of \mathcal{FL}^{m-1} . Indeed, re-writing implications as disjunctions in the obvious way, we have a formula in normal form: formulas (12) and (13) give the static conjunct; formula (14) gives the $2^t s$ existential conjuncts; and formula (15) gives the 2^t universal conjuncts. Exactly $2^t(s + 1)$ predicates of arity $(m - 2)$ have been added to the signature, and some m -ary predicates may have been lost. \dashv

LEMMA 4.9. *Let φ be a normal-form formula of \mathcal{FL}^m ($m \geq 3$) over a signature σ , and suppose that φ has s existential conjuncts ($s \geq 2$) and t universal conjuncts. If φ is satisfiable, then it is satisfiable over a domain of size at most $t(m - 2, 2|\sigma| + s + t)$.*

PROOF. We proceed by induction on m . If $m = 3$, Lemma 4.5 guarantees that φ has a model of size $s \cdot 2^{2|\sigma|} \leq 2^{2|\sigma|+s+t}$.

Suppose then $m > 3$, and that the lemma holds for smaller values of m . Let φ' be the formula guaranteed by Lemma 4.8, so that φ' is satisfiable. By inductive hypothesis φ' has a model of size at most $M = t(m - 3, 2(|\sigma| + 2^t(s + 1)) + 2^t s + 2^t)$, whence, by property (iv) of Lemma 4.8, φ has a model of size at most

$$s \cdot t(m - 3, 2(|\sigma| + 2^t(s + 1)) + 2^t s + 2^t) = s \cdot t(m - 3, 2|\sigma| + 3 \cdot 2^t(s + 1)) \leq t(m - 3, 2|\sigma| + 3 \cdot 2^t(s + 1) + s).$$

A routine calculation shows that, for $s \geq 2$, $|\sigma| \geq 1$ and $t \geq 0$, $2|\sigma| + 3 \cdot 2^t(s + 1) + s < 2^{2|\sigma|+s+t}$. This completes the inductive step, and proves the lemma. \dashv

Lemmas 4.1 and 4.9 thus yield the following upper complexity bounds.

THEOREM 4.10. *If $m \geq 3$, then any satisfiable formula of \mathcal{FL}^m has a model of $(m - 2)$ -tuply exponential size. Hence, the satisfiability problem for \mathcal{FL}^m is in $(m - 2)$ -NEXPTIME.*

PROOF. By Lemma 4.1 we may assume $\varphi \in \mathcal{FL}^m$ is in normal form; denote by σ the signature consisting of all predicate letters appearing in φ . Since φ is fluted

and features m variables, m is also the maximal arity of the predicate letters in σ . The finite model property just established allows one to test satisfiability of φ by checking existence of finite models of bounded cardinality. A σ -structure \mathfrak{A} can be guessed and verified to be a model of φ in time $O(|\varphi| \cdot |A|^m)$ [22]. \dashv

Thus, for $m \geq 3$, the satisfiability problem for \mathcal{FL}^m lies between $\lfloor m/2 \rfloor$ -NEXPTIME-hard and $(m - 2)$ -NEXPTIME. It is conceivable that, by using appropriate data-structures in place of connector-types, Lemma 4.8 might be generalized to yield an improved upper bound for all values of m . The present authors have, however, been unable to do so, even for the value $m = 5$. Small-model properties—and hence upper complexity bounds—for values of m up to 2 are easily derivable from known results. Trivially, all satisfiable \mathcal{FL}^0 -formulas have models with 1-element domains; and since \mathcal{FL}^1 is identical to the 1-variable fragment of first-order logic, any satisfiable \mathcal{FL}^1 -formula φ has a model of size bounded by $\|\varphi\|$. Moreover, \mathcal{FL}^2 is contained within the 2-variable fragment of first-order logic, which is shown to have the exponential-sized model property by Grädel, Kolaitis, and Vardi [4]. (Note that, for the fluted fragment, Lemma 4.5 above strengthens this to three variables.) Taking “0-NEXPTIME” to mean “NPTIME”, and noting that $4 - 2 = 4/2$, we see that \mathcal{FL}^m is $\lfloor m/2 \rfloor$ -NEXPTIME-complete for all m up to the value 4.

Finally, we mention that a measure of the complexity of the entire fluted fragment can be given within the framework of super-elementary classes developed in Schmitz [20]. Define the family of functions $F_\alpha : \mathbb{N} \rightarrow \mathbb{N}$, for $\alpha \in \mathbb{N}$, by

$$F_0(x) = x + 1, \quad F_{\alpha+1}(x) = F_\alpha^{[x+1]}(x),$$

where $f^{[k]}(x)$ denotes the k -fold iteration $f(\dots(f(x))\dots)$ of $f(x)$. This definition can be extended to transfinite ordinals α ; however, we need these functions only for very small finite α , and in particular, for $F_1(x) = 2x + 1$, $F_2(x) = 2^{x+1}(x + 1) - 1$, and $F_3(x)$, which satisfies $t(x + 1, x) \leq F_3(x) \leq t(2(x + 1), x)$. The complexity class TOWER is the class of languages recognizable by a Turing machine in time $F_3(p(n))$, where p is an elementary function. (It does not matter, at this level of complexity, whether the Turing machine in question is nondeterministic.) Completeness for TOWER is most naturally understood in terms of elementary-time many-one reductions. By a simple padding argument (*op. cit.*, p. 6), to show TOWER-hardness of a language Q , it suffices to find an elementary reduction from any P in $\text{TIME}(F_3(n))$ to Q . The proof of Theorem 3.2 provides just such a reduction. For let M be a Turing machine, running in $\text{TIME}(F_3(n))$, and recognizing the language P . Given an instance x of P with $|x| = n$, let (C, H, V) be a tiling system which has a solution over a grid of size $t(2(n + 1), n) \geq F_3(n)$ just in case M has a terminating run on input x . The fluted formula Ψ_w , as constructed in the proof of Theorem 3.2 (but with a grid of size $t(2(n + 1), n) \times t(2(n + 1), n)$), is satisfiable if and only if such a tiling exists. This completes the reduction.

On the other hand, the satisfiability problem for \mathcal{FL} is in TOWER. For let φ be a normal-form fluted formula. By Lemma 4.9, if φ is satisfiable, with $\|\varphi\| = n$, then φ has a model of size at most $t(n - 2, 4n)$, for sufficiently large n . All such models can be checked in time bounded by $F_3(n)$, again for sufficiently large n . Thus, we have

THEOREM 4.11. *The satisfiability problem for \mathcal{FL} is TOWER-complete.*

§5. Acknowledgments. The authors would like to thank the anonymous referees for their helpful suggestions. This work is supported by the Polish National Science Centre grants [2018/31/B/ST6/03662 to I.P.-H.] and [2013/09/B/ST6/01535 to W.S. and L.T.] and by the EPSRC grant [EP/K017438/1].

REFERENCES

- [1] H. ANDRÉKA, J. VAN BENTHEM, and I. NÉMETI, *Modal languages and bounded fragments of predicate logic*. *Journal of Philosophical Logic*, vol. 27 (1998), no. 3, pp. 217–274.
- [2] E. BÖRGER, E. GRÄDEL, and Y. GUREVICH, *The Classical Decision Problem*. Springer, Berlin, 1997.
- [3] E. GRÄDEL, *On the restraining power of guards*, this JOURNAL, vol. 64 (1999), no. 4, pp. 1719–1742.
- [4] E. GRÄDEL, P. KOLAITIS, and M. VARDI, *On the decision problem for two-variable first-order logic*. *Bulletin of Symbolic Logic*, vol. 3 (1997), no. 1, pp. 53–69.
- [5] A. HERZIG, *A new decidable fragment of first-order logic*. *Abstracts of the 3rd Logical Biennial Summer School and Conference in Honour of S. C. Kleene*, Varna, Bulgaria, June 1990.
- [6] U. HUSTADT, R. SCHMIDT, and L. GEORGIEVA, *A survey of decidable first-order fragments and description logics*. *Journal of Relational Methods in Computer Science*, vol. 1 (2004), no. 3, pp. 251–276.
- [7] C. LUTZ and U. SÄTTLER, *The complexity of reasoning with Boolean modal logics*, *Advances in Modal Logic*, vol. 3 (F. Wolter, H. Wansing, M. de Rijke, and M. Zakharyashev, editors), CLSI Publications, Menlo Park, 2002, pp. 329–348.
- [8] M. MORTIMER, *On languages with two variables*. *Mathematical Logic Quarterly*, vol. 21 (1975), no. 1, pp. 135–140.
- [9] A. NOAH, *Predicate-functors and the limits of decidability in logic*. *Notre Dame Journal of Formal Logic*, vol. 21 (1980), no. 4, pp. 701–707.
- [10] D. PLAISTED and S. GREENBAUM, *A structure-preserving clause form translation*. *Journal of Symbolic Computation*, vol. 2 (1986), no. 3, pp. 293–304.
- [11] I. PRATT-HARTMANN, W. SZWAST, and L. TENDERA, *Quine’s fluted fragment is nonelementary*, *25th EACSL Annual Conference on Computer Science Logic, CSL 2016* (J.-M. Talbot and L. Regnier, editors), Leibniz International Proceedings in Informatics, vol. 62, Schloß Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, 2016, pp. 39:1–39:21.
- [12] W. PURDY, *Decidability of fluted logic with identity*. *Notre Dame Journal of Formal Logic*, vol. 37 (1996), no. 1, pp. 84–104.
- [13] ———, *Fluted formulas and the limits of decidability*, this JOURNAL, vol. 61 (1996), no. 2, pp. 608–620.
- [14] ———, *Quine’s limits of decision*, this JOURNAL, vol. 64 (1999), no. 4, pp. 1439–1466.
- [15] ———, *Complexity and nicety of fluted logic*. *Studia Logica*, vol. 71 (2002), pp. 177–198.
- [16] W. V. QUINE, *On the limits of decision*, *Proceedings of the 14th International Congress of Philosophy, vol. III*, Herder, Vienna, 1969, pp. 57–62.
- [17] ———, *Algebraic logic and predicate functors*, *The Ways of Paradox*, revised and enlarged ed., Harvard University Press, Cambridge, MA, 1976, pp. 283–307.
- [18] ———, *The variable*, *The Ways of Paradox*, revised and enlarged ed., Harvard University Press, Cambridge, MA, 1976, pp. 272–282.
- [19] R. SCHMIDT and U. HUSTADT, *A resolution decision procedure for fluted logic*, *Automated Deduction—CADE-17* (D. McAllester, editor), Lecture Notes in Artificial Intelligence, vol. 1831, Springer-Verlag, Berlin, 2000, pp. 433–448.
- [20] S. SCHMITZ, *Complexity hierarchies beyond Elementary*. *ACM Transactions on Computation Theory*, vol. 8 (2016), no. 1, pp. 3:1–3:26.
- [21] L. STOCKMEYER, *The complexity of decision problems in automata and logic*, Ph.D. thesis, Massachusetts Institute of Technology, Computer Science Laboratory, 1974, Report MAC-TR-133.
- [22] M. VARDI, *On the complexity of bounded-variable queries*, *Proceedings of the Fourteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, ACM Digital Library, 1995, pp. 266–276.
- [23] M. VOIGT, *A fine-grained hierarchy of hard problems in the separated fragment*, *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS, 2017*, IEEE Xplore Digital Library, 2017, pp. 1–12.

SCHOOL OF COMPUTER SCIENCE
MANCHESTER UNIVERSITY
MANCHESTER, M13 9PL, UK

and

INSTYTUT INFORMATYKI
UNIwersytet OPOLSKI
45-040 OPOLE, POLAND

E-mail: ipratt@cs.man.ac.uk

INSTYTUT INFORMATYKI
UNIwersytet OPOLSKI
45-040 OPOLE, POLAND

E-mail: szwast@math.uni.opole.pl

E-mail: tendera@math.uni.opole.pl