Quine's Fluted Fragment is Non-Elementary*

Ian Pratt-Hartmann[1], Wiesław Szwast[2], and Lidia Tendera[3]

1 **School of Computer Science, University of Manchester, UK**
ipratt@cs.man.ac.uk
2 **Institute of Mathematics and Informatics, Opole University, Poland**
szwast@math.uni.opole.pl
2 **Institute of Mathematics and Informatics, Opole University, Poland**
tendera@math.uni.opole.pl

Abstract

We study the fluted fragment, a decidable fragment of first-order logic with an unbounded number of variables, originally identified by W.V. Quine. We show that the satisfiability problem for this fragment has non-elementary complexity, thus refuting an earlier published claim by W.C. Purdy that it is in NExpTime. More precisely, we consider, for all m greater than 1, the intersection of the fluted fragment and the m-variable fragment of first-order logic. We show that this sub-fragment forces $(m/2)$-tuply exponentially large models, and that its satisfiability problem is $(m/2)$-NExpTime-hard. We round off by using a corrected version of Purdy's construction to show that the m-variable fluted fragment has the m-tuply exponential model property, and that its satisfiability problem is in m-NExpTime.

1998 ACM Subject Classification F.4.1 Mathematical logic

Keywords and phrases Quine, fluted fragment, Purdy, non-elementary, satisfiability, decidability

Digital Object Identifier 10.4230/LIPIcs.CSL.2016.39

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
$$\forall x_1(\text{student}(x_1) \to \neg\forall x_2(\text{prof}(x_2) \to \text{admires}(x_1, x_2))) \tag{1}$$

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

Furthermore, as was observed in [4], various standard translations of multi-modal logic into first-order logic are also easily seen to yield only fluted formulas. The origins of the fluted fragment can be traced to a paper given by W.V. Quine to the 1968 *International Congress of Philosophy* [11], 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 $x_1, \ldots, x_m$. Boolean operators and quantifiers may be freely applied, except that the order of quantification must follow the order of arguments: a quantifier

* This work is supported by the Polish National Science Centre grant DEC-2013/09/B/ST6/01535 to W.S. and L.T., and by the EPSRC grant EP/K017438/1 to I.P-H.

binding an occurrence of x_i may only be applied to a subformula in which all occurrences of $x_{i+1}, \ldots, x_m$ are already bound. Quine explained how Herbrand's decision procedure for monadic first-order logic easily extends to cover all homogeneous m-adic formulas. The term *fluted logic* first appears (to the present authors' knowledge) in [13], 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 satisfiabilty. It seems that the allusion is architectural rather than musical: we are invited to think of arguments of predicates as being 'lined up' in columns. Quine's motivation for defining the fluted fragment 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' [12, 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* [5] maps, under the standard first-order translation, to $\mathcal{FL}$ – in fact, to $\mathcal{FL}^2$, the fluted fragment restricted to just two variables. (Thus, $\mathcal{FL}^2$ in effect subsumes the description logic $\mathcal{ALC}$.) On the other hand, even $\mathcal{FL}^2$ is not contained within the so-called *guarded fragment* of first-order logic [1]: the formula (1), 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 [4].

Noah [6] pointed out, however, that – contrary to Quine's assertion – Herbrand's technique does not obviously extend from homogeneous m-adic logic to the fluted fragment, and consequently, the decidability of the satisfiability problem for the latter should be regarded as open. This problem – together with the corresponding problems for various extensions of the fluted fragment – was considered in a series of papers in the 1990s by W.C. Purdy [7, 8, 9, 10]. The decidability of $\mathcal{FL}$ is proved in [8], while in [10] it is claimed (Corollary 10) 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 (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 non-elementary. On the other hand, we also show that any satisfiable formula of the m-variable fluted fragment has a model of m-tuply exponential size, so that the satisfiability problem for this sub-fragment is contained in m-NExpTime. Thus, $\mathcal{FL}$ has the finite model property, and its satisfiability (= finite satisfiability) problem is decidable, but not elementary. Note that the above complexity bounds for $\mathcal{FL}^m$ leave a gap of a factor of 2.

We mention at this point another incorrect claim by Purdy concerning an extension of the fluted fragment. In [9], the author considers what he calls *extended fluted logic*, 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

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

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. We also mention that an independent decision procedure for the fluted fragment – based on resolution theorem-proving – was presented in [15]. No complexity bounds are given there. Moreover, that paper omits detailed proofs, and these have, to the authors' knowledge, never been published.

The structure of this paper 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 [10]. In Section 4, we show how some of the constructions appearing in Purdy's paper can nevertheless be recycled to give a proof that the fluted fragment with m variables does indeed have the finite model property, and that its satisfiability problem is in m-NExpTime. This proof is shorter and more perspicuous than the original argument for the decidability of $\mathcal{FL}$ given in [8], and, moreover, yields detailed complexity information.

2 Preliminaries

Let $\bar{x}_\omega = x_1, x_2, \ldots$ be a fixed sequence of variables. We define the sets of formulas $\mathcal{FL}^{[k]}$ (for $k \geq 0$) by structural induction as follows: (i) any atom $p(x_\ell, \ldots, x_k)$, where $x_\ell, \ldots, x_k$ is a contiguous subsequence of $\bar{x}_\omega$, is in $\mathcal{FL}^{[k]}$; (ii) $\mathcal{FL}^{[k]}$ is closed under boolean combinations; (iii) if φ is in $\mathcal{FL}^{[k+1]}$, then $\exists x_{k+1}\varphi$ and $\forall x_{k+1}\varphi$ are in $\mathcal{FL}^{[k]}$. 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 suffix of some prefix $x_1, \ldots, x_k$ of $\bar{x}_\omega$; and when quantifying, only the last variable in this sequence may be bound, as illustrated by the fluted sentences in (1) and (2). Note that, in this paper, we consider only purely relational signatures.

Denote by $\mathcal{FL}^m$ the sub-fragment of $\mathcal{FL}$ consisting of those formulas featuring at most m variables, free or bound. Do not confuse $\mathcal{FL}^m$ (the set of fluted formulas with m variables, free or bound) with $\mathcal{FL}^{[m]}$ (the set of fluted formulas with m *free* variables). Thus, the formulas in (1) and (2) are in $\mathcal{FL}^{[0]}$; however (1) is in $\mathcal{FL}^m$ only for $m \geq 2$, and (2) is in $\mathcal{FL}^m$ only for $m \geq 3$. All formulas occurring in the remainder of the paper will be fluted.

In the sequel, we employ standard concepts and notation from first-order logic. Structures are denoted by Gothic capital letters and their domains by the corresponding Roman capitals. If $\mathfrak{A}$ is a structure, $\varphi(x_1, \ldots, x_k)$ a formula of $\mathcal{FL}^{[k]}$, and $\bar{a}$ a k-tuple of elements of A, then we write $\mathfrak{A} \models \varphi[\bar{a}]$ to indicate that $\bar{a}$ satisfies $\varphi(x_1, \ldots, x_k)$ in $\mathfrak{A}$; in the case where φ is a fluted sentence and $\mathfrak{A}$ is a model of φ, we write simply $\mathfrak{A} \models \varphi$. If φ is a formula, we write $\|\varphi\|$ to denote the number of symbols in φ. We use $\pm\varphi$ to stand either for φ or $\neg\varphi$, with multiple occurrences of $\pm$ in displayed material resolved *uniformly*: thus, for example $\pm p_i^1(x_1, x_2) \to \pm p_i^1(x_2)$ stands for a *pair* (not a quartet) of formulas, namely, $p_i^1(x_1, x_2) \to p_i^1(x_2)$ and $\neg p_i^1(x_1, x_2) \to \neg p_i^1(x_2)$.

3 Lower bound

In this section, we establish lower complexity bounds for the fluted fragment, which we express using the *tetration function* $\mathfrak{t}(k,n)$, defined, for $n,k \geq 0$, by induction as follows:

$$\mathfrak{t}(0,n) = n$$
$$\mathfrak{t}(k+1,n) = 2^{\mathfrak{t}(k,n)}.$$

Thus, $\mathfrak{t}(1,n) = 2^n$, $\mathfrak{t}(2,n) = 2^{2^n}$, and so on. Theorem 1 shows that an $\mathcal{FL}^{2m}$-formula of size $O(n^2)$ can force models of size at least $\mathfrak{t}(m,n)$, thus contradicting Corollary 10 of [10]. Theorem 2 shows that the satisfiability problem for $\mathcal{FL}^{2m}$ is m-NExpTime-hard, thus contradicting Theorem 11 of [10].

As a preliminary, for any $z \geq 0$, we take the (*canonical*) *representation* of any integer n in the range $(0 \leq j < 2^z)$ to be the bit-string $\bar{s} = s_{z-1}, \ldots, s_0$ 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 $s'_{z-1}, \ldots, s'_0$, then $n' = n - 1 \mod 2^z$ 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 1.** *For all $m \geq 1$, there exists a sequence of satisfiable sentences $\{\varphi_n\}_{n \in \mathbb{N}} \in \mathcal{FL}^{2m}$ such that $\|\varphi_n\|$ grows polynomially with m and n (and indeed quadratically in n for fixed m), but the smallest satisfying model of φ_n has at least $\mathfrak{t}(m,n)$ elements. Hence, there is no elementary bound on the size of models of satisfiable sentences in $\mathcal{FL}$.*

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

- unary predicates $p_0, \ldots, p_{n-1}$;
- 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 in_k, 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, $\mathrm{val}_k(b)$, between 0 and $\mathfrak{t}(k,n) - 1$. For $k = 1$, this value will be encoded by b's satisfaction of the unary predicates $p_0, \ldots, p_{n-1}$. Specifically, for any 1-integer b, define $\mathrm{val}_1(b)$ to be the integer canonically represented by the n-element bit-string $s_{n-1}, \ldots, s_0$, 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 $\mathrm{val}_{k+1}(b)$ will be encoded by how b is related to the various k-integers via the predicate in_k. Specifically, for any k $(1 \leq k < m)$ and any $(k+1)$-integer b, define $\mathrm{val}_{k+1}(b)$ to be the integer canonically represented by the bit-string $s_{N-1}, \ldots, s_0$ of length $N = \mathfrak{t}(k,n)$ where, for all i $(0 \leq i < N)$,

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

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

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

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

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

The proof proceeds by writing 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 $(1 \le k \le m)$. It follows from m-covering that $|A| \ge \mathfrak{t}(m,n)$, 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 $\mathrm{zero}_1, \ldots, \mathrm{zero}_m$; (ii) the binary predicates $\mathrm{pred}_{1,0}, \ldots, \mathrm{pred}_{m,0}$; and (iii) the ternary predicates $\mathrm{pred}_{1,1}, \ldots, \mathrm{pred}_{m-1,1}$. Furthermore, we take $\Sigma_{m,n}$ to contain, for all k $(1 \le k \le m)$ and all ℓ $(0 \le \ell \le 2(m-k))$, the $(\ell+2)$-ary predicate $\mathrm{eq}_{k,\ell}$. Observe that $\Sigma_{m,n}$ does not contain the predicate $\mathrm{pred}_{m,1}$. Observe also that, as k increases from 1 to m, the maximal value of the index ℓ in the predicates $\mathrm{eq}_{k,\ell}$ decreases, in steps of 2, from $2m-2$ down to 0; hence the maximal arity of these predicates decreases from $2m$ to 2.

Any model $\mathfrak{A} \models \Phi_{m,n}$ will be guaranteed to satisfy the following properties for all k $(1 \le k \le m)$ concerning the interpretation of these predicates.

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

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

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

The bounds on ℓ in the property k-predecessor 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 $\mathrm{pred}_{m,1}$.)

Thus, $\mathrm{zero}_k(x_1)$ can be read as "x_1 is zero", $\mathrm{pred}_{k,0}(x_1,x_2)$, as "$x_2$ is the predecessor of x_1", and $\mathrm{pred}_{k,1}(x_1,x_2,x_3)$ as "x_3 is the predecessor of x_1". Notice that, in the latter case, the argument x_2 is semantically inert. Similarly, $\mathrm{eq}_{k,\ell}(x_1,\ldots,x_{\ell+2})$ can be read as "x_1 is equal to $x_{\ell+2}$", with the ℓ arguments $x_2,\ldots,x_{\ell+1}$ 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.

We suppose that $\mathfrak{A} \models \Phi_{m,n}$, and establish the properties k-covering, k-harmony, k-zero, k-equality and k-predecessor for all k $(1 \le k \le m)$ by induction on k. For ease of reading, we introduce the various conjuncts of $\Phi_{m,n}$ as they are required. Appeals to the inductive hypothesis are indicated by the initials IH.

Base case ($k = 1$): Let b be a 1-integer, and recall that $\mathrm{val}_1(b)$ is defined by b's satisfaction of the predicates $p_0,\ldots,p_{n-1}$. 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

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

Thus, if a is a 1-integer, $\mathfrak{A} \models \mathrm{zero}_1[b] \Leftrightarrow \mathrm{val}_1(b) = 0$. 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

$$\bigwedge_{i=0}^{n-1} \bigwedge_{\ell=0}^{L} \forall x_1(\mathrm{int}_1(x_1) \wedge \pm p_i(x_1) \to \forall x_2 \cdots \forall x_{\ell+1} \pm p_i^\ell(x_1, \ldots, x_{\ell+1})). \tag{Φ_2}$$

Note that this really is a *pair* of formulas: the two 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 had 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 ($0 \leq \ell < 2m-2$), let $\varepsilon_{1,\ell}(x_1, \ldots, x_{\ell+2})$ abbreviate the formula: $\bigwedge_{i=0}^{n-1}(p_i^{\ell+1}(x_1, \ldots, x_{\ell+2}) \leftrightarrow p_i(x_{\ell+2}))$. We see from (3) that $\varepsilon_{1,\ell}(x_1, \ldots, x_{\ell+2})$ 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

$$\begin{aligned}\bigwedge_{\ell=0}^{2m-2} &\forall x_1(\mathrm{int}_1(x_1) \to \\ &\quad \forall x_2 \cdots \forall x_{\ell+2}(\mathrm{int}_1(x_{\ell+2}) \to \mathrm{eq}_{1,\ell}(x_1, \ldots, x_{\ell+2}) \leftrightarrow \varepsilon_{1,\ell}(x_1, \ldots, x_{\ell+2}))).\end{aligned} \tag{Φ_3}$$

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

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

$$\begin{aligned}\bigwedge_{i=0}^{n-1} &([\bigwedge_{j=0}^{i-1} \neg p_j^{\ell+1}(x_1, \ldots, x_{\ell+2}) \to (p_i^{\ell+1}(x_1, \ldots, x_{\ell+2}) \leftrightarrow \neg p_i(x_{\ell+2}))] \wedge \\ &\quad [\bigvee_{j=0}^{i-1} p_j^{\ell+1}(x_1, \ldots, x_{\ell+2}) \to (p_i^{\ell+1}(x_1, \ldots, x_{\ell+2}) \leftrightarrow p_i(x_{\ell+2}))]).\end{aligned}$$

From our preliminary remarks on the canonical representations of numbers by bit-strings, we see that $\pi_{1,\ell}(x_1, \ldots, x_{\ell+2})$ 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 $x_1 \mod 2^n$. We then add to $\Phi_{m,n}$ the conjuncts

$$\begin{aligned}\bigwedge_{\ell=0}^{1} &\forall x_1(\mathrm{int}_1(x_1) \to \\ &\quad \forall x_2 \cdots \forall x_{\ell+2}(\mathrm{int}_1(x_{\ell+2}) \to (\mathrm{pred}_{1,\ell}(x_1, \ldots, x_{\ell+2}) \leftrightarrow \pi_{1,\ell}(x_1, \ldots, x_{\ell+2})))),\end{aligned} \tag{Φ_4}$$

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 $\ell = 0$; this suffices to satisfy the property 1-predecessor. Observe that care is required in this case, because $\Sigma_{1,n}$ does not feature the ternary predicate $\mathrm{pred}_{1,1}$ – indeed, it features no ternary predicates at all.

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

$$\exists x_1(\mathrm{int}_1(x_1) \wedge \mathrm{zero}_1(x_1)) \tag{Φ_5}$$

$$\forall x_1(\mathrm{int}_1(x_1) \rightarrow \exists x_2(\mathrm{int}_1(x_2) \wedge \mathrm{pred}_{1,0}(x_1, x_2))). \tag{Φ_6}$$

Observe that (Φ_6) features only $\mathrm{pred}_{1,0}$, and not $\mathrm{pred}_{1,1}$, 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 $\mathrm{val}_k : \mathrm{int}_k^{\mathfrak{A}} \rightarrow [0, \mathfrak{t}(k, n) - 1]$ 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:

$$\forall x_1(\mathrm{int}_k(x_1) \rightarrow \forall x_2(\mathrm{int}_{k+1}(x_2) \wedge \pm\mathrm{in}_k(x_1, x_2) \rightarrow \\ \forall x_3(\mathrm{int}_k(x_3) \wedge \mathrm{eq}_{k,1}(x_1, x_2, x_3) \rightarrow \pm\mathrm{out}_k(x_2, x_3)))). \tag{Φ_7}$$

If a, a' are k-integers such that $\mathrm{val}_k(a) = \mathrm{val}_k(a')$, and b is any $(k + 1)$-integer, then, by k-equality (IH), $\mathfrak{A} \models \mathrm{eq}_{k,1}[a, b, a']$, whence ($\Phi_7$) 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 $\mathrm{val}_{k+1}(b)$ 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 $\mathrm{val}_k(a) = i$; and by $(k + 1)$-harmony (just established), all such k-integers a agree on what the ith bit in $\mathrm{val}_{k+1}(b)$ should be.

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

$$\forall x_1(\mathrm{int}_{k+1}(x_1) \rightarrow (\mathrm{zero}_{k+1}(x_1) \leftrightarrow \forall x_2(\mathrm{int}_k(x_2) \rightarrow \neg\mathrm{out}_k(x_1, x_2)))). \tag{Φ_8}$$

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

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

$$\mathfrak{A} \models \mathrm{in}_k^{\triangleleft}[a, b] \Leftrightarrow (\text{for any } k\text{-integer } a', \mathrm{val}_k(a') < \mathrm{val}_k(a) \Rightarrow \mathfrak{A} \not\models \mathrm{in}_k[a', b]); \tag{4}$$

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

Condition (4) allows us to read $\mathrm{in}_k^{\triangleleft}(x_1, x_2)$ 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

$$\forall x_1(\mathrm{int}_k(x_1) \rightarrow \forall x_2(\mathrm{int}_{k+1}(x_2) \wedge \pm\mathrm{in}_k^{\triangleleft}(x_1, x_2) \rightarrow \\ \forall x_3(\mathrm{int}_k(x_3) \wedge \mathrm{eq}_{k,1}(x_1, x_2, x_3) \rightarrow \pm\mathrm{out}_k^{\triangleleft}(x_2, x_3)))). \tag{Φ_9}$$

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

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

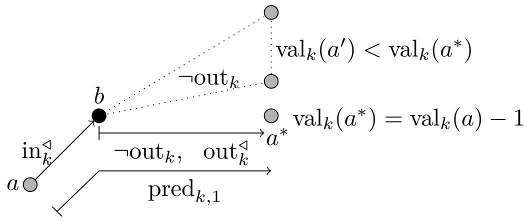

Figure 1 Fixing the interpretation of $\mathrm{in}^{\triangleleft}_k$.

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

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

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

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

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

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

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

In the sequel, we shall employ a formula denoted $\varrho_{k,0}(x_1,\ldots,x_4)$, which will be of interest where x_1 and x_4 are k-integers *with equal values*, while x_2 and x_3 are $(k+1)$-integers. Intuitively this formula says: "the $\mathrm{val}_k(x_1)$th – equivalently, $\mathrm{val}_k(x_4)$th – digits of x_2 and x_3 are as they should be if $\mathrm{val}_{k+1}(x_3) = \mathrm{val}_{k+1}(x_2) - 1 \mod \mathfrak{t}(k+1,n)$." Similarly, we shall employ the formula $\varrho_{k,1}(x_1,\ldots,x_5)$ which will be of interest where x_1 and x_5 are k-integers with equal values, while x_2 and x_4 are $(k+1)$-integers. Intuitively this formula says: "the $\mathrm{val}_k(x_1)$th – equivalently, $\mathrm{val}_k(x_5)$th – digits of x_2 and x_4 are as they should be if $\mathrm{val}_{k+1}(x_4) = \mathrm{val}_{k+1}(x_2) - 1 \mod \mathfrak{t}(k+1,n)$." Note that, in the latter case, x_3 is a dummy variable. Formally, for $0 \leq \ell \leq 1$, define $\varrho_{k,\ell}(x_1,\ldots,x_{\ell+4})$ to be the formula

$$\begin{aligned}
&[\mathrm{in}_k^{\triangleleft\ell+2}(x_1,\ldots,x_{\ell+4}) \to (\mathrm{in}_k^{\ell+2}(x_1,\ldots,x_{\ell+4}) \leftrightarrow \neg\mathrm{out}_k(x_{\ell+3},x_{\ell+4}))] \wedge \\
&[\neg\mathrm{in}_k^{\triangleleft\ell+2}(x_1,\ldots,x_{\ell+4}) \to (\mathrm{in}_k^{\ell+2}(x_1,\ldots,x_{\ell+4}) \leftrightarrow \mathrm{out}_k(x_{\ell+3},x_{\ell+4}))].
\end{aligned}$$

Suppose that a, a' are k-integers, b, b' $(k+1)$-integers and $\bar{c}$ an ℓ-tuple of elements. From (Φ_{12}), $\mathfrak{A} \models \mathrm{in}_k^{\ell+4}[a,b,\bar{c},b',a'] \Leftrightarrow \mathfrak{A} \models \mathrm{in}_k[a,b]$, and from (Φ_{13}), $\mathfrak{A} \models \mathrm{in}_k^{\triangleleft\ell+2}[a,b,\bar{c},b',a'] \Leftrightarrow \mathfrak{A} \models \mathrm{in}_k^{\triangleleft}[a,b]$. Furthermore, by $(k+1)$-harmony, $\mathfrak{A} \models \mathrm{out}_k[b',a'] \Leftrightarrow \mathfrak{A} \models \mathrm{in}_k[a',b']$. 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 $\mathrm{val}_k(a')$-th digit in the encoding of $\mathrm{val}_{k+1}(b')$ is the same as the $\mathrm{val}_k(a)$-th digit in the encoding of $\mathrm{val}_{k+1}(b) - 1$, modulo $\mathfrak{t}(k+1,n)$.

Now add to $\Sigma_{m,n}$ the ternary predicate $\mathrm{predDig}_{k+1,0}$ and quaternary predicate $\mathrm{predDig}_{k+1,1}$, and add to $\Phi_{m,n}$ the conjunct

$$\begin{aligned}
\bigwedge_{\ell=0}^{1} \forall x_1(\mathrm{int}_k(x_1) \to \forall x_2(\mathrm{int}_{k+1}(x_2) \to \forall x_3 \cdots \forall x_{\ell+3}(\mathrm{int}_{k+1}(x_{\ell+3}) \to \\
\forall x_{\ell+4}(\mathrm{int}_k(x_{\ell+4}) \wedge \mathrm{eq}_{k,\ell+2}(x_1,\ldots,x_{\ell+4}) \to \\
(\mathrm{predDig}_{k+1,\ell}(x_2,\ldots,x_{\ell+4}) \leftrightarrow \varrho_{k,\ell}(x_1,\ldots,x_{\ell+4}))))). \qquad (\Phi_{14})
\end{aligned}$$

This formula is illustrated in the left-hand diagram of Fig. 2 in the case $\ell = 1$: here, $\varrho_{k,1}$ holds of the tuple a, b, c, b', a, and $\mathrm{predDig}_{k+1,1}$ of the tuple b, c, b', a, just in case the $\mathrm{val}_k(a)$th digit of $\mathrm{val}_{k+1}(b')$ agrees with the $\mathrm{val}_k(a)$th digit of $\mathrm{val}_{k+1}(b) - 1$. (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), $\mathfrak{A} \models \mathrm{eq}_{k,\ell+2}[a,b,\bar{c},b',a]$, and from the properties of $\varrho_{k,\ell}$ just established (setting $a' = a$), $\mathfrak{A} \models \mathrm{predDig}_{k+1,\ell}[b,\bar{c},b',a]$ just in case the $\mathrm{val}_k(a)$th digit of $\mathrm{val}_{k+1}(b')$ is equal to the $\mathrm{val}_k(a)$th digit of $\mathrm{val}_{k+1}(b) - 1$, modulo $\mathfrak{t}(k+1,n)$.

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

$$\begin{aligned}
\bigwedge_{\ell=0}^{1} \forall x_1(\mathrm{int}_{k+1}(x_1) \to \forall x_2 \cdots \forall x_{\ell+2}(\mathrm{int}_{k+1}(x_{\ell+2}) \to \\
(\mathrm{pred}_{k+1,\ell}(x_1,\ldots,x_{\ell+2}) \leftrightarrow \forall x_{\ell+3}(\mathrm{int}_k(x_{\ell+3}) \to \mathrm{predDig}_{k+1,\ell}(x_1,\ldots,x_{\ell+3}))))). \quad (\Phi_{15})
\end{aligned}$$

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

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, $\mathrm{in}_k^{\triangleleft 2}$, $\mathrm{predDig}_{k+1,0}$ (not in_k^3, $\mathrm{in}_k^{\triangleleft 3}$ or $\mathrm{predDig}_{k+1,1}$), 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 $\Sigma_{m,n}$.

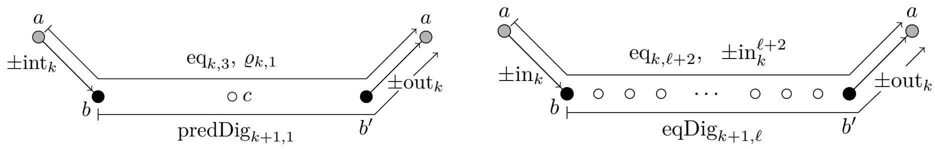

Figure 2 Fixing the interpretations of $\mathrm{predDig}_{k+1,1}$ (left) and $\mathrm{eqDig}_{k+1,\ell}$ (right).

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

$$\exists x_1(\mathrm{int}_{k+1}(x_1) \wedge \mathrm{zero}_{k+1}(x_1)) \qquad (\Phi_{16})$$

$$\forall x_1(\mathrm{int}_{k+1}(x_1) \to \exists x_2(\mathrm{int}_{k+1}(x_2) \wedge \mathrm{pred}_{k+1,0}(x_1, x_2))). \qquad (\Phi_{17})$$

Note that (Φ_{17}) features only $\mathrm{pred}_{k+1,0}$, and not $\mathrm{pred}_{k+1,1}$, 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 $\mathrm{eq}_{k+1,\ell}$ for 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

$$\bigwedge_{\ell=0}^{L} \forall x_1(\mathrm{int}_k(x_1) \to \forall x_2(\mathrm{int}_{k+1}(x_2) \wedge \pm\mathrm{in}_k(x_1, x_2) \to \forall x_3 \cdots \forall x_{\ell+2} \pm \mathrm{in}_k^\ell(x_1, x_2, \ldots, x_{\ell+1}, x_{\ell+2}))), \qquad (\Phi_{18})$$

thus fixing in_k^ℓ to be the result of adding ℓ semantically inert arguments to in_k^ℓ. (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 $\mathrm{eqDig}_{k+1,\ell}$, and add to $\Phi_{m,n}$ the conjuncts

$$\begin{aligned}\bigwedge_{\ell=0}^{L} \forall x_1(\mathrm{int}_k(x_1) \to \forall x_2(\mathrm{int}_{k+1}(x_2) \to \forall x_3 \cdots \forall x_{\ell+3}(\mathrm{int}_{k+1}(x_{\ell+3}) \to\\ \forall x_{\ell+4}(\mathrm{int}_k(x_{\ell+4}) \wedge \mathrm{eq}_{k,\ell+2}(x_1, \ldots, x_{\ell+4}) \to\\ (\mathrm{eqDig}_{k+1,\ell}(x_2, \ldots, x_{\ell+4}) \leftrightarrow \eta_{k,\ell+2}(x_1, \ldots, x_{\ell+4})))))),\end{aligned} \qquad (\Phi_{19})$$

where $\eta_{k,\ell+2}(x_1, \ldots, x_{\ell+4})$ is the formula: $\mathrm{in}_k^{\ell+2}(x_1, \ldots, x_{\ell+4}) \leftrightarrow \mathrm{out}_k(x_{\ell+3}, x_{\ell+4})$.

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 $\mathfrak{A} \models \mathrm{eqDig}_{k+1,\ell}[b, \bar{c}, b', a]$ just in case $\mathrm{val}_{k+1}(b)$ and $\mathrm{val}_{k+1}(b')$ agree on their $\mathrm{val}_k(a)$th bit. For, by k-equality (IH), $\mathfrak{A} \models \mathrm{eq}_{k,\ell+2}[a, b, \bar{c}, b', a]$. Hence, by (Φ_{19}) $\mathfrak{A} \models \mathrm{eqDig}_{k+1,\ell}[b, \bar{c}, b', a]$ holds just in case $\mathfrak{A} \models \eta_{k,\ell+2}[a, b, \bar{c}, b', a]$. But, by (Φ_{18}), $\mathfrak{A} \models \mathrm{in}_k^{\ell+2}[a, b, \bar{c}, b', a]$ if and only if $\mathfrak{A} \models \mathrm{in}_k[a, b]$, i.e. if and only if the $\mathrm{val}_k(a)$th bit of $\mathrm{val}_{k+1}(b)$ is 1. That is, $\mathfrak{A} \models \mathrm{eqDig}_{k+1,\ell}[b, \bar{c}, b', a]$ is equivalent to the statement that $\mathfrak{A} \models \mathrm{in}_k[a, b]$ if and only if $\mathfrak{A} \models \mathrm{out}_k[b', a]$. The situation is illustrated (for the case where $\mathfrak{A} \models \mathrm{eqDig}_{k+1,\ell}[b, \bar{c}, b', a]$ holds) in the right-hand diagram of Fig. 2, where all polarity alternatives $\pm$ are assumed to be resolved in the same way.

But, by $(k+1)$-harmony, $\mathfrak{A} \models \mathrm{out}_k[b', a]$ if and only if $\mathfrak{A} \models \mathrm{in}_k[a, b']$. This establishes the claim.

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

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

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 $\mathfrak{t}(m,n)$. We claim that $\Phi_{m,n}$ is satisfiable. Let $A = A_1 \dot\cup \cdots \dot\cup A_m$, where $A_k = \{\langle k,i\rangle \mid 0 \le i < \mathfrak{t}(k,n)\}\rangle$. (That is, A is the disjoint union of the various sets of integers $[0, \mathfrak{t}(k,n)-1]$.) Let $\text{int}_k^{\mathfrak{A}} = A_k$ for all k $(1 \le k \le m)$, and interpret the other predicates of $\Sigma_{m,n}$ as described. 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 $\ell = 0$ 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 $\max(5, 2m) \le 2m$ variables. If, however, $m = 2$, 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 $\ell = 0$, which feature only 4 variables. Either way, (Φ_7)–(Φ_{20}) are in $\mathcal{FL}^{2m}$. ◀

Now that we can enforce m-tuply exponentially large models in $\mathcal{FL}^{2m}$, it is a simple matter to establish that the satisfiability problem for this fragment is m-NExpTime-hard. The technique involves the encoding of tiling problems over a grid of m-tuply exponential size using formulas of $\mathcal{FL}^{2m}$, the coordinates of the various positions in this grid being represented as pairs of m-integers. The details are roughly analogous to the NExpTime-hardness proof for the two-variable fragment of first-order logic (see e.g. [2], pp. 253 ff.), and are relegated to the Appendix.

▶ **Theorem 2.** *The satisfiability problem for $\mathcal{FL}^{2m}$ is m-NExpTime-hard. Hence, the satisfiability problem for $\mathcal{FL}$ is non-elementary.*

4 Upper bound

In this section we show that $\mathcal{FL}$ is decidable, thus confirming the Purdy's original article [8]. The strategy adopted here, however, employs a variant of a construction found in the (as we now know) flawed article [10]. As well as providing m-NExpTime upper complexity bounds for the sub-fragments $\mathcal{FL}^m$ for all $m \ge 1$, the argument below yields a much shorter and more perspicuous proof than that given in [8].

The 1-variable fluted fragment, $\mathcal{FL}^1$, coincides with the 1-variable fragment of first-order logic, and so its satisfiability (= finite satisfiability) problem is in NPTime (and hence certainly in NExpTime). Furthermore, the 2-variable fluted fragment, $\mathcal{FL}^2$, is a proper subset of the 2-variable fragment of first-order logic, whose satisfiability (= finite satisfiability) problem is known to be in NExpTime [3] (and hence certainly in 2-NExpTime). In the sequel, therefore, we may confine attention to the case $m \ge 3$.

We begin by narrowing the range of fluted formulas we need to consider. An $\mathcal{FL}^m$-sentence Φ is in *normal form* if it is a conjunction of sentences of the forms:

$$\forall x_1 \forall x_2 \ldots \forall x_k(\theta \to \exists x_{k+1}\xi) \qquad \forall x_1 \forall x_2 \ldots \forall x_k(\vartheta \to \forall x_{k+1}\psi),$$

where $1 \le k < m$, and θ, ξ, ϑ, ψ are quantifier-free fluted formulas such that $\theta, \vartheta \in \mathcal{FL}^{[k]}$, and $\xi, \psi \in \mathcal{FL}^{[k+1]}$. The proof of the following lemma is routine:

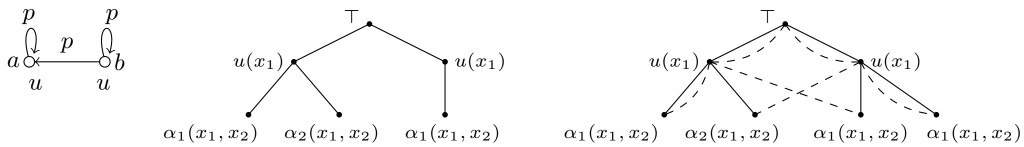

Figure 3 A structure $\mathfrak{A}$, its characteristic fluted (2,0)-constituent, $\mu = \mathrm{fc}_2^{\mathfrak{A}}[\,]$, depicted as a tree, and a double-tree formed from μ featuring a duplicated node.

▶ **Lemma 3.** *Let φ be an $\mathcal{FL}^m$-sentence over a signature Σ. We can compute, in exponential time, a disjunction $\Psi = \bigvee_{i \in I} \Psi_i$ of normal form $\mathcal{FL}^m$-sentences over a signature Σ' such that φ is satisfiable over a given domain A if and only if Ψ is satisfiable over A, $\|\Psi_i\| = O(\|\varphi\| \log\|\varphi\|)$ $(i \in I)$ and Σ' consists of Σ together with some additional predicates of arity at most $m-1$.*

We require one or two additional technical preliminaries. An *atomic fluted k-type* (over a given signature) is a maximal consistent conjunction of atomic or negated atomic $\mathcal{FL}^{[k]}$-formulas. For example, over a signature Σ featuring a single unary predicate u and a single binary predicate p, the formulas $\alpha_1(x_1, x_2) = u(x_2) \wedge p(x_1, x_2)$ and $\alpha_2(x_1, x_2) = u(x_2) \wedge \neg p(x_1, x_2)$ are atomic fluted 2-types. We take $\top$ to be the unique atomic fluted 0-type. Where the signature is clear from context, we denote the set of all atomic fluted k-types by $\boldsymbol{\alpha}^{(k)}$, and write $\boldsymbol{\alpha} = \bigcup_{k=0}^{m} \boldsymbol{\alpha}^{(k)}$. Observe that $|\boldsymbol{\alpha}^{(k)}|$ for each k is bounded by $2^{|\Sigma|}$, hence $|\boldsymbol{\alpha}| \leq m \cdot 2^{|\Sigma|}$. For a given Σ-structure $\mathfrak{A}$ and $\bar{a} \in A^k$, we denote by $\mathrm{tp}^{\mathfrak{A}}[\bar{a}]$ the unique atomic fluted k-type t such that $\mathfrak{A} \models t[\bar{a}]$.

We employ an adaptation, to the fluted case, of the familiar notion of Hintikka constituent (see, e.g. [14]). Again, fix some signature Σ. Define a *fluted (m, m)-constituent* to be an atomic fluted m-type. Let k satisfy $m > k \geq 0$. Define a *fluted (k, m)-constituent* to be a formula

$$\lambda = t(x_1, \ldots, x_k) \wedge \bigwedge_{\lambda' \in \Lambda} \exists x_{k+1}.\lambda' \wedge \forall x_{k+1} \bigvee_{\lambda' \in \Lambda} \lambda', \tag{6}$$

where Λ is some set of fluted $(k+1, m)$-constituents, and t an atomic fluted k-type. We call the elements of Λ the *successors* of λ, and t the *atomic part of* λ. Any fluted (k, m)-constituent λ has a natural representation as a labelled tree, with the children of each node given by its successors and the label by its atomic part.

We illustrate these notions for the signature Σ considered above, and for the value $m = 2$. The fluted $(2, 2)$-constituents are, by definition, the atomic fluted 2-types. Hence, the formulas $\lambda_1(x_1) = u(x_1) \wedge \exists x_2.\alpha_1(x_1, x_2) \wedge \exists x_2.\alpha_2(x_1, x_2) \wedge \forall x_2(\alpha_1(x_1, x_2) \vee \alpha_2(x_1, x_2))$ and $\lambda_2(x_1) = u(x_1) \wedge \exists x_2.\alpha_1(x_1, x_2) \wedge \forall x_2.\alpha_1(x_1, x_2)$ are fluted (1,2)-constituents. In the same way, the sentence $\mu = \exists x_1 \lambda_1(x_1) \wedge \exists x_1 \lambda_2(x_1) \wedge \forall x_1(\lambda_1(x_1) \vee \lambda_2(x_1))$ is a fluted (0,2)-constituent. The tree corresponding to this fluted (0,2)-constituent is depicted in the middle diagram of Fig. 3. By elementary combinatorics, one can show that the number of fluted (k, m)-constituents is bounded by $\mathfrak{t}(m - k + 1, n + m - k)$, where $n = |\Sigma|$.

The following lemma states that, for fixed k and m, the fluted (k, m)-constituents over a given signature form a partition. The proof proceeds in exactly the same way as for Hintikka constituents (cf. Theorem 2 in [8] and Theorem 3.10 in [14]).

▶ **Lemma 4.** *Fix some signature Σ and integers $m \geq k \geq 0$. Denote by Λ the set of fluted (k, m)-constituents over Σ. Then: (i) $\models \bigvee \Lambda$; and (ii) $\models \lambda \rightarrow \neg\mu$ for all distinct $\lambda, \mu \in \Lambda$.*

It follows that, in any structure $\mathfrak{A}$, and for any m, k ($m \geq k \geq 1$), a k-tuple $\bar{a}$ from A satisfies a unique fluted (k, m)-constituent. We denote this fluted (k, m)-constituent by $\mathrm{fc}^{\mathfrak{A}}_m[\bar{a}]$. In the case $k = 0$, we obtain the fluted $(0, m)$-constituent $\mathrm{fc}^{\mathfrak{A}}_m[\,]$, which we call the *characteristic* fluted $(0, m)$-constituent of $\mathfrak{A}$. The left-hand diagram in Fig. 3 shows a 2-element structure $\mathfrak{A}$ over domain $\{a, b\}$ whose characteristic fluted $(0, m)$-constituent is the sentence μ in our example above. Indeed, we see that $\mathrm{fc}^{\mathfrak{A}}_2[a] = \lambda_1$ and $\mathrm{fc}^{\mathfrak{A}}_2[b] = \lambda_2$.

Let λ be a fluted (k, m)-constituent. If $k > 0$, define $\varphi^{\leftarrow}$ to be the formula obtained from φ by deleting all literals of λ containing x_1 and then shifting remaining variables left, i.e. replacing each variable x_{i+1} by x_i, for $i > 1$. Thus, continuing the example of the previous paragraph, we have (removing unnecessarily duplicated conjuncts) $\lambda_1^{\leftarrow} = \exists x_1.u(x_1) \wedge \forall x_1.u(x_1)$. If, on the other hand, $k < m$, define $\lambda^{\uparrow}$ to be the formula obtained by removing from λ all literals containing x_m and all subformulas starting with a quantifier binding x_m. (If $k = m = 1$, we take $\lambda^{\leftarrow}$ to be $\top$; similarly, if $k = 0$ and $m = 1$, we take $\lambda^{\uparrow}$ to be $\top$.) The following Lemma is completely routine.

▶ **Lemma 5.** *Let λ be a fluted (k, m)-constituent. If $k > 0$, then $\lambda^{\leftarrow}$ is a fluted $(k-1, m-1)$-constituent; if, on the other hand, $k < m$, then $\lambda^{\uparrow}$ is a fluted $(k, m-1)$-constituent. Moreover, if $\lambda = \mathrm{fc}^{\mathfrak{A}}_m[\,]$ for some Σ-structure $\mathfrak{A}$ and λ' is a fluted $(1, m)$-constituent that is a successor of λ, then $\lambda'^{\leftarrow} = \lambda^{\uparrow}$.*

By way of motivation, suppose $\mathfrak{A}$ is a structure, with $\mathrm{fc}^{\mathfrak{A}}_m[a_1, \ldots, a_k] = \lambda$. If $k > 0$, then $\mathrm{fc}^{\mathfrak{A}}_{m-1}[a_2, \ldots, a_k] = \lambda^{\leftarrow}$; likewise, if $k < m$, then $\mathrm{fc}^{\mathfrak{A}}_{m-1}[a_1, \ldots, a_k] = \lambda^{\uparrow}$.

For the purposes of this section it is expedient to view a (finite, non-empty) tree as a tuple $T = \langle V, \varepsilon, p\rangle$, where $\varepsilon \in V$ and $p : (V \setminus \{\varepsilon\}) \to V$ is a function such that, for all $v \in V$, $p(\cdots p(v)) = \varepsilon$ for some finite (possibly zero) number of applications of p. The elements of V are the nodes of T, ε is the root, and, for every other node v, $p(v)$ is the parent of v. The number $h(v)$ of applications of p required to reach ε from v is the *height* of v; thus, $h(\varepsilon) = 0$. As usual, for a subset $X \subseteq V$, $p[X]$ denotes the image of X, and $p^{-1}[X]$, the inverse image of X under p. We write $p^{-1}[v]$ for $p^{-1}[\{v\}]$ – that is, the children of v. As usual, a *leaf* of a tree is a node v with no children, and we say that a tree is *uniform* if all leaves have the same height. The *height* of a uniform tree is the height of any of its leaves.

We now define the main notions for this section. Let $T_f = \langle V, \varepsilon, f\rangle$ and $T_g = \langle V, \varepsilon, g\rangle$ be uniform trees of height m sharing the same nodes and root. We say $\mathbb{T} = \langle V, \varepsilon, f, g\rangle$ is a *double tree* if the following hold:

$$\text{for every } v \in V,\ f^{-1}[\varepsilon] = g^{-1}[\varepsilon]; \tag{D1}$$

$$\text{for every } v \in V \setminus \{\varepsilon\}, \text{ if } f^{-1}[v] \neq \emptyset, \text{ then } g[f^{-1}[v]] = f^{-1}[g(v)]. \tag{D2}$$

A routine induction shows that, if $\mathbb{T} = \langle V, \varepsilon, f, g\rangle$ is a double tree, then, for all $v \in V$, v has the same height in the trees $\langle V, \varepsilon, f\rangle$ and $\langle V, \varepsilon, g\rangle$. We call this number the *height* of v in $\mathbb{T}$, and denote it by $h(v)$. Note that (D2) is, in essence, a *confluence condition*: if m is the height of $\mathbb{T}$, then, for any node v (with $1 \leq h(v) < m$) the children of v in T_f map under g onto the children of $u = g(v)$ in T_f (see Fig. 4).

A moment's thought shows that not all trees can be made into double trees. This is true, for example, of the tree depicted in the middle diagram of Fig. 3: the right-hand node labelled $u(x_1)$ has exactly one child and one proper sister, which makes it impossible to satisfy (D2). On the other hand, by duplicating nodes if necessary, double trees can be created, as shown in the right-hand diagram of Fig. 3, where the function g is indicated by dashed edges. This process of transforming trees (specifically, trees depicting characteristic fluted constituents of structures) into double trees plays a crucial role in the proof of Lemma 7, below.

Let $\mathbb{T} = \langle V, \varepsilon, f, g \rangle$ be a double tree and $\tau : V \to \boldsymbol{\alpha}$ a function. We say $\mathbb{T} = \langle V, \varepsilon, f, g, \tau \rangle$ is an $\boldsymbol{\alpha}$-*double tree* if the following conditions hold:

$$\text{for every } v \in V,\ \tau(v) \in \boldsymbol{\alpha}^{(k)}, \text{ where } k = h(v); \tag{D3}$$

$$\text{for every } v \in V \setminus \{\varepsilon\},\ \tau(g(v)) = \tau(v)^{\leftarrow}. \tag{D4}$$

Thus, in an $\boldsymbol{\alpha}$-*double tree*, every node v is labelled with some atomic fluted $h(v)$-type, α. Notice that, if $h(v) \geq 1$, then $g(v)$ must be labelled with $\alpha^{\leftarrow}$.

Before we actually construct any $\boldsymbol{\alpha}$-double trees, let us first see what we can do with them. Let $\varphi \in \mathcal{FL}^m$ be a sentence in normal form, and $\mathbb{T} = \langle V, \varepsilon, f, g, \tau \rangle$ be an $\boldsymbol{\alpha}$-double tree. We say $\mathbb{T}$ *satisfies* φ, and write $\mathbb{T} \models \varphi$, when the following conditions hold for all $k < m$ and for every $v \in V$ with $h(v) = k$: (i) for every conjunct of φ of the form $\forall x_1 \forall x_2 \ldots \forall x_k(\theta \to \exists x_{k+1}\xi)$, if $\tau(v) \models \theta$ then there is a $v' \in f^{-1}[v]$ such that $\tau(v') \models \xi$; and (ii) for every conjunct of φ of the form $\forall x_1 \forall x_2 \ldots \forall x_k(\vartheta \to \forall x_{k+1}\psi)$, if $\tau(v) \models \vartheta$ then for every $v' \in f^{-1}[v]$, $\tau(v') \models \psi$. Thus, any $\boldsymbol{\alpha}$-double tree of height m endows any normal form formula of $\mathcal{FL}^m$ with a truth-value in a natural way. Observe that, in this definition, a conjunct of φ with variables $x_1, \ldots, x_{k+1}$ imposes constraints *only* on the labels of nodes with height k and $k+1$. However, conditions (D1)–(D4) on $\boldsymbol{\alpha}$-double trees ensure that these constraints affect adjacent pairs of nodes throughout $\mathbb{T}$.

There is a close – though subtle – relationship between $\boldsymbol{\alpha}$-double trees and models of normal-form $\mathcal{FL}$-formulas, encapsulated in the following two lemmas.

▶ **Lemma 6.** *Let $\varphi \in \mathcal{FL}^m$ be in normal form, and suppose $\mathbb{T}$ is an $\boldsymbol{\alpha}$-double tree of height m, s.t. $\mathbb{T} \models \varphi$. Then there is a model $\mathfrak{A} \models \varphi$, with $|A|$ equal to the number of leaves in $\mathbb{T}$.*

Proof. Let A be the set of leaves of $\mathbb{T}$. We assign to each node v of $\mathbb{T}$ a subset $A_v \subseteq A$ with the property that, for each non-leaf node v, the family $\{A_w \mid w \text{ a child of } v\}$ partitions A into non-empty subsets, making crucial use of the properties (D1) and (D2) of double trees. In this way, any k-tuple $a_1, \ldots a_k$ from A is naturally associated with a node v of height k, namely, that node v with ancestors $\varepsilon, v_1, \ldots, v_k = v$ where, for all i $(1 \leq i \leq k)$, $a_i \in A_{v_i}$. This assignment of sequences to nodes allows us to turn A into a structure $\mathfrak{A}$: if $a_1, \ldots a_k$ is associated with the node v, we set $\mathrm{tp}^{\mathfrak{A}}_m[a_1, \ldots, a_k]$ to be whatever atomic fluted k-type $\mathbb{T}$ labels v with. Using the properties (D3)–(D4), this assignment can be shown to be consistent, and to result in a structure satisfying φ. The details are given below.

Let $\mathbb{T} = \langle V, \varepsilon, f, g, \tau \rangle$ be a finite $\boldsymbol{\alpha}$-double tree satisfying φ. We decompose V as a union of disjoint subsets $V = \bigcup_{k=0}^{m} V^k$, where V^k $(0 \leq k \leq h)$ is the set of nodes of height k. Let A be the set of leaves of $\mathbb{T}$ – this will be the domain of the model $\mathfrak{A}$ we are going to construct. First we define two labelling functions $d_0, d : V \mapsto \mathcal{P}(A)$ satisfying the following properties for each k $(0 \leq k \leq m)$:

(i) for every $w \in V_k$, $d_0(w) \neq \emptyset$ and $d_0(w) \subseteq d(w)$,
(ii) if $0 < k$ then, for every $w \in V^k$, $d(w) \subseteq d(g(w))$,
(iii) $d(\varepsilon) = A$, and if $0 < k$, then, for every $w \in V^k$, the family $\{d(w') \mid w' \in f^{-1}[f(w)]\}$ is a partition of A.

We call $d(w)$ the *local domain* of w and $d_0(w)$ the *initial* local domain of w. We remark that, in (iii), $f^{-1}[f(w)]$ is the set of (reflexive) siblings of w in T_f.

For each element $w \in V$, we define $d_0(w)$ to be the set of leaves in the subtree of T_g rooted at w. Thus for any w with $h(w) \geq 1$,

$$d_0(w) \subseteq d_0(g(w)). \tag{7}$$

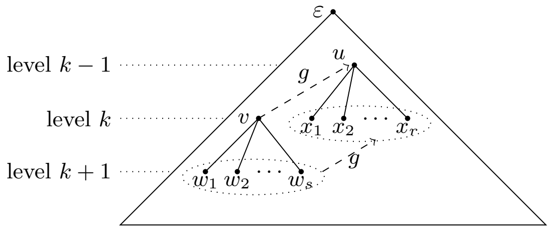

Figure 4 A double tree: $v \in V^k$, $w_1, \ldots, w_s$ are f-children of v; $x_1, \ldots, x_r$ are f-children of $u = g(v)$ and $g[\{w_1, \ldots, w_s\}] = \{x_1, \ldots, x_r\}$.

Now define $d(\varepsilon) = d_0(\varepsilon) = A$ and for each $w \in V^1$, define $d(w) = d_0(w)$. This ensures properties (i)–(iii) for $k = 0$ and $k = 1$. The function d is defined for remaining nodes by induction on the height. Assume properties (i)–(iii) hold for all nodes $v \in V^k$, where $1 \leq k < m$; we extend d to nodes $w \in V^{k+1}$ as follows. Consider first any element $v \in V^k$: we proceed to define $d(w)$ for every $w \in f^{-1}[v]$. Remembering that $k \geq 1$, let $u = g(v)$. Condition (D2) in the definition of double trees tells us that $g(f^{-1}[v]) = f^{-1}(u)$; we may therefore consider the various nodes in $x \in f^{-1}(u)$ one by one, defining $d(w)$ for all those $w \in f^{-1}[v]$ such that $g(w) = x$. Suppose $g^{-1}[x] = \{w_1, \ldots, w_l\}$. We define $d(w)$ for each of these nodes w by starting with the set $d_0(w)$, and then distributing the elements of $d(x) \setminus (d_0(w_1) \cup \ldots \cup d_0(w_l))$ among these sets in any way. Property (i) is thus secured trivially for all these w. From (7), we see that, for each such w, $d(w) \subseteq d(x) = d(g(w))$, thus securing (ii). By executing this procedure for all $x \in f^{-1}[u]$, we will have defined $d(w)$ for all $w \in f^{-1}[v]$. From property (iii) applied to level k, the family $\{d(x) \mid x \in f^{-1}[u]\}$ is a partition of A, whence the family $\{d(w) \mid w \in f^{-1}[v]\}$ will be a partition of A, thus securing property (iii) for all children of v in T_f. Therefore, after considering every $v \in V$ of height k, we obtain properties (i)–(iii) for all elements of V of height $k+1$.

We use one more piece of notation: for every $v \in V$ define S_v, the *sequence domain of* v, to be the set of k-tuples $S_v = d(v_1) \times \ldots \times d(v_k)$, where $\varepsilon, v_1, \ldots, v_k = v$ is the path in T_f from the root to v. Simple induction using property (iii) of the construction shows that sequence domains at each level k $(0 < k \leq m)$ form a partition of A^k:

(iv) $\bigcup_{v \in V^k} S_v = A^k$ and, if for some $\bar{a} \in A^k$ we have $\bar{a} \in S_v$ and $\bar{a} \in S_w$ then $v = w$.

Now we set the interpretation of the predicate letters on A. Namely, for every $p \in \Sigma$ of arity k, for every $v \in V$ with $h(v) \geq k$, and for every $\bar{a}\bar{b} \in S_v$ with $|\bar{b}| = k$, define:

$$\bar{b} \in p^{\mathfrak{A}} \quad \text{iff} \quad p(x_{l+1}, \ldots, x_{l+k}) \in \tau(v).$$

This is well defined. For, suppose in addition that $\bar{c}\bar{b} \in S_w$ for some $w \in V$ with $h(w) \geq k$. Let u and x be the ancestors of, respectively, v and w in T_g at level k. By property (ii) of the d-labelling, $\bar{b} \in S_u$ and $\bar{b} \in S_x$. Now, property (iv) implies $u = x$. And so applying repeatedly condition (D4) of the definition of an $\boldsymbol{\alpha}$-double tree to $\tau(v)$ and $\tau(w)$ we get:

$$p(x_{l+1}, \ldots, x_{l+k}) \in \tau(v) \quad \text{iff} \quad p(x_1, \ldots, x_k) \in \tau(u) \quad \text{iff} \quad p(x_{j+1}, \ldots, x_{j+k}) \in \tau(w).$$

It remains to check that $\mathfrak{A} \models \varphi$. Observe that for each $v \in V$ with $v \neq \varepsilon$ S_v is non-empty, and moreover for each $\bar{a} \in S_v$ $tp^{\mathfrak{A}}(\bar{a}) = \tau(v)$. Since φ is in normal form and $\mathbb{T}$ satisfies φ, it is obvious that all conjuncts of φ are true in $\mathfrak{A}$. ◀

▶ **Lemma 7.** *Let $\varphi \in \mathcal{FL}^m$ be in normal form, and suppose $\mathfrak{A} \models \varphi$. Then there exists a finite $\boldsymbol{\alpha}$-double tree satisfying φ with the number of leaves bounded by $\mathfrak{t}(m, O(mn))$.*

Proof. We begin by considering the characteristic fluted $(m, 0)$-constituent $\mu = \mathrm{fc}^{\mathfrak{A}}_m[\,]$. As we have seen, μ can be viewed as a tree $T_\lambda = \langle V, f, \varepsilon \rangle$. Any node v at level k $(0 \leq k \leq m)$ in T_λ corresponds to a position in the syntax tree of μ defining a fluted (k, m)-constituent v_λ. Note that we may have $\lambda_v = \lambda_w$ for distinct nodes v, w, because there may be repeated subformulas. Our task is to make T_λ into an $\boldsymbol{\alpha}$-double tree satisfying φ. We begin by adding a second parent function g to T_λ satisfying

$$\lambda^{\uparrow}_{g(v)} = \lambda^{\leftarrow}_v \text{ for all } v \in V \setminus \{\varepsilon\}, \tag{8}$$

which we now proceed to define.

We start by setting, for every v with height 1, $g(v) = f(v) = \varepsilon$. Note that Lemma 5 implies that for all nodes v with height 1, $\lambda^{\leftarrow}_v = \lambda^{\uparrow}_\varepsilon$, as required by (8). Now suppose that $g(v) = u$ has been defined for some node v of height $k < m$ in such a way that (8) holds, and suppose that w is a child of v, that is: $f(w) = v$. Thus, λ_w is simply a subformula of λ_v, whence, by (8), $\lambda_{g(v)}$ must have some subformula π such that $\pi^{\uparrow} = \lambda^{\leftarrow}_w$. Indeed, π is a fluted (k, m)-constituent. Let x be the child of u (i.e. $f(x) = u$) such that $\lambda_x = \pi$. Thus, $\lambda^{\leftarrow}_w = \lambda^{\uparrow}_x$, and we may set $g(w) = x$. Proceeding in this way, we can define g for all of $V \setminus \{\varepsilon\}$ satisfying (8). The resulting construction is illustrated in Fig. 4. Unfortunately, this construction does not *quite* ensure (D2); for, while g maps the children (under f) of v to the children (under f) of $g(v)$, the latter set may not be covered by this mapping. That is, we have $g[f^{-1}[v]] \subseteq f^{-1}[g(v)]$ rather than the desired $g[f^{-1}[v]] = f^{-1}[g(v)]$. However, this matter can be rectified by creating duplicates of the children of v which are then free to be mapped to any children of $g(v)$ not yet accounted for (c.f. Fig. 3).

To make this construction precise we need one more notion. Let k satisfy $0 < k < m$. A fluted (k, m)-*semi-constituent* is defined in the same way as a fluted (k, m)-constituent via the recursion in (6), except that Λ is now a finite *multiset* (rather than *set*) of fluted $(k + 1, m)$-semi-constituents. In other words, fluted semi-constituents are just like fluted constituents, with the difference that repeated successors are allowed. Let η and ζ be fluted (k, m)-semi-constituents. We write $\eta \approx \zeta$ if, after removing repeated successors from η and ζ (at all levels in the recursion), the same fluted (k, m)-constituent is obtained. We represent fluted semi-constituents as trees, just as we do fluted constituents; moreover, the operations $^{\leftarrow}$ and $\uparrow$ on fluted constituents are extended to fluted semi-constituents in the obvious way.

Now, we are ready to give the details. To define the function g in T_λ, we start as mentioned above by setting, for every $w \in V^1$, $g(w) = f(w) = \varepsilon$. For nodes $w \in V$ with $h(w) > 1$, we define $g(w)$ assuming that $g(f(w))$ has already been defined and that the following condition holds for $v = f(w)$

$$\lambda^{\uparrow}_{g(v)} \approx \lambda^{\leftarrow}_v. \tag{9}$$

The basic step, for a single node w, is performed as follows.

Define $g(w)$: suppose $v = f(w)$ and (9) holds for v. Since λ_w is a successor of λ_v there exists at least one conjunct η of $\lambda_{g(v)}$ such that $\eta^{\uparrow} \approx \lambda^{\leftarrow}_w$. So we may pick a node x such that $\lambda_x = \eta$ and we set $g(w) = x$.

Observe that after defining $g(w)$ as above we have $\lambda^{\uparrow}_{g(w)} \approx \lambda^{\leftarrow}_w$. So, we may extend the definition of g to a complete f-subtree $T_f(w)$ of w, traversing the subtree level by level. This is performed as follows.

Complete$(g, T_f(w))$: suppose $v = f(w)$ and (9) holds for v.
Let $z_1 \ldots, z_r$ be all the nodes of $T_f(w)$, excluding the root, ordered so that nodes

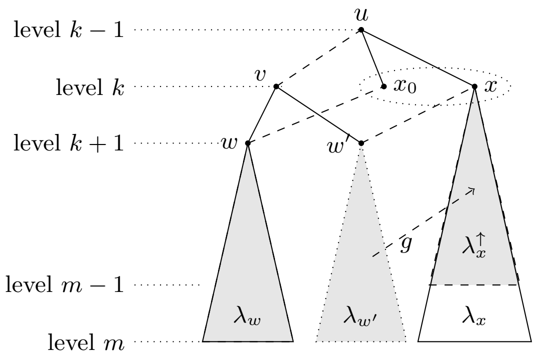

Figure 5 Covering x with respect to v: solid lines correspond to f and dashed lines – to g. $\lambda_w^{\leftarrow} \approx \lambda_x^{\uparrow}$ and $\lambda_{w'}$ is a copy of λ_w.

on lower levels appear before nodes on higher levels (e.g. according to breadth-first search). For every i ($0 \le i \le r$) call *Define*$(g(z_i))$.

Now, we are ready to extend the definition of g to all nodes of V by calling the above procedure for every node $w \in V^1$. Denote the resulting tree by $T = \langle V, \varepsilon, f, g, \tau\rangle$. Evidently, in T the following properties hold for every k ($0 \le k \le m$)

(i) if $k > 0$, then for all $w \in V^k$, $\lambda^{\uparrow}_{g(w)} \approx \lambda^{\leftarrow}_w$,

(ii) if $0 < k < m$ then, for every $w \in V^k$, $g[f^{-1}[w]] \subseteq f^{-1}[g(w)]$.

In the required $\boldsymbol{\alpha}$-double tree to ensure (D2) the inclusions in condition (ii) are supposed to become equalities. This requires one more operation.

Suppose x and v are two nodes such that $g(v) = f(x)$ (refer to Figure 5). Suppose there is no node w such that $f(w) = v$ and $g(w) = x$. We then proceed as follows.

Covering x with respect to v: suppose $x \in V$, $1 \le h(x) < m$, $g(v) = f(x) = u$, and $g^{-1}[x] = \emptyset$, where $h(x)$ is the height of x in the f-tree. By (i), $\lambda^{\uparrow}_u \approx \lambda^{\leftarrow}_v$. So, there is a conjunct η in λ_v such that $\eta^{\leftarrow} \approx \lambda^{\uparrow}_x$. Pick a node w such that $f(w) = v$ and $\lambda_w = \eta$. Add to v a copy $T_{w'}$ of the f-tree T_w. (We remind the reader that this is equivalent to add to λ_v a copy of the conjunct λ_w.) Define $g(w') = x$ and run the procedure *Complete*$(g, T_f(w'))$.

We note that the above operation maintains conditions (i)–(ii).

Now, we proceed on all nodes of the tree constructed, level by level, covering all the nodes x such that $h(x) < m$ with respect to all nodes v such that $g(v) = f(x)$. Denote the resulting tree by $\mathbb{T}$ and the set of nodes of $\mathbb{T}$ by $\mathbb{V}$.

Evidently, $\mathbb{T}$ is an $\boldsymbol{\alpha}$-double tree. Moreover, since we have added only identical conjuncts to T, $\mathbb{T}$ also satisfies φ.

It remains to estimate the number of leaves in $\mathbb{T}$. Let $C(k,m)$ be the number of fluted (k,m)-constituents and let $N(k)$ be the maximal number of children of a node $v \in \mathbb{V}$ with $h(v) = k$. By elementary combinatorics, one can show that $C(k,m) \le \mathfrak{t}(m-k+1, n+m-k)$, where $n = |\Sigma|$. We have $N(0) = C(1,m)$, $N(1) \le C(2,m) + C(1,m)$, and $N(m-1) \le C(m,m) + \ldots + C(1,m) \le m \cdot C(1,m)$. Hence, the number of leaves in $\mathbb{T}$ is bounded by $(m \cdot C(m-1))^m$, which is bounded by $\mathfrak{t}(m, O(mn))$. ◀

Lemmas 3, 6 and 7 instantly imply the main theorem of this section:

▶ **Theorem 8.** *$\mathcal{FL}$ has the finite model property. Moreover, the satisfiability problem for $\mathcal{FL}^m$ is in m-NExpTime.*

Discussion

When restricting attention to $\mathcal{FL}$ with a fixed number of variables, the upper complexity bounds given in Theorem 8 are not tight, even for $\mathcal{FL}^1$. As mentioned earlier, the 1-variable fluted fragment, $\mathcal{FL}^1$, coincides with the 1-variable fragment of first-order logic, and so its satisfiability problem is NPTime-complete (that is: 0-NExpTime-complete). Furthermore, the 2-variable fluted fragment, $\mathcal{FL}^2$, is a proper subset of the 2-variable fragment of first-order logic, whose satisfiability (= finite satisfiability) problem is known to be in NExpTime. Hence, setting $m = 1$ in Theorem 2, the satisfiability problem for $\mathcal{FL}^2$ is NExpTime-complete. By considering more closely the normal forms yielded by Lemma 3 and using a more complicated construction, the present authors have been able to strengthen Theorem 8 to show that, for $m \geq 3$, $\mathcal{FL}^m$ is in fact in $(m-2)$-NExpTime. Together with Theorem 2, this implies that $\mathcal{FL}^3$ is NExpTime-complete and $\mathcal{FL}^4$ is 2-NExpTime-complete. However, the additional gain in the upper complexity-bound is purchased at the cost of a less perspicuous proof, and anyway fails to close the gap with the lower complexity-bounds for $\mathcal{FL}^m$ when $m \geq 5$.

References

1 H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.
2 E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem.* Springer, 1997.
3 E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
4 U. Hustadt, R. Schmidt, and L. Georgieva. A survey of decidable first-order fragments and description logics. *Journal of Relational Methods in Computer Science*, 1(3):251–276, 2004.
5 C. Lutz and U. Sattler. The complexity of reasoning with Boolean modal logics. In F. Wolter, H. Wansing, M. de Rijke, and M. Zakharyaschev, editors, *Advances in Modal Logic*, volume 3, pages 1–20, Menlo Park, 2001. CLSI Publications.
6 A. Noah. Predicate-functors and the limits of decidability in logic. *Notre Dame Journal of Formal Logic*, 21(4):701–707, 1980.
7 W. C. Purdy. Decidability of fluted logic with indentity. *Notre Dame Journal of Formal Logic*, 37(1):84–104, 1996.
8 W. C. Purdy. Fluted formulas and the limits of decidability. *Journal of Symbolic Logic*, 61(2):608–620, 1996.
9 W. C. Purdy. Quine's limits of decision. *Journal of Symbolic Logic*, 64(4):1439–1466, 1999.
10 W. C. Purdy. Complexity and nicety of fluted logic. *Studia Logica*, 71:177–198, 2002.
11 W. V. Quine. On the limits of decision. In *Proceedings of the 14th International Congress of Philosophy*, volume III, pages 57–62. University of Vienna, 1969.
12 W. V. Quine. Algebraic logic and predicate functors. In *The Ways of Paradox*, pages 283–307. Harvard University Press, revised and enlarged edition, 1976.
13 W. V. Quine. The variable. In *The Ways of Paradox*, pages 272–282. Harvard University Press, revised and enlarged edition, 1976.
14 V. Rantala. Constituents. In R. Bogdan, editor, *Jaakko Hintikka*, volume 8 of *Profiles*, pages 43–76. Springer Netherlands, 1987.
15 R. Schmidt and U. Hustadt. A resolution decision procedure for fluted logic. In D. McAllester, editor, *Proceedings, CADE*, number 1831 in LNAI, pages 433–448. Springer-Verlag, 2000.

5 Appendix

Proof of Theorem 2. If $m = 1$, the result may be obtained by simple adaptation of the familiar proof that the satisfiability problem for the two-variable (non-fluted) fragment of first-order logic is NExpTime-hard [2, pp. 253, ff.]. We therefore assume in the sequel that $m \geq 2$. This avoids a tedious special case.

We employ the apparatus of tiling systems. A *tiling system* is a triple $\langle C, H, V\rangle$, where C is a non-empty, finite set and H, V are binary relations on C. The elements of C are referred to as *colours*, and the relations H and V as the *horizontal* and *vertical* constraints, respectively. For any integer N, a *tiling* for $\langle C, H, V\rangle$ of size N is a function $f : \{0, \ldots, N-1\}^2 \to C$ such that, for all i, j $(0 \leq i, j < N)$, the pair $\langle f(i, j), f(i+1, j)\rangle$ is in H and the pair $\langle f(i, j), f(i, j+1)\rangle$ is in V, with addition in arguments taken to be modulo N. A tiling of size N is to be pictured as a colouring of an $N \times N$ toroidal grid by the colours in C; the horizontal constraints H thus specify which colours may appear 'to the right of' which other colours; the vertical constraints V likewise specify which colours may appear 'above' which other colours. An n-tuple $\bar{c}$ of elements of C is an *initial configuration* for the tiling f if $\bar{c} = f(0, 0), \ldots, f(n-1, 0)$. An initial configuration for f is to be pictured as a row of n colours occupying the 'bottom left-hand' corner of the grid.

The *m-tuply exponential tiling problem* for a tiling system (C, H, V) is the following problem: given an n-tuple $\bar{c}$ from C, determine whether there exists a tiling for (C, H, V) of size $\mathfrak{t}(m, n)$ with initial configuration $\bar{c}$. Because of the close connection between runs of Turing machines and solutions of tiling systems, it is straightforward to see that there exist tiling systems for which the m-tuply exponential tiling problem is m-NExpTime-complete.

Let (C, H, V) be a tiling system and $m \geq 2$. We construct, for any n-tuple $\bar{c}$ from C, a formula $\Phi_{\bar{c}}$ over a signature $\Sigma_{\bar{c}}$, with the property that there exists a tiling for (C, H, V) of size $\mathfrak{t}(m, n)$ with initial configuration $\bar{c}$ if and only if $\Phi_{\bar{c}}$ is satisfiable. For ease of reading, we add predicates to $\Sigma_{\bar{c}}$ as they are encountered, and we specify the conjuncts of $\Phi_{\bar{c}}$ as and when they are required in the course of the reduction.

We begin by setting $\Sigma_{\bar{c}}$ to be the signature $\Sigma_{m,n}$ from Theorem 1, and we add to $\Phi_{\bar{c}}$ all the conjuncts of $\Phi_{m,n}$. In fact, it is the predicates established in the *penultimate* stage of the induction that interest us here. Specifically, in any model of $\Phi_{\bar{c}}$, we have a valuation function val_{m-1} defined on $(m-1)$-integers with values in the range $[1, \mathfrak{t}(m-1, n)]$, and satisfying the properties $(m-1)$-harmony, $(m-1)$-zero, $(m-1)$-predecessor, $(m-1)$-covering and $(m-1)$-equality. Using these predicates, we construct objects that may be thought of as, in essence, *pairs* of m-integers.

We start by adding to $\Sigma_{\bar{c}}$ the unary predicate vtx and the binary predicates in_X and in_Y. If a structure $\mathfrak{A}$ is clear from context, we call any element of A satisfying vtx a *vertex*. Define the function $\mathrm{val}_X : \mathrm{vtx}^{\mathfrak{A}} \to [0, \mathfrak{t}(m, n) - 1]$ by setting $\mathrm{val}_X(b)$, for any vertex b, to be the integer with canonical representation $s_{N-1}, \ldots, s_0$ of length $N = \mathfrak{t}(m-1, n)$ where, for all i $(0 \leq i < N)$,

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

Think of $\mathrm{val}_X(b)$ as the *horizontal coordinate* of b. The *vertical coordinate* of b, $\mathrm{val}_Y(b)$ is defined similarly, using the binary predicates in_Y. We rely on $(m-1)$-covering to ensure that, for any i in the range $[0, \mathfrak{t}(m-1, n) - 1]$, there *is* an $(m-1)$-integer having any value

i; and below we establish harmony-like properties showing that it does not matter *which* such $(m-1)$-integer we choose, if there are many.

Add to $\Sigma_{\bar{c}}$ the binary predicates out_X, out_Y, eq_X, eq_Y pred_X and pred_Y. Using D to stand for either of the letters X or Y, add to $\Phi_{\bar{c}}$ conjuncts ensuring the following properties

D-harmony: For all vertices b and all $(m-1)$-integers a, a' in $\mathfrak{A}$ such that $\text{val}_{m-1}(a) = \text{val}_{m-1}(a')$, $\mathfrak{A} \models \text{in}_D[a,b] \Leftrightarrow \mathfrak{A} \models \text{out}_D[a',b]$.

D-zero: For all vertices b in $\mathfrak{A}$, $\mathfrak{A} \models \text{zero}_D[a] \Leftrightarrow \text{val}_D(b) = 0$.

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 $\mathfrak{A}$, $\mathfrak{A} \models \text{pred}_D[b,b'] \Leftrightarrow \text{val}_D(b') = \text{val}_D(b) - 1$ modulo $\mathfrak{t}(m,n)$.

The conjuncts in question are trivial adaptations of those used in the proof of Theorem 1 to establish m-harmony m-zero, m-equality and m-predecessor. All require at most $2m$ variables.

The following conjuncts of $\Phi_{\bar{c}}$ now establish that, for all pairs of integers i, j in the range $[0, \mathfrak{t}(m,n)-1]$, there exists a vertex a with coordinates (i,j):

$$\exists x_1(\text{vtx}(x_1) \wedge \text{zero}_X(x_1) \wedge \text{zero}_Y(x_1))$$
$$\forall x_1(\text{vtx}(x_1) \to \exists x_2(\text{vtx}(x_2) \wedge \text{pred}_Y(x_1,x_2) \wedge \text{eq}_X(x_1,x_2)))$$
$$\forall x_1(\text{vtx}(x_1) \to \exists x_2(\text{vtx}(x_2) \wedge \text{pred}_X(x_1,x_2) \wedge \text{eq}_Y(x_1,x_2))).$$

Note that there is no requirement that vertices be uniquely defined by their horizontal and vertical coordinates.

Treating each colour in C as a unary predicate in $\Sigma_{\bar{c}}$, we colour each vertex uniquely by adding to $\Phi_{\bar{c}}$ the conjunct

$$\forall x_1(\text{vtx}(x_1) \to \left(\bigvee_{c \in C} c(x_1) \right) \wedge \bigwedge_{c,d \in C}^{c \neq d} \neg(c(x_1) \wedge d(x_1))).$$

To obtain a well-defined grid-colouring, we wish models of $\Phi_{\bar{c}}$ to satisfy the following 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 $c \in C$, $\mathfrak{A} \models c[b] \Leftrightarrow \mathfrak{A} \models c[b']$.

And this we ensure by adding to $\Phi_{\bar{c}}$ the conjunct

$$\forall x_1(\text{vtx}(x_1) \wedge c(x_1) \to \forall x_2(\text{vtx}(x_2) \wedge \text{eq}_X(x_1,x_2) \wedge \text{eq}_Y(x_1,x_2) \to c(x_2))).$$

Thus, any model of $\Phi_{\bar{c}}$ defines a colouring of the $\mathfrak{t}(m,n) \times \mathfrak{t}(m,n)$ toroidal grid. To ensure that this colouring is a tiling for (C,H,V), we simply add to $\Phi_{\bar{c}}$ the conjuncts

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

Finally, we need to ensure that $\bar{c}$ is written in the bottom left configuration. But this is routine. If c_0 is the initial element of $\bar{c}$, we write

$$\exists x_1(\text{vtx}(x_1) \wedge \text{zero}_X(x_1) \wedge \text{zero}_Y(x_1) \wedge c_0(x_1)).$$

thus setting the 'bottom left' tile to have colour c_0. The remaining elements of $\bar{c}$ may be easily set using the predicates predDig_X and eq_Y. This completes the construction of $\Phi_{\bar{c}}$.

We have shown that, if $\Phi_{\bar{c}}$ is satisfiable, then the $\mathfrak{t}(m,n) \times \mathfrak{t}(m,n)$-grid colouring problem (C, H, V) has a solution. Conversely, a simple check shows that, if $\mathfrak{t}(m,n) \times \mathfrak{t}(m,n)$-grid colouring problem (C, H, V) has a solution, then by interpreting the predicates involved in $\Phi_{\bar{c}}$ as suggested above, we obtain a model of $\Phi_{\bar{c}}$. This completes the reduction. It is evident that, for fixed (C, H, V) and m, this reduction runs in time polynomially bounded by the length of $\bar{c}$. ◀