Twin-width IV: ordered graphs and matrices

Édouard Bonnet
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Ugo Giocanti
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Patrice Ossona de Mendez
Centre d'Analyse et de Mathématique Sociales CNRS UMR 8557, France
and Computer Science Institute of Charles University (IUUK), Praha, Czech Republic

Pierre Simon
UC Berkeley

Stéphan Thomassé
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Szymon Toruńczyk
University of Warsaw

Abstract

We establish a list of characterizations of bounded twin-width for hereditary, totally ordered binary structures. This has several consequences. First, it allows us to show that a (hereditary) class of matrices over a finite alphabet either contains at least $n!$ matrices of size $n \times n$, or at most c^n for some constant c. This generalizes the celebrated Stanley-Wilf conjecture/Marcus-Tardos theorem from permutation classes to any matrix class over a finite alphabet, answers our small conjecture [SODA '21] in the case of ordered graphs, and with more work, settles a question first asked by Balogh, Bollobás, and Morris [Eur. J. Comb. '06] on the growth of hereditary classes of ordered graphs. Second, it gives a fixed-parameter approximation algorithm for twin-width on ordered graphs. Third, it yields a full classification of fixed-parameter tractable first-order model checking on hereditary classes of ordered binary structures. Fourth, it provides a model-theoretic characterization of classes with bounded twin-width.

2012 ACM Subject Classification Mathematics of computing → Discrete mathematics → Combinatorics → Enumeration; Theory of computation → Logic → Finite Model Theory; Theory of computation → Design and analysis of algorithms → Parameterized complexity and exact algorithms

Keywords and phrases Twin-width, matrices, ordered graphs, enumerative combinatorics, model theory, algorithms, computational complexity, Ramsey theory

Funding This paper is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 810115 – DYNASNET).

Acknowledgments We thank Colin Geniet, Eunjung Kim, Jarik Nešetřil, Sebastian Siebertz, Michał Pilipczuk, and Rémi Watrigant for fruitful discussions.

1 Introduction

A common goal in combinatorics, structural graph theory, parameterized complexity, and finite model theory is to understand *tractable* classes of graphs, or other structures. Depending on the context, tractability may mean e.g. few structures of any given size, small chromatic number, efficient algorithms, for example for the clique problem; or structural properties of sets definable by logical formulas.

Twin-width is a recently introduced graph width parameter [8, 7, 6] which, despite its generality, ensures many of those properties. Intuitively, a graph has twin-width d if it can be constructed by merging larger and larger parts so that at any moment, every part has a non-trivial interaction with at most d other parts (two parts have a trivial interaction if either no edges, or all edges span across the two parts).

Many well-studied graph classes have bounded twin-width: planar graphs, and more generally, any class of graphs excluding a fixed minor; cographs, and more generally, any class of bounded clique-width.

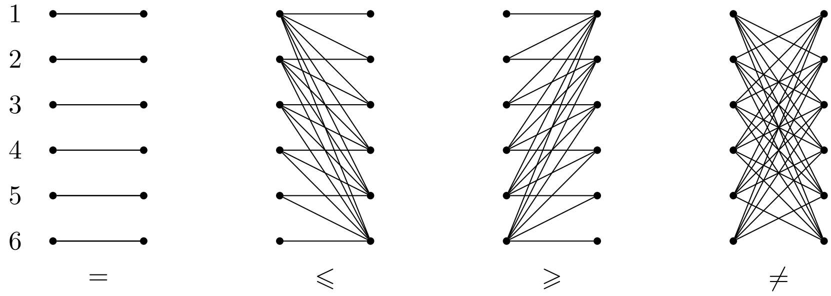

Figure 1 The four bipartite graphs $G^n_=$, $G^n_\leqslant$, $G^n_\geqslant$, and $G^n_\neq$, for $n = 6$.

Twin-width can be generalized to structures equipped with several binary relations. Then posets of bounded width, tree orders, and permutations omitting a fixed permutation also have bounded twin-width.

Classes of bounded twin-width enjoy many remarkable properties of combinatorial, algorithmic, and logical nature. For instance, classes of bounded twin-width are small (contain $n! \cdot 2^{O(n)}$ graphs with vertex set $\{1, \ldots, n\}$), are χ-bounded (the chromatic number is bounded in terms of the clique number) [7], and have the NIP property from model theory (every first-order formula $\varphi(x, y)$ defines a binary relation of bounded VC-dimension). Furthermore, model checking first-order logic (FO) is fixed-parameter tractable (FPT) on classes of bounded twin-width, assuming a *witness* of G having bounded twin-width is provided. This means that there is an algorithm which, given an FO sentence φ, a graph G together with a witness that its twin-width is at most d, decides whether φ holds in G in time $f(\varphi, d) \cdot |G|^c$ for some computable function f and universal constant c.

For each of the classes $\mathscr{C}$ mentioned above there is actually an algorithm which, given a graph $G \in \mathscr{C}$, computes the required witness of low twin-width, in polynomial time [8]. Hence FO model checking is FPT on these classes (without requiring the witness), generalizing many previous results, while it is AW[$*$]-hard (thus, unlikely to be FPT) on the class of all graphs [17]. It is however unknown whether such witness can be computed efficiently for every class $\mathscr{C}$ of bounded twin-width.

It transpires from the mere definitions that every graph can be equipped with some total order, resulting in an ordered graph of the same twin-width [8]. Such orders, while elusive to efficiently find, are crucial to most of the combinatorial and algorithmic applications. This suggests that *ordered graphs* of bounded twin-width are a more fundamental object than unordered graphs of bounded twin-width.

Main result. Our main result, Theorem 1 below, gives multiple characterizations of hereditary (i.e., closed under induced substructures) classes of ordered graphs of bounded twin-width, connecting notions from various areas of mathematics and theoretical computer science – enumerative combinatorics, model theory, parameterized complexity, matrix theory, and graph theory – and solving several open problems on the way. Furthermore, we show that if a class $\mathscr{C}$ of ordered graphs has bounded twin-width, then for each $G \in \mathscr{C}$, a witness that G has twin-width bounded by a constant can be computed in polynomial time. Consequently, FO model checking is FPT on $\mathscr{C}$. We also prove that the converse holds, under common complexity-theoretic assumptions.

We now briefly discuss some notions which are involved in our characterization. One of them involves certain forbidden patterns, which are obfuscated matchings between ordered sets, defined as follows. Fix a number $n \geqslant 1$ and let L and R be two copies of $[n] = \{1, \ldots, n\}$. Consider the four bipartite graphs $G^n_=, G^n_\leqslant, G^n_\geqslant, G^n_\neq$ with vertices $L \cup R$ corresponding to the binary relations $=, \leqslant, \geqslant, \neq$ on $L \times R$, as depicted in Figure 1. Fix parameters $P \in \{=, \neq, \leqslant, \geqslant\}$, $S \in \{l, r\}$ and $\lambda, \rho \in \{0, 1\}$, and define $\mathscr{M}_{P_S,\lambda,\rho}$ as the hereditary closure of the class of all ordered graphs obtained from the graphs G^n_P, for $n \geqslant 1$, as follows:

- if $S = l$, order $L \cup R$ by ordering $L \simeq [n]$ as usual, followed by the vertices of R ordered arbitrarily,
- if $S = r$, order $L \cup R$ by ordering L arbitrarily, followed by the vertices of $R \simeq [n]$ in reverse order,
- if $\lambda = 1$, create a clique on L, otherwise L remains an independent set,

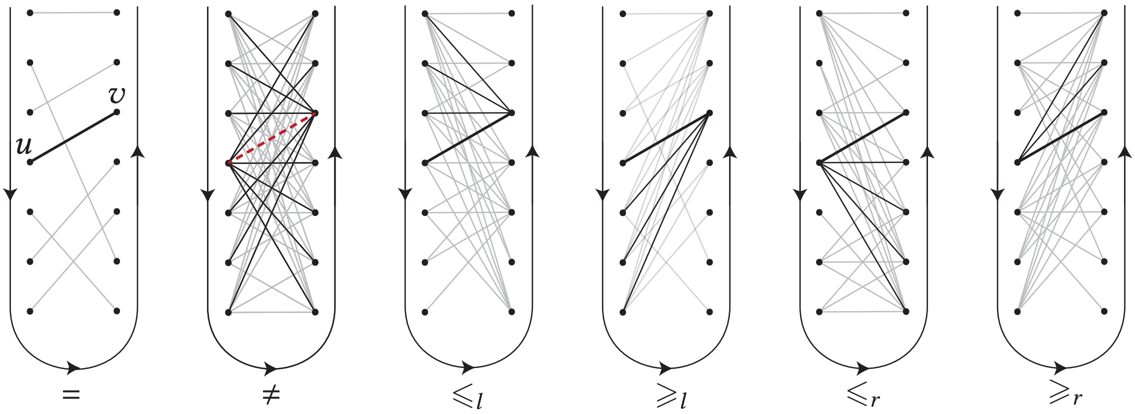

Figure 2 The six ordered graphs of $\mathscr{M}_{\mathrm{s},0,0}$ for $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ corresponding to the same ordered matching of $\mathscr{M}_{=,0,0}$ represented to the left. In each ordered graph the black edges are those implied by the bold edge uv in the matching. To picture the other classes $\mathscr{M}_{\mathrm{s},\rho,\lambda}$ with $(\lambda, \rho) \neq (0, 0)$, one just needs to turn the left part and/or the right part of each graph into cliques.

- if $\rho = 1$, create a clique on R, otherwise R remains an independent set.

As $\mathscr{M}_{=_l,\lambda,\rho}$ and $\mathscr{M}_{=_r,\lambda,\rho}$ are the same class, we denote it by $\mathscr{M}_{=,\lambda,\rho}$. Similarly $\mathscr{M}_{\neq_l,\lambda,\rho} = \mathscr{M}_{\neq_r,\lambda,\rho}$ is denoted $\mathscr{M}_{\neq,\lambda,\rho}$. Altogether there are 24 classes: $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ for each $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ and $\lambda, \rho \in \{0, 1\}$. For example, $\mathscr{M}_{=,0,0}$ consists of all induced subgraphs of ordered matchings with vertices $a_1 < \ldots < a_n < b_1 < \ldots < b_n$, whose edges form a matching between the a_i's and the b_j's, while $\mathscr{M}_{\neq,1,1}$ is the class of their edge complements. See Figure 2 for another representation of those classes.

In addition to these 24 classes, we will need to define one more class, $\mathscr{P}$. An *ordered permutation graph* associated with a permutation π of $[n] = \{1, \ldots, n\}$ is the ordered graph with vertices $[n]$ ordered naturally, such that two vertices $i < j$ are adjacent if and only if $\pi(i) > \pi(j)$. Let $\mathscr{P}$ be the class of all ordered permutation graphs.

Another characterization is in terms of adjacency matrices of the considered ordered graphs. A *d-division* of a matrix is a partition of its entries into d^2 zones using $d-1$ vertical lines and $d-1$ horizontal lines. A *rank-k d-division* is a d-division where each zone has at least k non-identical rows or at least k non-identical columns. A matrix has *grid rank* at least k if it has a rank-k k-division, simply called rank-k division. Note that this notion depends on the order of rows and columns in the matrix. The grid rank of an ordered graph is the grid rank of its adjacency matrix along the order of the graph.

A further characterization involves (simple first-order) *interpretations*, which is a notion originating in logic. Interpretations are a means of producing new structures out of old ones, using formulas. The new structure has the same domain as the old one (or a subset of it, defined by a formula $\delta(x)$) while each of its relations is defined by a formula $\varphi(\overline{x})$ interpreted in the old structure. For example, there is an interpretation which transforms a given graph G into its edge complement (using the formula $\neg E(x, y)$), and an interpretation which transforms G into its square (using the formula $\exists z. E(x, z) \wedge E(z, y)$). Transductions are a similar notion, but additionally allow to arbitrarily color the old structure before applying the interpretation and then use the colors in the formulas. Say that $\mathscr{C}$ *interprets* the class of all graphs if there is an interpretation I such that every (finite) graph G can be obtained as the result of I applied to some structure in $\mathscr{C}$. Replacing interpretations with transductions, we say that $\mathscr{C}$ *transduces* the class of all graphs. It is known that no class of bounded twin-width transduces the class of all graphs [8].

Finally, the *growth* of a class of structures $\mathscr{C}$ is the function counting, for a given n, the number of n-element structures in $\mathscr{C}$, up to isomorphism. It was previously shown that every class of ordered graphs of bounded twin-width has growth $2^{O(n)}$ [7].

We now state our main result, characterizing hereditary classes of ordered graphs with bounded twin-width. It provides a dichotomy result for all such classes: Either they have bounded twin-width, and are therefore well-behaved, or otherwise, they are very untamable.

▶ **Theorem 1.** *Let $\mathscr{C}$ be a hereditary class of ordered graphs. Then either $\mathscr{C}$ satisfies conditions (i)-(v), or $\mathscr{C}$ satisfies conditions (i')-(v') below:*

(i) $\mathscr{C}$ has bounded twin-width
(ii) $\mathscr{C}$ has bounded grid rank
(iii) $\mathscr{C}$ has growth $2^{O(n)}$
(iv) $\mathscr{C}$ does not transduce the class of all graphs
(v) FO *model checking is* FPT *on $\mathscr{C}$*

(i') $\mathscr{C}$ has unbounded twin-width
(ii') $\mathscr{C}$ contains $\mathscr{P}$ or one of the 24 classes $\mathscr{M}_{s,\lambda,\rho}$
(iii') $\mathscr{C}$ has growth at least $\sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k! \geqslant \lfloor \frac{n}{2} \rfloor !$
(iv') $\mathscr{C}$ interprets the class of all graphs
(v') FO *model checking is* AW[∗]*-hard on $\mathscr{C}$.*

The above result connects notions from graph theory (*i*), enumerative combinatorics (*iii*),(*iii'*), parameterized complexity (*v*),(*v'*), model theory (*iv*),(*iv'*), matrix theory (*ii*) and Ramsey theory (*ii'*). The lower bound in (*iii'*) is optimal.

Theorem 1 is proved in greater generality for arbitrary classes of ordered, binary structures. We prove an analogous result for classes $\mathcal{M}$ of 0, 1-matrices (and more generally matrix classes over finite alphabets). In this result (see Theorem 5), the lower bound on the number of $n \times n$-matrices in $\mathcal{M}$ in (*iii'*) is replaced by $n!$, and the 25 classes in (*ii'*) are reduced to six classes $\mathcal{F}_s$ of matrices, indexed by a single parameter $s \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$ (see Figure 3). Those are the class of all permutation matrices, the class obtained from permutation matrices by exchanging 0's with 1's, and four classes obtained from permutation matrices by propagating each 1 entry downward/upward/leftward/rightward, respectively.

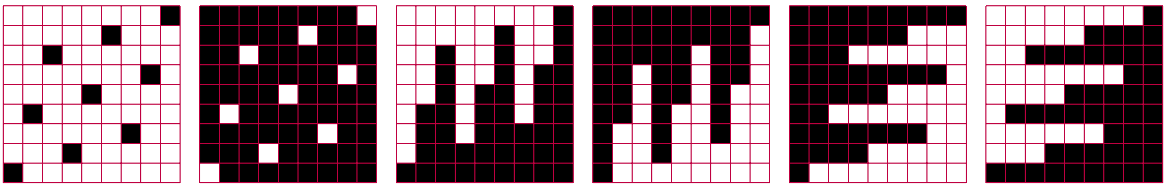

Figure 3 The matrices in $\mathcal{F}_=, \mathcal{F}_{\neq}, \mathcal{F}_{\leqslant_R}, \mathcal{F}_{\geqslant_R}, \mathcal{F}_{\leqslant_C}, \mathcal{F}_{\geqslant_C}$ (from left to right) for the same permutation matrix (the one to the left). The 1 entries are represented in black, the 0 entries, in white. As is standard with permutation patterns, we always place the first row of the matrix at the bottom.

As a consequence or by-product of Theorem 1, we settle a handful of questions in combinatorics and algorithmic graph theory. The main by-product is an approximation algorithm for twin-width in totally ordered binary structures.

▶ **Theorem 2.** *There is a fixed-parameter algorithm that, given a totally ordered binary structure G of twin-width k, outputs a witness that G has twin-width at most $2^{O(k^4)}$.*

As our second main result, we provide further characterizations of bounded twin-width classes in terms of model-theoretic notions, but which also transpire in algorithmic and structural graph theory. We consider arbitrary *monadically NIP* classes of relational structures, which are not necessarily finite, ordered, or binary. Those can be equivalently characterized as classes which do not transduce the class of all finite graphs. They include all graph classes of bounded twin-width (with or without an order), but also all transductions of *nowhere dense* classes [34], such as classes of bounded maximum degree.

The following theorem generalizes some notions and implications appearing in Theorem 1.

▶ **Theorem 3.** *For any class of structures $\mathscr{C}$, consider the following statements:*

(1) $\mathscr{C}$ does not transduce the class of all graphs,
(2) $\mathscr{C}$ is monadically NIP,
(3) $\mathscr{C}$ does not define large grids (see Definition 52),
(4) $\mathscr{C}$ is 1*-dimensional (see Definition 60),*
(5) $\mathscr{C}$ is a restrained class (see Definition 50).

Then the implications (1) ↔ (2) ↔ (3) ↔ (4) → (5) hold. For hereditary classes of binary, ordered structures, the above conditions are all equivalent to $\mathscr{C}$ having bounded twin-width.

Defining large grids generalizes the property of containing one of the classes $\mathscr{M}_{s,\lambda,\rho}$ or $\mathscr{P}$ to arbitrary structures, while the notion of a restrained class implies, for classes of ordered graphs, bounded grid rank. In particular, those notions do not require the structures to be ordered, finite, or binary. The notion of 1-dimensionality has a somewhat geometric flavor. It is defined in terms of a variant of forking independence – a central concept in stability theory, generalizing independence in vector spaces or algebraic independence. The equivalence (1) ↔ (2) is due to Baldwin and Shelah [2], the implications (2) → (3) → (4) are due to Shelah [38] and the implication (4) → (2) is a very recent result of Braunfeld and Laskowski [11]. Our contribution is the implication (4) → (5), and the overall equivalence in the case of ordered binary structures. Indeed, the implication (5) → (1) follows from the implication (*ii*)→(*iv*) in Theorem 1.

We now detail the consequences of our results.

Stanley-Wilf conjecture, Marcus-Tardos theorem. In the late 80's, Stanley and Wilf independently conjectured that every proper permutation class has single-exponential growth. To be more specific, every set of permutations closed under taking subpermutations (definition of permutation *class*) and not equal to the set of all permutations (definition of *proper*) contains at most c^n permutations over $[n]$, for some constant c depending solely on the class. This conjecture was confirmed in 2004 by Marcus and Tardos [33] (see Section 3.5 for more details). What Marcus and Tardos actually showed is that there is a function f such that every square $0,1$-matrix with at least $f(k)$ 1's in average per row (or column) admits a k-division where every zone contains a 1. This result is a milestone in combinatorics, and is at the core of the twin-width theory. We will use it again in the current paper.

Let us call *permutation matrix class* any (hereditary) class consisting of all the submatrices of a set of permutation matrices. In the language of matrices, the Stanley-Wilf conjecture/Marcus-Tardos theorem says that the growth of permutation matrix classes is (at least) $n!$ (for the class of all permutation matrices) or at most $2^{O(n)}$. Theorem 1 (*iii*),(*iii'*) extends that result to any matrix class over a finite alphabet, and the dividing line is the boundedness of twin-width.

We also obtain the following classification of all inclusion-minimal classes of superexponential growth. Call a (submatrix-closed) class of matrices a *Stanley-Wilf* class if it has superexponential growth, but every its proper subclass has at most exponential growth, that is, growth $2^{O(n)}$. Then the class of permutation matrices is a Stanley-Wilf class, as shown by Marcus and Tardos. By a similar argument, each of the classes $\mathcal{F}_s$ for $s \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$ is a Stanley-Wilf class. Moreover, these six classes are precisely *all* the Stanley-Wilf classes of $0,1$-matrices, and every matrix class of superexponential growth contains one of those classes. This is a consequence of our result for matrices, and the fact that the six classes are mutually incomparable (see Corollary 35).

In the same way we may define Stanley-Wilf classes of ordered graphs, as those hereditary classes of superexponential growth whose every proper, hereditary subclass has at most exponential growth. Then the 25 classes $\mathscr{M}_{s,\lambda,\rho}$ and $\mathscr{P}$ are precisely all the Stanley-Wilf classes of ordered graphs (see Corollary 48).

Speed gap in hereditary classes of ordered graphs. A couple of years after Marcus and Tardos proved the Stanley-Wilf conjecture, Balogh, Bollobás, and Morris [4, 3] analyzed the growth of ordered structures, and more specifically, ordered graphs, in an attempt to generalize Marcus and Tardos's ideas to new settings. They conjectured [4, Conjecture 2] that a hereditary class of (totally) ordered graphs has, up to isomorphism, either at most $2^{O(n)}$ graphs with n elements, or at least $n^{n/2+o(n)}$ such graphs, and proved it for weakly sparse graph classes, that is, without arbitrarily large bicliques (as subgraphs). In a concurrent work, Klazar [30] repeated that question, and more recently, Gunby and Pálvölgyi [26] observe that establishing the first superexponential jump in the growth of hereditary ordered graph classes is still an open question.

Theorem 1 (*iii*),(*iii'*) settles that question, and yields the optimal bound $\sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k!$, as conjectured by Balogh et al [4].

Small conjecture. Classes of bounded twin-width are small [7], that is, they contain at most $n!c^n$ distinct labeled n-vertex structures, for some constant c. (Actually they further contain at most c^n pairwise non-isomorphic structures [9].) The converse was conjectured for hereditary classes [7]. In the context of classes of totally ordered structures, it is simpler to drop the labeling and to count up to isomorphism. Indeed, every ordered structure has no non-trivial automorphism. Then a class is *small* if, up to isomorphism, it contains $2^{O(n)}$ distinct n-vertex structures. With that in mind, Theorem 1 (*i*),(*iii*) resolves the conjecture in the particular case of ordered graphs.

Ramsey theory. Our proofs are based on multiple Ramsey-theoretic arguments, but also our main result, Theorem 1, has a bearing on Ramsey theory. For example, we can conclude the following: *For every ordered matching H there is some cubic graph G such that for every total order $\leqslant$ on $V(G)$, the resulting ordered graph $G_\leqslant$ contains H as an induced subgraph.* Indeed, otherwise there is some ordered matching H such that every cubic graph G can be equipped with an order in a way which avoids H as an induced subgraph. This way, we obtain a class $\mathscr{C}$ of ordered subcubic graphs which does not contain the class $\mathscr{M}_{=,0,0}$, as it already fails to contain H. Clearly, $\mathscr{C}$ does not contain any of the remaining 24 classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ and $\mathscr{P}$, as those have unbounded degree. By Theorem 1 (*ii'*),(*i*), $\mathscr{C}$ has bounded twin-width. This implies that the class of all (unordered) subcubic graphs also has bounded twin-width, which we know is false (see [7]). More directly based on Theorem 1, a contradiction can be reached by observing that $\mathscr{C}$ does *not* have growth $2^{O(n)}$.

Transductions and interpretations. The study of transductions in theoretical computer science originates from the study of word-like and tree-like structures, such as graphs of bounded treewidth [1] or graphs of bounded clique-width [14]. Classes of bounded twin-width are closed under transductions [8]; in particular, no class of bounded twin-width transduces (nor interprets) the class of all graphs.

Theorem 1 (*i*),(*iv*) characterizes hereditary classes of ordered graphs of bounded twin-width as precisely those which do not transduce the class of all graphs. This is not unlike a result [15] characterizing classes of bounded clique-width as precisely those which do not transduce the class of all graphs via some transduction of counting monadic second-order logic (CMSO, an extension of first-order logic).

Grid theorems. Grid theorems are dichotomy results in structural graph theory which state that either a structure has a small *width*, or otherwise, a grid-like obstruction can be found in the structure. For example, this applies to the treewidth parameter and planar grids occurring as minors [36]. It also applies to clique-width and grids being definable in CMSO [15].

Theorem 3 proves an appropriate grid theorem for classes of ordered graphs of bounded twin-width, and more generally, for all classes which are not restrained. Indeed, such classes define large grids, which, intuitively, allows to define the 'same row' and 'same column' relations of arbitrarily large grids using first-order formulas in the graphs from the class. From this (also, from Theorem 1 (*i'*),(*iv'*)) it follows that if a hereditary class has unbounded twin-width then it interprets the class of all graphs.

There are other known grid theorems, including the Marcus-Tardos theorem itself.

Monadic NIP. Model theory typically classifies infinite structures according to the combinatorial complexity of families of definable sets. This is usually done through the introduction of tameness properties. The most important such notion is that of stability. A structure is stable if no formula $\varphi(x, y)$ encodes arbitrary large half-graphs (the graphs $G^n_\leqslant$ in Figure 1), which roughly means that there is no definable order on large subsets of the structure. A related, weaker, notion is that of NIP: a structure is *NIP* (or *dependent*) if no formula $\varphi(x, y)$ encodes arbitrary bipartite graphs. This captures the tameness properties of families of sets arising from geometric settings (for instance sets of points defined by polynomials of bounded degree).

The notion of monadically NIP (or monadically dependent) is a stronger requirement which says that the structure is NIP even if every subset of the domain can be used as a unary predicate. This is closely related to notions studied in finite model theory and structural graph theory. A class $\mathscr{C}$ of structures is monadically NIP if and only if it does not transduce the class of all graphs. By the recent result of Braunfeld and Laskowski [11], this is equivalent to not defining large grids.

Thus Theorem 1 (*i*),(*iv*) proves that a class of ordered graphs is monadically NIP if, and only if it has bounded twin-width. Furthermore, our results provide a model-theoretic characterization of classes of *unordered* graphs of bounded twin-width: a class of graphs has bounded twin-width if and only if it it can be obtained from some monadically NIP class of *ordered* graphs by forgetting the order.

Fixed-parameter tractable first-order model checking. Testing if a given FO sentence φ holds in a given structure G takes time $O(|G|^{|\varphi|})$ using a naive algorithm, and it is conjectured that the exponential dependency on $|\varphi|$ cannot be avoided. More precisely, it is conjectured that FO model checking is not *fixed-parameter tractable* (FPT) on the class of all graphs, i.e., does *not* admit an algorithm with running time $f(\varphi) \cdot |G|^c$, for some computable function $f\colon \mathbb{N} \to \mathbb{N}$ and constant c. That statement would indeed hold if FPT≠AW[$*$], as it is conjectured in parameterized complexity theory [19].

There is an ongoing program aiming to classify all the hereditary graph classes on which FO model checking is FPT. Currently such an algorithm is known for nowhere dense classes [25], for interpretations of bounded-degree classes [21, 22], for map graphs [18], for some families of intersection and visibility graphs [27], for transductions of bounded expansion classes when a suitable witness is given [23], and finally, for classes with bounded twin-width for which a witness of bounded twin-width can be efficiently computed [8]. Those include proper minor-closed classes, classes of bounded clique-width, posets of bounded width and permutations excluding a fixed permutation pattern.

All known tractable hereditary[1] classes are monadically NIP. This observation is the basis of the following conjecture:[2]

▶ **Conjecture 4.** *Let $\mathscr{C}$ be a hereditary class of structures. Then* FO *model checking is* FPT *on $\mathscr{C}$ if and only if $\mathscr{C}$ is monadically NIP.*

Both implications in the conjecture are open. This conjecture is now confirmed in two prominent cases, assuming FPT≠AW[$*$]:

- subgraph-closed graph classes [25], where monadically NIP classes are precisely nowhere dense classes,
- hereditary classes of ordered graphs, where monadically NIP classes are precisely classes with bounded twin-width.

The second item is by our main result, Theorem 1 (*i*),(*iv*),(*v*),(*v'*). Both the AW[$*$]-hardness result in the case of unbounded twin-width, and the FPT algorithm[3] in the case of bounded twin-width, are new. Furthermore Theorem 1 (*v*),(*v'*) indicates that for ordered graphs, twin-width is the dividing line of the tractability of FO model checking, as for (unordered) graphs, bounded treewidth is with MSO_2 (where quantifying over edge subsets is allowed) [32], and rank-width is with MSO_1 (where quantifying over edge subsets is disallowed).

Conjecture 4 in particular predicts tractability of every class of *unordered* graphs of bounded twin-width. Note that from [8] this is only known if the graph is provided with a witness of bounded twin-width. Our Theorem 2 gives the desired missing link for ordered graphs, that is, an FPT algorithm which either concludes that the twin-width is at least k, or provides a witness of the twin-width being bounded by some computable function of k. This is interesting on its own and gives some hope for the unordered case.

Related to this, we believe that Theorem 3 may be of independent interest, and possibly of broader applicability than just in the context of ordered, binary structures. For example, by Theorem 3, all graph classes of bounded twin-width (without an order) and all transductions of nowhere dense classes are restrained, generalizing the fact that classes of ordered graphs of bounded twin-width have bounded grid

[1] Tractable classes that are not hereditary include for example the class of all finite Abelian groups [10]
[2] This conjecture has been circulating in the community for some time, see e.g. the open problem session at the workshop on Algorithms, Logic and Structure in Warwick in 2016. See also [20, Conjecture 8.2].
[3] Recall that [8] only provides an FPT algorithm when a witness of bounded twin-width is given. Here we lift this requirement for classes of ordered graphs, thanks to Theorem 2.

rank. We remark that although our proof of Theorem 1 is purely combinatorial, an alternative proof can be derived from Theorem 3 (see our unpublished report [40]). This demonstrates that model-theoretic methods can be used in the context of algorithmic and structural graph theory, and that those two areas are intimately related.

Expressive power of least fixed-point logic. It is well-known that least fixed-point logic LFP captures the complexity class P on ordered structures. The *ordered conjecture* [31] asserts that LFP is more expressive than first-order logic on every infinite class $\mathscr{C}$ of finite ordered structures, in the sense that there is a sentence $\varphi \in \mathsf{LFP}$ which is not equivalent to any sentence $\varphi' \in \mathsf{FO}$ on $\mathscr{C}$.

The conjecture holds for every class $\mathscr{C}$ of ordered structures for which FO model checking is FPT (see [12, Cor. III.2]). In particular, it follows from Theorem 1 (iv),(v) that the ordered conjecture holds for every subclass of a hereditary dependent class of finite ordered binary structures.

2 Roadmap

For the most part, the proof will be carried out in the language of matrices over a finite alphabet. Matrices are considered ordered, in the sense that they are equipped with a total order on the rows and on the columns. A *class* of matrices is, by definition, assumed to be closed under taking submatrices, that is, removing rows and/or columns.

For the sake of simplicity, the description below concerns matrices with entries 0 or 1, called *0,1-matrices*. A 0,1-matrix can be seen as a relational structure whose domain consists of its rows and columns, equipped with two unary predicates marking the rows and the columns, respectively, a total order which places the rows before the columns, and a binary, symmetric relation which relates a row with a column if the entry at their intersection is equal to 1.

Recall that the six matrix classes of unbounded twin-width which arise are: the class $\mathcal{F}_=$ of all permutation matrices, the class $\mathcal{F}_{\neq}$ obtained from permutation matrices by exchanging 0's with 1's, and four classes $\mathcal{F}_{\leqslant R}, \mathcal{F}_{\geqslant R}, \mathcal{F}_{\leqslant C}, \mathcal{F}_{\geqslant C}$ obtained from permutation matrices by propagating each value 1 downward/upward/leftward/rightward, respectively (see Figure 3). The *growth* of a class of matrices is the function counting the number of distinct (square) $n \times n$-matrices in the class, for a given $n \geqslant 1$.

Our main result concerning (hereditary) classes of matrices over a finite alphabet is as follows.

▶ **Theorem 5.** *Given a class $\mathcal{M}$ of 0,1-matrices, the following are equivalent.*

(i) *$\mathcal{M}$ has bounded twin-width.*
(ii) *$\mathcal{M}$ has bounded grid rank.*
(iii) *$\mathcal{M}$ does not contain any of the six classes $\mathcal{F}_=, \mathcal{F}_{\neq}, \mathcal{F}_{\leqslant R}, \mathcal{F}_{\geqslant R}, \mathcal{F}_{\leqslant C}, \mathcal{F}_{\geqslant C}$.*
(iv) *$\mathcal{M}$ does not interpret the class of all graphs.*
(v) *$\mathcal{M}$ does not transduce the class of all graphs.*
(vi) *$\mathcal{M}$ does not have growth at least $n!$.*
(vii) *$\mathcal{M}$ has growth at most $2^{O(n)}$.*
(viii) FO *model checking is* FPT *on $\mathcal{M}$. (The implication from (viii) holds if* FPT≠AW[∗]*.)*
(ix) *there is some $r \in \mathbb{N}$ such that no matrix $M \in \mathcal{M}$ admits a r-rich division.*

The last condition, (ix), is a technical one, whose definition is deferred to Section 3. This will be a key intermediate step in proving that (ii) implies (i), as well as in getting an approximation algorithm for the twin-width of a matrix. Theorem 5 reads the same for matrices over a finite alphabet A, except that (iii) is replaced by: No selection of $\mathcal{M}$ contains any of the six classes $\mathcal{F}_=, \mathcal{F}_{\neq}, \mathcal{F}_{\leqslant R}, \mathcal{F}_{\geqslant R}, \mathcal{F}_{\leqslant C}, \mathcal{F}_{\geqslant C}$, where for $a \in A$, the *a-selection* of a matrix class $\mathcal{M}$ is the class obtained from the matrices of $\mathcal{M}$ by replacing the letter $a \in A$ with 1 and the remaining letters with 0. In Figure 4, a class satisfying (iii) is called *pattern-avoiding* (the definition will be recalled in Section 3.9).

As mentioned in the introduction, we prove an analogous result (see Theorem 1) for classes of ordered graphs, or more generally for classes of ordered binary structures. In an informal nutshell, the high points of the paper read: For hereditary ordered binary structures, bounded twin-width, small, subfactorial

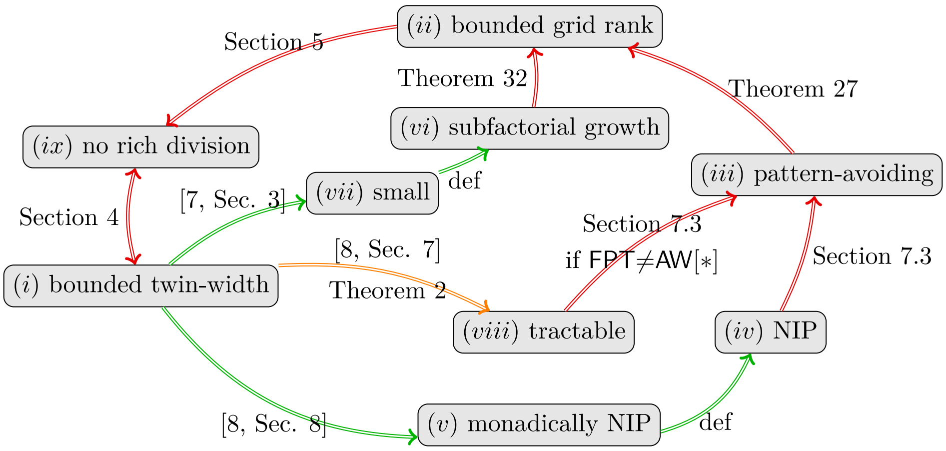

Figure 4 A bird's eye view of the paper. In green (arrows without a reference to a part of the paper), the implications that were already known for general binary structures. In red (other arrows except $(i) \Rightarrow (viii)$), the new implications for matrices on finite alphabets, or ordered binary structures. The effective implication $(i) \Rightarrow (ix)$ is useful for Theorem 2. See Figure 6 for a more detailed proof diagram, distinguishing what is done in the language of matrices and what is done in the language of ordered graphs.

growth, and tractability of FO model checking are all equivalent. We conclude by giving a more detailed statement of the approximation algorithm.

▶ **Theorem 2** (more precise statement). *There is a fixed-parameter algorithm, which, given an ordered binary structure G, encoded by a matrix M, and a parameter k, either outputs*

- *a $2^{O(k^4)}$-sequence of G, implying that $\text{tww}(G) = 2^{O(k^4)}$, or*
- *a $2k(k+1)$-rich division of M, implying that $\text{tww}(G) > k$.*

We now outline the proof of Theorems 1 and 5.

Proof outline. Bounded twin-width is already known to imply interesting properties: FO model checking is FPT if a witness of small twin-width is part of the input [8], monadic dependence [8], smallness [7] (see the green and orange arrows in Figures 4 and 6). For a characterization of some sort in the particular case of ordered structures, the challenge is to find interesting properties implying bounded twin-width. A central characterization in the first paper of the series [8] goes as follows.

A graph class $\mathscr{C}$ has bounded twin-width if and only if there is a constant d such that every graph $G \in \mathscr{C}$ can be ordered so that the adjacency matrix along that order has no d-division where each zone contains two non-identical rows and two non-identical columns. The backward direction is effective: From such an ordering, we obtain a witness of bounded twin-width in polynomial time.

Now that we consider *ordered* graphs it is tempting to try this order to get a witness of low twin-width. Things are not that simple. Consider the graph G with vertices $1, \ldots, n$ ordered naturally, where two vertices are adjacent if and only if they have different parity. The adjacency matrix M of G is the checkerboard matrix to the left in Figure 5. This matrix is fairly simple and indeed has bounded twin-width (this will be evident once we formally define twin-width). Yet for $d = \lfloor n/2 \rfloor$, the matrix M has a d-division where each zone has two different rows and columns. Now a *good reordering* of G would put all the odd-indexed vertices together, followed by all the even-indexed vertices. Then the adjacency matrix M' of G along the new order (Figure 5, right), where the entries of M' now encode the edges of G as well as the original order, is such that every 4-division contains a constant zone.

Can we find such reorderings automatically? Eventually we can, but a crucial opening step is precisely to nullify the importance of the reordering. We show that matrices have bounded twin-width exactly when they have bounded grid rank, that is, they do not admit rank-k divisions for arbitrary k. This natural strengthening on the condition that cells should satisfy (from rank 2 to rank k) exempts us from the need to reorder. Note that the checkerboard matrix does not have any rank-k division already for $k = 3$.

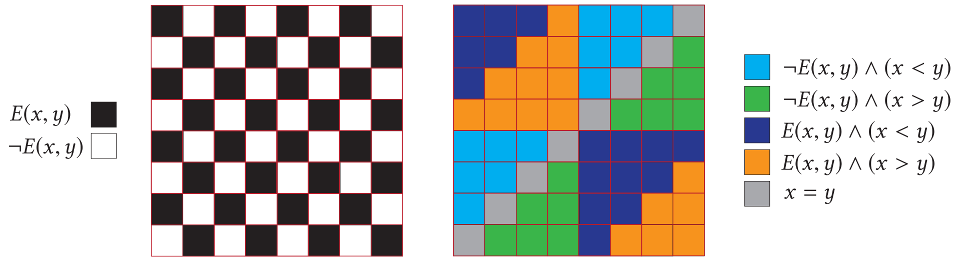

Figure 5 *Left:* The adjacency matrix of the ordered graph G with vertices $1, \ldots, n$ and edges ij such that $i + j$ is odd, along the usual order. (The first row is at the bottom.) *Right:* The adjacency matrix along another order, encoding the adjacency as well as the original order. Every 4-division contains some constant zone.

An important intermediate step is provided by the concept of rich divisions (see Section 3.4 for a definition). We first prove that a greedy strategy to find a potential witness of bounded twin-width can only be stopped by the presence of a large rich division; thus, unbounded twin-width implies the existence of arbitrarily large rich divisions. This brings a theme developed in [8] to the ordered world. In turn, leveraging the Marcus-Tardos theorem, we show that huge rich divisions contain large rank divisions (this is almost entirely summarized by Figure 9).

By a series of Ramsey-like arguments, we find in large rank divisions more and more structured submatrices encoding universal permutations. Eventually we find at least one of six encodings of all permutations. More precisely, each class of unbounded grid rank contains one of the classes $\mathcal{F}_{\mathrm{s}}$, for some $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$. As each of the classes $\mathcal{F}_{\mathrm{s}}$ has growth $n!$, this chain of implications shows that hereditary classes with unbounded grid rank have growth at least $n!$. Conversely, classes of matrices of bounded twin-width have growth $2^{O(n)}$ by [7]. That establishes the announced speed gap for matrix classes. Moreover, as each of the classes $\mathcal{F}_{\mathrm{s}}$ interprets the class of all graphs and has an AW[$*$]-hard model checking, we obtain Theorem 5.

Finally we translate the permutation encodings in the language of ordered graphs. This allows us to refine the growth gap specifically for ordered graphs. We also prove that including a family $\mathcal{F}_{\mathrm{s}}$, or its ordered-graph equivalent $\mathscr{M}_{\mathrm{s},\lambda,\rho}$, is an obstruction to being NIP. This follows from the fact that the class of all permutation graphs is independent. As we get an effectively constructible interpretation to the class of all structures (matrices or ordered graphs), we conclude that FO model checking is not FPT on hereditary classes of unbounded twin-width. This is the end of the road. The remaining implications to establish the equivalences of Theorems 1 and 5 come from [8, Sections 7 and 8], [7, Section 3], and Theorem 2 (see Figure 6).

Theorem 3 is proved using model-theoretic methods. In particular, it relies on a suitable analogue of forking independence for monadically NIP classes.

Organization. The rest of the paper is organized as follows. In Section 3, we formally define all the notions used in the rest of the paper. In Section 4, we show that *(i)* and *(ix)* are equivalent. As a by-product, we obtain a fixed-parameter $f(\mathrm{OPT})$-approximation algorithm for the twin-width of ordered matrices. In Section 5, we prove the implication *(ii)* $\rightarrow$ *(ix)*. In Section 6, we introduce so-called *rank Latin divisions* and show that large rank divisions contain large rank Latin divisions. In Section 7, we further clean the rank Latin divisions in order to show that *(iii)* $\rightarrow$ *(ii)* and *(vi)* $\rightarrow$ *(ii)*. In Section 8, we show that *(viii)* $\rightarrow$ *(iii)* and *(iv)* $\rightarrow$ *(iii)* transposed to the language of ordered graphs. We also refine the lower bound on the growth of ordered graph classes with unbounded twin-width, to completely settle Balogh et al.'s conjecture [4]. See Figure 6 for a visual outline. Finally, in Section 9 we prove Theorem 3.

3 Preliminaries

Everything which is relevant to the rest of the paper will now be properly defined. We may denote by $[i, j]$ the set of integers that are at least i and at most j, and $[i]$ is a short-hand for $[1, i]$. We start with

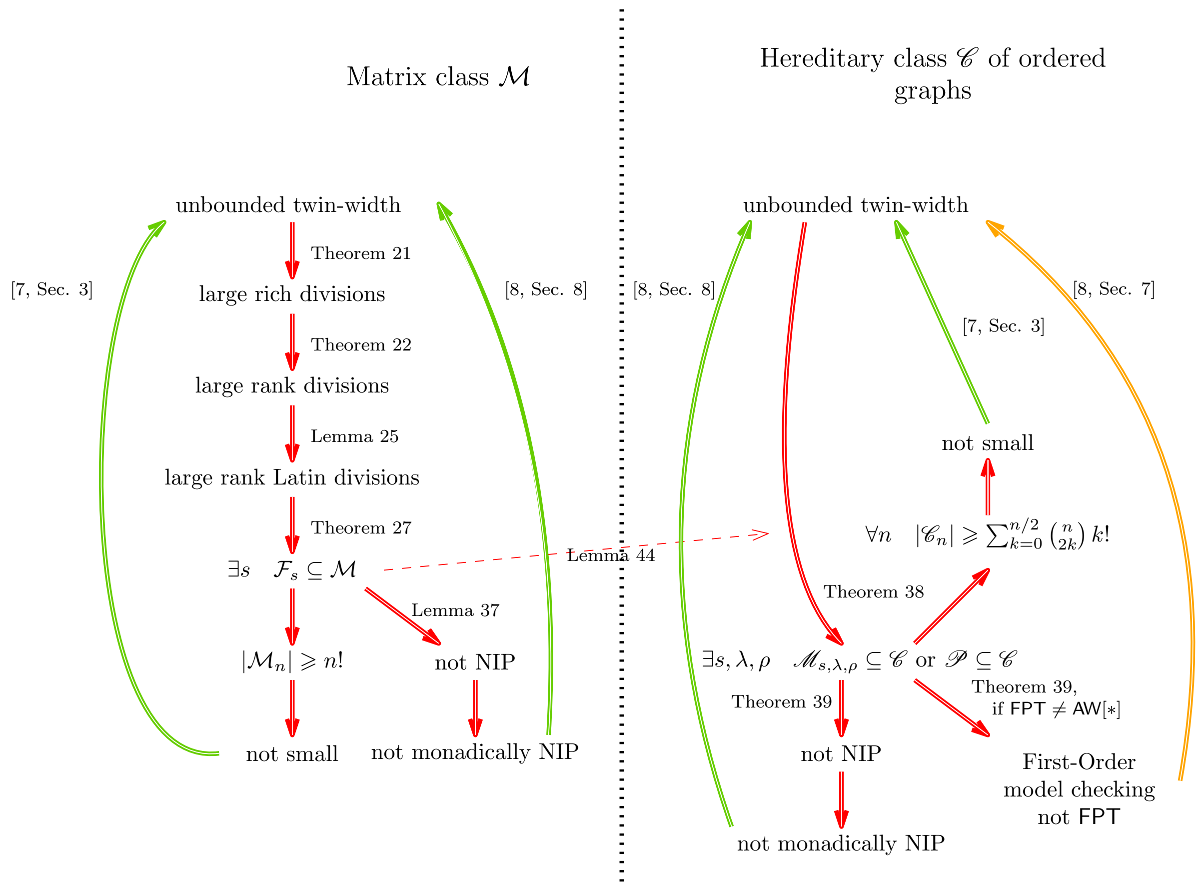

Figure 6 A more detailed proof diagram.

the combinatorial objects.

3.1 Graphs, orders, matrices, permutations

By *graph* we mean a simple, undirected graph G, and denote its set of vertices $V(G)$ and set of edges $E(G)$. An edge with endpoints u and v is denoted uv or vu. A *total order* on a set X is a binary relation $<$ which is transitive, irreflexive, such that for all $x, y \in X$ either $x < y$ or $y < x$ holds. An *ordered graph* is a graph together with a total order on its vertices. The *edge complement* of a graph (resp. ordered graph) G is the graph (resp. ordered graph) obtained from G by replacing edges by non-edges, and vice-versa.

A *matrix* M over a finite alphabet A is a function $M\colon R \times C \to A$, where R is a totally ordered set of rows and C is a totally ordered set of columns. The value $M(r, c)$, that we will often denote $M_{r,c}$, is the *entry* of M at position (r, c), or in row r and column c. We may say that M is an $R \times C$ matrix, or an $n \times m$ matrix, where $n = |R|$ and $m = |C|$.

A $0,1$-matrix is a matrix over the alphabet $\{0, 1\}$. A $0,1$-matrix with rows R and columns C can be viewed as an ordered graph with vertices $R \uplus C$, total order $<$ obtained from the orders on R and C by making all the columns larger than all the rows, and edges rc such that $r \in R$, $c \in C$ and $M(r, c) = 1$.

We distinguish matrices only up to isomorphisms which preserve the order of the rows and columns. A *submatrix* of a matrix M is any matrix obtained from M by deleting a (possibly empty) set of rows and columns. Analogously to permutation classes which are by default supposed closed under taking subpermutations (or patterns), we will define a *class* of matrices as a set of matrices closed under taking submatrices. The *submatrix closure* of a matrix M is the set of all submatrices of M (including M itself). Thus our matrix classes include the submatrix closure of every matrix they contain. On the contrary, classes of (ordered) graphs are only assumed to be closed under isomorphism. A *hereditary*

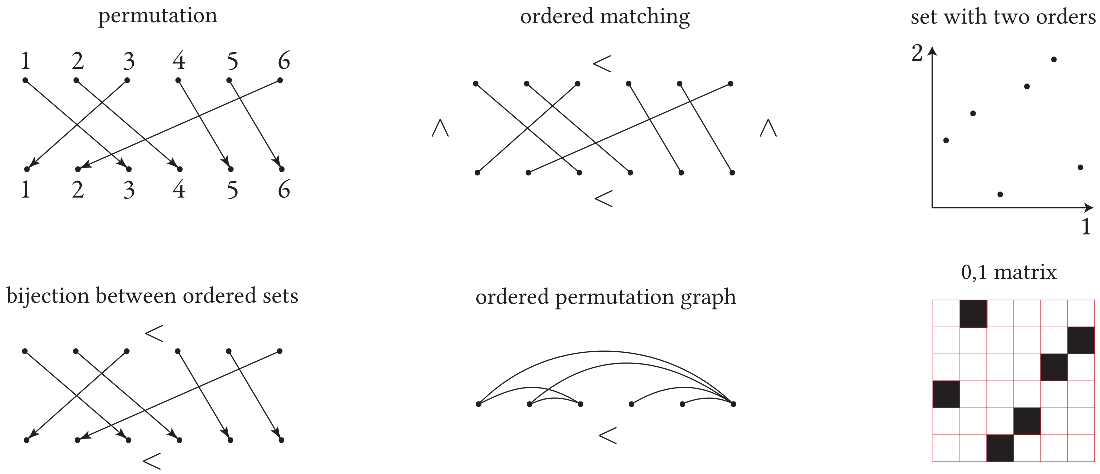

Figure 7 Six different views on the same permutation. We use the convention that the first row of a matrix is at the bottom.

class of (ordered) graphs (resp. binary structures) is one that is closed under taking induced subgraphs (resp. induced substructures).

An *n-permutation*, for $n \geqslant 1$, is a bijection $\pi\colon [n] \to [n]$. The set of all n-permutations is denoted $\mathfrak{S}_n$. Permutations turn out to be of central importance in the theory developed here, and indeed, twin-width has its origins in the Stanley-Wilf conjecture which is precisely about permutations. As we will see, classes of ordered graphs or matrices with unbounded twin-width are exactly those which contain encodings of all permutations, under a suitable encoding.

We will use several views on permutations (see Figure 7): as bijections between two ordered sets, as sets equipped with two total orders, as ordered matchings, as ordered permutation graphs, and as $0,1$-matrices.

An n-permutation π may be viewed as a bijection π between two totally ordered sets, namely $X = ([n], <)$ and $Y = ([n], <)$. Conversely, for every bijection $f\colon X \to Y$ between two totally ordered sets of size n there is a unique n-permutation π such that $f = i_Y^{-1} \circ \pi \circ i_X$ holds for the unique order-preserving bijections $i_X\colon X \to [n]$ and $i_Y\colon Y \to [n]$. Using this correspondence, we may define the notion of a *subpermutation*. A subpermutation of an n-permutation π induced by a set $U \subseteq [n]$ is the unique $|U|$-permutation which corresponds to the restriction $\pi|_U$, treated as a bijection between the ordered sets $U \subseteq [n]$ and $\pi(U) \subseteq [n]$, via the correspondence described above.

Similarly, an n-permutation π defines two orders on $[n]$, namely the usual order $<_1$, and the order $<_2$ such that $i <_2 j$ if and only if $\pi(i) < \pi(j)$. Conversely, every finite set equipped with two total orders is isomorphic to one obtained from a permutation as described above. Via this correspondence, subpermutations correspond exactly to induced substructures of sets equipped with two total orders.

An *ordered matching* is an ordered graph with vertices $a_1 < \ldots < a_n < b_1 < \ldots < b_n$ such that each a_i is adjacent with exactly one b_j, and vice-versa. Hence, there is a unique n-permutation π such that a_i is adjacent with $b_{\pi(i)}$, for $i \in [n]$.

An *ordered permutation graph* associated with an n-permutation π is the ordered graph G_π with vertices $[n]$ ordered naturally, such that $i < j$ are adjacent if and only if $\pi(i) > \pi(j)$. Note that the isomorphism type of G_π determines the permutation π uniquely. If σ is a subpermutation of π induced by $U \subseteq [n]$ then G_σ is the subgraph of G_π induced by U. Observe that the edge complement of a permutation graph G_π is also a permutation graph. Namely, if G_π corresponds to two total orders $<_1, <_2$ on $[n]$, as explained above, then the edge complement of G_π corresponds to the orders $<_1, >_2$ on $[n]$.

Finally, n-permutations correspond to $n \times n$ $0,1$-matrices with exactly one 1 in each row and in each column. We adopt the standard conventions that the 1 entries in the matrix of permutation $\sigma \in \mathfrak{S}_n$ are at positions $(i, \sigma(i))$ for $i \in [n]$, and that, in this context of patterns, the first row is placed at the bottom.

A permutation σ is a subpermutation of π if the matrix of σ is a submatrix of the matrix of π.

In fact, we will see even more representations of permutations as matrices or ordered graphs, namely five further matrix classes and twenty-three further classes of ordered graphs.

3.2 Structures

A relational *signature* Σ is a finite set of relation symbols R, each with a specified arity $r \in \mathbb{N}$. A *Σ-structure* $\mathbf{A}$ is defined by a set A (the *domain* of $\mathbf{A}$) together with a relation $R^{\mathbf{A}} \subseteq A^r$ for each relation symbol $R \in \Sigma$ with arity r. The syntax and semantics of first-order formulas over Σ, or *Σ-formulas* for brevity, are defined as usual.

A graph is viewed as a structure over the signature with one binary relation E indicating the adjacency between vertices. A total order is viewed as a structure over the signature with one binary relation $<$. An ordered graph is viewed as a structure over the signature with two binary relations, E and $<$. More generally, an *ordered binary structure* is a structure $\mathbf{A}$ over a signature Σ consisting of unary and binary relation symbols which includes the symbol $<$, and such that $<$ defines in $\mathbf{A}$ a total order on the domain of $\mathbf{A}$.

A matrix M over a finite alphabet A with rows R and columns C is viewed as an ordered binary structure with domain $R \uplus C$, equipped with the following relations:

- unary relations R and C, interpreted as the set of rows and set of columns, respectively,
- a binary relation $<$ which defines a total order on $R \uplus C$, extending the total orders on the rows and columns of M in such a way that the rows precede the columns,
- one binary relation E_a, for each $a \in A$, where $E_a(r, c)$ holds if and only if r is a row, c is a column, and a is the entry of M at row r and column c.

3.3 Twin-width

In the first paper of the series [8], we define twin-width for general binary structures via *unordered* matrices. The twin-width of (ordered) matrices can be defined this way by encoding the total orders on the rows and on the columns with two binary relations. However we will give an equivalent definition, tailored to ordered structures. This slight shift is already a first step in understanding these structures better, with respect to twin-width. We insist that matrices are always ordered objects, in the current paper. Thus the twin-width of a matrix does *not* coincide with the twin-width of unordered matrices, as defined in [8].

Let M be an $n \times m$ matrix with entries ranging over a fixed finite set. We denote by $R := \{r_1, \dots, r_n\}$ its set of rows and by $C := \{c_1, \dots, c_m\}$ its set of columns. Let S be a non-empty subset of columns, c_a be the column of S with minimum index a, and c_b, the column of S with maximum index b. The *span* of S is the set of columns $\{c_a, c_{a+1}, \dots, c_{b-1}, c_b\}$. We say that a subset $S \subseteq C$ is in *conflict* with another subset $S' \subseteq C$ if their spans intersect. A partition $\mathcal{P}$ of C is *k-overlapping* if every part of $\mathcal{P}$ is in conflict with at most k other parts of $\mathcal{P}$. The definitions of *span*, *conflict*, and *k-overlapping partition* similarly apply to sets of rows. With that terminology, a *division* is a 0-overlapping partition.

A partition $\mathcal{P}$ is a *contraction* of a partition $\mathcal{P}'$ (defined on the same set) if it is obtained by merging two parts of $\mathcal{P}'$. A *contraction sequence* of M is a sequence of partitions $\mathcal{P}_1, \dots, \mathcal{P}_{n+m-1}$ of the set $R \cup C$ such that $\mathcal{P}_1$ is the partition into $n + m$ singletons, $\mathcal{P}_{i+1}$ is a contraction of $\mathcal{P}_i$ for all $i \in [n + m - 2]$, and $\mathcal{P}_{n+m-1} = \{R, C\}$. In other words, we merge at every step two *column parts* (made exclusively or columns) or two row parts (made exclusively or rows), and terminate when all rows and all columns both form a single part. We denote by $\mathcal{P}_i^R$ the partition of R induced by $\mathcal{P}_i$ and by $\mathcal{P}_i^C$ the partition of C induced by $\mathcal{P}_i$. A *contraction sequence* is *k-overlapping* if all partitions $\mathcal{P}_i^R$ and $\mathcal{P}_i^C$ are k-overlapping partitions. Note that a 0-overlapping sequence is a sequence of divisions.

If S^R is a subset of R, and S^C is a subset of C, we denote by $S^R \cap S^C$ the submatrix at the intersection of the rows of S^R and of the columns of S^C. Given some column part C_a of $\mathcal{P}_i^C$, the *error value* of C_a is the number of row parts R_b of $\mathcal{P}_i^R$ for which the submatrix $C_a \cap R_b$ of M is not constant. The error value is defined similarly for rows, by switching the role of columns and rows. The *error value* of $\mathcal{P}_i$ is the maximum error value of some part in $\mathcal{P}_i^R$ or in $\mathcal{P}_i^C$. A contraction sequence is a *(k, e)-sequence* if all partitions $\mathcal{P}_i^R$ and $\mathcal{P}_i^C$ are k-overlapping partitions with error value at most e. Strictly speaking, to be consistent with the definitions in the first paper [8], the *twin-width* of a matrix M, denoted by $\mathrm{tww}(M)$, is the minimum $k+e$ such that M has a (k,e)-sequence. This matches, setting $d := k+e$, what we called a *d-sequence* for the binary structure encoding M [8]. We will however not worry about the exact value of twin-width, but merely whether it is bounded or unbounded on a class of structures. Thus for the sake of simplicity, we often consider the minimum integer k such that M has a (k,k)-sequence. This integer is indeed sandwiched between $\mathrm{tww}(M)/2$ and $\mathrm{tww}(M)$.

The twin-width of a matrix class $\mathcal{M}$, denoted by $\mathrm{tww}(\mathcal{M})$, is simply defined as the supremum of $\{\mathrm{tww}(M) \mid M \in \mathcal{M}\}$. We say that $\mathcal{M}$ has *bounded twin-width* if $\mathrm{tww}(\mathcal{M}) < \infty$, or equivalently, if there is a finite integer k such that every matrix $M \in \mathcal{M}$ has twin-width at most k. A class $\mathscr{C}$ of ordered graphs has *bounded twin-width* if all the adjacency matrices of graphs $G \in \mathscr{C}$ along their vertex ordering, or equivalently their submatrix closure, form a set/class with bounded twin-width.

We can more generally define the twin-width of ordered binary structures via matrices. The *matrix encoding* of an ordered binary structure $\mathbf{A}$ with relations $<, E_1, \ldots, E_p$ is the $|A| \times |A|$ matrix over the alphabet $\{-1,0,1,2\}^p$ whose entry at position (x,y), is the vector $(b_1,\ldots,b_p) \in \{-1,0,1,2\}^p$ such that $b_i = 1$ if $E_i(x,y) \wedge \neg E_i(y,x)$ holds, $b_i = -1$ if $\neg E_i(x,y) \wedge E_i(y,x)$ holds, $b_i = 2$ if $E_i(x,y) \wedge E_i(y,x)$ holds, and $b_i = 0$ otherwise. Then the twin-width of an ordered binary structure $\mathbf{A}$ is simply the twin-width of the matrix encoding of $\mathbf{A}$. We choose this particular encoding so that the vector $(b_1, b_2, \ldots, b_p)$ at position (x,y) and the one $(b'_1, b'_2, \ldots, b'_p)$ at position (y,x) satisfies $b_i = \pm b'_i$ for every $i \in [p]$. We then say that the matrix is *mixed-symmetric* as in each vector some coordinates are symmetric while others are skew-symmetric. This technicality allows to turn a contraction sequence of a matrix encoding into a contraction sequence of its associated ordered binary structure. For more details, see [8, Section 5, Theorem 14].

3.4 Rank division and rich division

We recall that a division $\mathcal{D}$ of a matrix M is a pair $(\mathcal{D}^R, \mathcal{D}^C)$, where $\mathcal{D}^R$ (resp. $\mathcal{D}^C$) is a partition of the rows (resp. columns) of M into (contiguous) intervals, or equivalently, a 0-overlapping partition. A *d-division* is a division satisfying $|D^R| = |D^C| = d$. For every pair $R_i \in \mathcal{D}^R$, $C_j \in \mathcal{D}^C$, the submatrix $R_i \cap C_j$ may be called *zone* (or *cell*) of $\mathcal{D}$ since it is, by definition, a contiguous submatrix of M. We observe that a d-division defines d^2 zones.

A *rank-k d-division* of M is a d-division $\mathcal{D}$ such that for every $R_i \in \mathcal{D}^R$ and $C_j \in \mathcal{D}^C$ the zone $R_i \cap C_j$ has at least k distinct rows or at least k distinct columns. A *rank-k division* is simply a short-hand for a rank-k k-division. The *grid rank* of a matrix M is the largest integer k such that M admits a rank-k division. A class $\mathcal{M}$ has *bounded grid rank* if there is some integer k such that every matrix $M \in \mathcal{M}$ has grid rank less than k, or equivalently, for every k-division $\mathcal{D}$ of M, there is a zone of $\mathcal{D}$ with less than k distinct rows and less than k distinct columns.

Closely related to rank divisions, a *k-rich division* is a division $\mathcal{D}$ of a matrix M on rows and columns $R \cup C$ such that:

- for every part R_a of $\mathcal{D}^R$ and for every subset Y of at most k parts in $\mathcal{D}^C$, the submatrix $R_a \cap (C \setminus \cup Y)$ has at least k distinct row vectors, and symmetrically
- for every part C_b of $\mathcal{D}^C$ and for every subset X of at most k parts in $\mathcal{D}^R$, the submatrix $(R \setminus \cup X) \cap C_b$ has at least k distinct column vectors.

Informally, in a large rich division (that is, a k-rich division for some large value of k), the diversity in the column vectors within a column part cannot drop too much by removing a controlled number of row parts. And the same applies to the diversity in the row vectors.

3.5 Enumerative combinatorics

In the context of unordered structures, a graph class $\mathcal{C}$ is *small* if there is a constant c, such that its number of n-vertex graphs bijectively labeled by $[n]$ is at most $n!c^n$. When considering totally ordered structures, for which the identity is the unique automorphism, one can advantageously drop the labeling and the $n!$ factor. Indeed, on these structures, counting up to isomorphism or up to equality is the same.

Thus a matrix class $\mathcal{M}$ is *small* if there exists a real number c such that the total number of $n \times m$ matrices in $\mathcal{M}$ is at most $c^{\max(n,m)}$.

Marcus and Tardos [33] showed the following central result, henceforth referred to as *Marcus-Tardos theorem*, which by an argument due to Klazar [29] was known to imply the Stanley-Wilf conjecture, that permutation classes avoiding any fixed pattern are small.

▶ **Theorem 6.** *There exists a function* $\mathrm{mt} : \mathbb{N} \to \mathbb{N}$ *such that every* $n \times m$ *matrix* M *with at least* $\mathrm{mt}(k)\max(n,m)$ *non-zero entries has a* k*-division in which every zone contains a non-zero entry.*

We call $\mathrm{mt}(\cdot)$ the *Marcus-Tardos bound.* The current best bound is $\mathrm{mt}(k) = \frac{8}{3}(k+1)^2 2^{4k} = 2^{O(k)}$ [13]. Among other things, the Marcus-Tardos theorem is a crucial tool in the development of the theory around twin-width. In the second paper of the series [7], the Stanley-Wilf conjecture/Marcus-Tardos theorem was generalized to classes with bounded twin-width. We showed that every graph class with bounded twin-width is small (while proper subclasses of permutation graphs have bounded twin-width [8]). This can be readily extended to every bounded twin-width class of binary structures. It was further conjectured that the converse holds for hereditary classes: Every hereditary small class of binary structures has bounded twin-width. We will confirm this conjecture, in the current paper, for the special case of totally ordered binary structures.

We denote by $\mathcal{M}_n$, the *n-slice* of a matrix class $\mathcal{M}$, that is the set of all $n \times n$ matrices of $\mathcal{M}$ (up to isomorphism which preserves the order on the rows and columns). The *growth* (or *speed*) of a matrix class is the function $n \in \mathbb{N} \mapsto |\mathcal{M}_n|$. A class $\mathcal{M}$ has *subfactorial* growth if there is a finite integer beyond which the growth of $\mathcal{M}$ is strictly less than $n!$; more formally, if there is n_0 such that for every $n \geqslant n_0$, $|\mathcal{M}_n| < n!$. Similarly, $\mathscr{C}$ being a class of ordered graphs, the *n-slice* of $\mathscr{C}$, $\mathscr{C}_n$, is the set of n-vertex ordered graphs in $\mathscr{C}$, up to isomorphism. And the *growth* (or *speed*) of a class $\mathscr{C}$ of ordered graphs is the function $n \in \mathbb{N} \mapsto |\mathscr{C}_n|$.

The following result follows from [7].

▶ **Theorem 7.** *Let* $\mathscr{C}$ *be a class of ordered binary structures of bounded twin-width. Then* $\mathscr{C}$ *has growth* c^n *for some constant* c*.*

3.6 Computational complexity

First-order (FO) matrix model checking asks, given a matrix M (or a totally ordered binary structure $\mathcal{S}$) and a first-order sentence φ (i.e., a formula without any free variable), if $M \models \varphi$ holds. The atomic formulas in φ are of the kinds described in Section 3.2.

We then say that a matrix class $\mathcal{M}$ is *tractable* if FO model checking is fixed-parameter tractable (FPT) when parameterized by the sentence size and the input matrices are drawn from $\mathcal{M}$. That is, $\mathcal{M}$ is tractable if there exists a constant c and a computable function f, such that $M \models \varphi$ can be decided in time $f(\ell)\,(m+n)^c$, for every $n \times m$ matrix $M \in \mathcal{M}$ and FO sentence φ of quantifier depth ℓ. We may denote the size of M, $n+m$, by $|M|$, and the quantifier depth (i.e., the maximum number of nested quantifiers) of φ by $|\varphi|$. Similarly a class $\mathscr{C}$ of binary structures is *tractable* if FO model checking is FPT on $\mathscr{C}$.

FO model checking of general (unordered) graphs is AW[$*$]-complete [17], and thus very unlikely to be FPT. Indeed FPT≠AW[$*$] is a much weaker assumption than the already widely-believed Exponential Time Hypothesis [28], and if false, would in particular imply the existence of a subexponential algorithm solving 3-SAT. In the first paper of the series [8], we show that FO model checking of general binary structures of bounded twin-width given with an $O(1)$-sequence can even be solved in linear FPT time $f(|\varphi|)\,|U|$, where U is the universe of the structure. In other words, bounded twin-width classes admitting a $g(\mathrm{OPT})$-approximation for the contraction sequences are tractable. It is known for (unordered) graph classes that the converse does not hold. For instance, the class of all subcubic graphs (i.e., graphs with degree at most 3) is tractable [37] but has unbounded twin-width [7]. Theorem 2 will show that, on every class of ordered graphs, a fixed-parameter approximation algorithm for the contraction sequence exists. Thus every bounded twin-width class of ordered graphs is tractable. We will also see that the converse holds for hereditary classes of ordered graphs.

3.7 Interpretations and transductions

Let Σ, Γ be signatures. A *simple interpretation* $\mathsf{I}\colon \Sigma \to \Gamma$ consists of the following Σ-formulas: a *domain* formula $\nu(x)$, and for each relation symbol $R \in \Gamma$ of arity r, a formula $\rho_R(x_1, \ldots, x_r)$. If $\mathbf{A}$ is a Σ-structure, the Γ-structure $\mathsf{I}(\mathbf{A})$ has domain $\nu(\mathbf{A}) = \{v \in A : \mathbf{A} \models \nu(v)\}$ and the interpretation of a relation symbol $R \in \Sigma$ of arity r is $\rho_R(\mathbf{A}) \cap \nu(\mathbf{A})^r$, that is:

$$R^{\mathsf{I}(\mathbf{A})} = \{(v_1, \ldots, v_r) \in \nu(\mathbf{A})^r : \mathbf{A} \models \rho_R(v_1, \ldots, v_r)\}.$$

If $\mathscr{C}$ is a class of Σ-structures then denote $\mathsf{I}(\mathscr{C}) = \{\mathsf{I}(\mathbf{A}) \mid \mathbf{A} \in \mathscr{C}\}$.

An important property of (simple) interpretations is that they can be composed: if $\mathsf{I}\colon \Sigma \to \Gamma$ and $\mathsf{J}\colon \Gamma \to \Delta$ are interpretations, then there is an interpretation $\mathsf{J} \circ \mathsf{I}\colon \Sigma \to \Delta$ (computable from I and J) such that $(\mathsf{J} \circ \mathsf{I})(\mathbf{A}) = \mathsf{J}(\mathsf{I}(\mathbf{A}))$ for every Σ-structure $\mathbf{A}$. Similarly, for every Σ-sentence φ there is a sentence $\mathsf{I}^*(\varphi)$ computable from I and φ such that for every Σ-structure $\mathbf{A}$ and we have

$$\mathsf{I}(\mathbf{A}) \models \varphi \quad \iff \quad \mathbf{A} \models \mathsf{I}^*(\varphi).$$

A class $\mathscr{C}$ *interprets* a class $\mathscr{D}$ if there is an interpretation I such that $\mathsf{I}(\mathscr{C}) \supseteq \mathscr{D}$. We say that $\mathscr{C}$ *efficiently interprets* $\mathscr{D}$ if additionally there is an algorithm which, given $\mathbf{D} \in \mathscr{D}$, computes in time polynomial in the size of $\mathbf{D}$ a structure $\mathbf{C} \in \mathscr{C}$ such that $\mathsf{I}(\mathbf{C})$ is isomorphic to $\mathbf{D}$. (A structure is represented by the size of its domain written in unary, followed by the adjacency matrices representing each of its relations.) By composition of interpretations, we conclude that if $\mathscr{C}$ efficiently interprets $\mathscr{D}$ and $\mathscr{D}$ efficiently interprets $\mathscr{E}$, then $\mathscr{C}$ efficiently interprets $\mathscr{E}$.

Efficient interpretations are a convenient way for obtaining FPT reductions, as expressed by the following straightforward lemma.

▶ **Lemma 8.** *Suppose that $\mathscr{C}$ efficiently interprets a class $\mathscr{D}$. Then there is an* FPT *reduction of* FO *model checking on $\mathscr{D}$ to* FO *model checking on $\mathscr{C}$: there is a computable function f, a constant c, and an algorithm which given a structure $\mathbf{D} \in \mathscr{D}$ and an* FO *sentence φ computes in time $f(\varphi) \cdot |\mathbf{D}|^c$ a structure $\mathbf{C} \in \mathscr{C}$ and an* FO *sentence ψ such that*

$$\mathbf{D} \models \varphi \quad \Leftrightarrow \quad \mathbf{C} \models \psi.$$

Since FO model checking on the class of all graphs is AW[$*$]-hard [17], we get:

▶ **Corollary 9.** *If $\mathscr{C}$ efficiently interprets the class of all graphs then model checking on $\mathscr{C}$ is* AW[$*$]*-hard.*

An important class of ordered graphs which efficiently interprets the class of all graphs is the class $\mathscr{M}$ of all ordered matchings. This is expressed by the following folklore result, whose proof is included in Appendix A for completeness.

▶ **Lemma 10.** *The class $\mathscr{M}$ of ordered matchings efficiently interprets the class of all graphs.*

Let $\Sigma \subseteq \Sigma^+$ be relational signatures. The *Σ-reduct* of a Σ^+-structure $\mathbf{A}$ is the structure obtained from $\mathbf{A}$ by "forgetting" all the relations not in Σ. We denote this interpretation as $\mathsf{Reduct}_\Sigma\colon \Sigma^+ \to \Sigma$, or simply Reduct, when Σ is clear from context.

A class $\mathscr{C}$ of Σ-structures *transduces* a class $\mathscr{D}$ if there is a class $\mathscr{C}^+$ of Σ^+-structures, where Σ^+ is the union of Σ and some unary relation symbols such that $\mathsf{Reduct}_\Sigma(\mathscr{C}^+) = \mathscr{C}$ and $\mathscr{C}^+$ interprets $\mathscr{D}$.

The following result follows from [8].

▶ **Theorem 11.** *Let $\mathscr{C}$ be a class of ordered, binary structures, and suppose that $\mathscr{C}$ has bounded twin-width. Then $\mathscr{C}$ does not transduce the class of all graphs.*

This result more generally holds for (non necessarily ordered) binary structures. We only state it in the ordered case, since the definition of twin-width we gave in Section 3.3 only fits ordered binary structures.

Fix a binary signature Σ containing the symbol $<$. An *atomic type* $\tau(x_1, \ldots, x_n)$ over Σ is a maximal conjunction of atomic formulas or negated atomic formulas with variables $x_1, \ldots, x_n$, which is satisfiable in some ordered Σ-structure. (It is sufficient to verify this condition for structures with n elements, since the formulas are quantifier-free.) If $\bar{a}$ is an n-tuple of elements of an ordered Σ-structure $\mathbf{A}$ then *the atomic type of* $\bar{a}$ is the unique (up to equivalence) atomic type $\tau(x_1, \ldots, x_n)$ satisfied by $\bar{a}$ in $\mathbf{A}$. For an atomic type $\tau(x, y)$ and ordered Σ-structure $\mathbf{A}$ let $\mathsf{I}_\tau(\mathbf{A})$ be the ordered graph whose domain and order are the same as in $\mathbf{A}$, and where two vertices $u < v$ are adjacent if and only if $\tau(u, v)$ holds in $\mathbf{A}$. Then I_τ is an interpretation from Σ to the signature of ordered graphs.

We formulate a standard lemma reducing the model checking problem for adjacency matrices of structures from a class $\mathscr{C}$ to the model checking problem for $\mathscr{C}$. Let us view here the adjacency matrix $M(\mathbf{A})$ of an ordered Σ-structure $\mathbf{A}$ as the matrix $|\mathbf{A}| \times |\mathbf{A}|$ matrix whose entry at position (a, b), for $a, b \in \mathbf{A}$, is the atomic type of the pair (a, b). Hence, $M(\mathbf{A})$ is a matrix over the alphabet A_Σ consisting of all atomic types $\tau(x, y)$ with two variables. See Appendix B for a proof of the lemma.

▶ **Lemma 12.** *Let $\mathscr{C}$ be a class of ordered binary structures and let $\mathcal{M} = \{M(\mathbf{A}) \mid \mathbf{A} \in \mathscr{C}\}$ be the class of adjacency matrices of structures in $\mathscr{C}$. Then there is an* FPT *reduction of the* FO *model checking problem for $\mathcal{M}$ to the* FO *model checking problem for $\mathscr{C}$. In particular, if the former is* AW[∗]*-hard, so is the latter.*

3.8 Model theory

Let $\varphi(\overline{x}, \overline{y})$ be a Σ-formula and let $\mathscr{C}$ be a class of Σ-structures. The formula φ is *independent* over $\mathscr{C}$ if for every binary relation $R \subseteq A \times B$ between two finite sets A and B there exists a Σ-structure $\mathbf{C} \in \mathscr{C}$, some tuples $(\overline{u}_a)_{a \in A}$ in $C^{|\overline{x}|}$, and $(\overline{v}_b)_{b \in B}$ in $C^{|\overline{y}|}$ such that

$$\mathbf{C} \models \varphi(\overline{u}_a, \overline{v}_b) \iff R(a, b) \qquad \text{for all } a \in A \text{ and } b \in B.$$

The class $\mathscr{C}$ is *independent* if there is a Σ-formula $\varphi(\overline{x}, \overline{y})$ that is independent over $\mathscr{C}$. Otherwise, the class $\mathscr{C}$ is *dependent* (or *NIP*, for Not the Independence Property). Note that if a class $\mathscr{C}$ interprets the class of all graphs, then it is independent.[4]

A *monadic lift* of a class $\mathscr{C}$ of Σ-structures is a class $\mathscr{C}^+$ of Σ^+-structures, where Σ^+ is the union of Σ and a set of unary relation symbols, and $\mathscr{C} = \{\mathsf{Reduct}_\Sigma(\mathbf{A}) : \mathbf{A} \in \mathscr{C}^+\}$. A class $\mathscr{C}$ of Σ-structures is *monadically dependent* (or *monadically* NIP) if every monadic lift of $\mathscr{C}$ is dependent (or NIP).

The following theorem witnesses that transductions are particularly fitting to the study of monadic dependence:

▶ **Theorem 13** (Baldwin and Shelah [2]). *A class $\mathscr{C}$ of Σ-structures is monadically dependent if and only if for every monadic lift $\mathscr{C}^+$ of $\mathscr{C}$ (in Σ^+-structures), every Σ^+-formula $\varphi(\overline{x}, \overline{y})$ with $|\overline{x}| = |\overline{y}| = 1$ is dependent over $\mathscr{C}^+$. Consequently, $\mathscr{C}$ is monadically dependent if and only if $\mathscr{C}$ does not transduce the class $\mathscr{G}$ of all finite graphs.*

3.9 Patterns

For a permutation $\sigma \in \mathfrak{S}_n$ and parameter $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$, we define the $n \times n$ matrix $F_{\mathrm{s}}(\sigma)$ with entry at row i and column j equal to (see Figure 3):

- $[\sigma(i) = j]$, *if* s *is* '$=$',
- $[\sigma(i) \neq j]$, *if* s *is* '$\neq$',
- $[i \leqslant \sigma^{-1}(j)]$, *if* s *is* '$\leqslant_R$',
- $[i \geqslant \sigma^{-1}(j)]$, *if* s *is* '$\geqslant_R$',
- $[j \leqslant \sigma(i)]$, *if* s *is* '$\leqslant_C$',
- $[j \geqslant \sigma(i)]$, *if* s *is* '$\geqslant_C$'.

Here, $[\alpha]$ is the Iverson bracket, with value 1 if α holds and 0 otherwise. Let $\mathcal{F}_{\mathrm{s}}$ denote the submatrix closure of the matrices $F_{\mathrm{s}}(\sigma)$, for all permutations σ.

We can also define classes $\mathcal{F}_{\mathrm{s}}$ for $\mathrm{s} \in \{<_R, >_R, <_C, >_C\}$ analogously as above, but replacing the non-strict inequalities $\leqslant$ and $\geqslant$ by the strict variants $<$ and $>$. While changing the subscript s in $F_{\mathrm{s}}(\sigma)$

[4] The converse also holds if the interpretations can use constant symbols, and we can take induced substructures after performing the interpretation, see [39].

from a non-strict inequality to its strict variant affects the matrix entries, we nevertheless have:

$$\mathcal{F}_{<R} = \mathcal{F}_{\leqslant R}, \quad \mathcal{F}_{>R} = \mathcal{F}_{\geqslant R}, \quad \mathcal{F}_{<C} = \mathcal{F}_{\leqslant C}, \quad \mathcal{F}_{>C} = \mathcal{F}_{\geqslant C}.$$

A class $\mathcal{M}$ of 0, 1-matrices is *pattern-avoiding* if it does not include any of the six matrix classes $\mathcal{F}_{\mathrm{s}}$, for $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$. We now lift this notion to arbitrary alphabets.

For a finite alphabet A, letter $a \in A$, and matrix M over A, the *a-selection* of M is the 0, 1-matrix $s_a(M)$ obtained from M by replacing each occurrence of a by 1 and each other letter by 0. The a-selection of a class $\mathcal{M}$ of matrices is the class $s_a(\mathcal{M})$ of a-selections of matrices in $\mathcal{M}$. Say that $\mathcal{M}$ is *pattern-avoiding* if every its a-selection $s_a(\mathcal{M})$ is pattern-avoiding.

In the introduction we define the 25 classes of ordered graphs $\mathscr{P}$ and $\mathscr{M}_{\mathrm{s},\lambda,\rho}$, for $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ and $\lambda, \rho \in \{0, 1\}$. The parameters $=, \neq, \leqslant_R, \geqslant_R, \leqslant_C$, and $\geqslant_C$ used for matrices are renamed to $=, \neq, \leqslant_l, \geqslant_l, \leqslant_r$, and $\geqslant_r$ in the case of ordered graphs, since rows and columns are interpreted as left and right vertices, respectively. The classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ can be alternatively defined as follows.

Let H be an ordered matching with vertices $a_1 < \ldots < a_n < b_1 < \ldots < b_n$, so that there is some $\sigma \in \mathfrak{S}_n$ such that a_i is matched with $b_{\sigma(i)}$, for $1 \leqslant i \leqslant n$. Then for $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$, define an ordered graph $H[\mathrm{s}, \lambda, \rho]$ with vertices $a_1 < \ldots < a_n < b_1 < \ldots < b_n$ such that $[E(a_i, b_j)]$ (the truth value of the adjacency between a_i and b_j) is equal to:

$$[\sigma(i) = j], \quad [\sigma(i) \neq j], \quad [\sigma(i) \leqslant j], \quad [\sigma(i) \geqslant j], \quad [i \leqslant \sigma^{-1}(j)], \text{ or } [i \geqslant \sigma^{-1}(j)],$$

depending on the parameter $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$, and for $1 \leqslant i < j \leqslant n$, $[E(a_i, a_j)] = \lambda$ and $[E(b_i, b_j)] = \rho$. Note that $([E(a_i, b_j)])_{1 \leqslant i,j \leqslant n} = F_{\mathrm{s}}(\sigma)$, where s is now treated as an element of $\{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$.

The class $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ is the hereditary closure of the class of all ordered graphs $H[\mathrm{s}, \lambda, \rho]$, where H is an ordered matching. The class $\mathscr{P}$ is the class of all ordered permutation graphs, and satisfies the following properties (see Section 3.1).

▶ **Lemma 14.** *The class $\mathscr{P}$ is hereditary, closed under edge complements, and has growth $n!$.*

3.10 Ramsey theory

We recall Ramsey's theorem.

▶ **Theorem 15** (Ramsey's theorem [35]). *There exists a function $\mathrm{R}_\cdot(\cdot) : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that for every $k \geqslant 1$, $t \geqslant 1$ the complete graph $K_{\mathrm{R}_t(k)}$ with edges colored by t distinct colors contains a monochromatic clique on k vertices, i.e., a clique whose edges all have the same color.*

For every $p \geqslant 0$ we will denote with $\mathrm{R}_t^{(p)}(\cdot)$ the function $\mathrm{R}_t(\cdot)$ iterated p times. A well-known variant of Theorem 15 for complete bipartite graphs is the following:

▶ **Theorem 16** (Bipartite Ramsey's theorem). *There exists a function $\mathrm{b}_\cdot(\cdot) : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that for every $k \geqslant 1$, $t \geqslant 1$ the complete graph $K_{\mathrm{b}_t(k),\mathrm{b}_t(k)}$ with edges colored by t distinct colors contains a monochromatic biclique $K_{k,k}$.*

The *order type* of a pair (x, y) of elements of a totally ordered set is the integer $\mathrm{ot}(x, y)$ defined by

$$\mathrm{ot}(x, y) = \begin{cases} -1 & \text{if } x > y \\ 0 & \text{if } x = y \\ 1 & \text{if } x < y. \end{cases}$$

For structures with two orders, we will use a convenient specific result from Ramsey theory, which is a special case of the so-called *product Ramsey theorem* (see e.g. Proposition 3 in [5] in the special case of the full product of two copies of $(\mathbb{Q}, <)$. See also the historical comment following it).

For a set with two total orders $(X, <_1, <_2)$, let $\mathrm{ot}_1(x, y)$ denote the order type of x, y with respect to $<_1$, while $\mathrm{ot}_2(x, y)$ the order type with respect to $<_2$.

▶ **Lemma 17.** *Fix a finite set of colors Γ with $t := |\Gamma|$. There exists a function $\mathrm{g}_t(\cdot)\ : \mathbb{N} \to \mathbb{N}$ such that for every finite set with two total orders $\mathbf{M} = ([k], <_1, <_2)$ there is another finite set with two total orders $\mathbf{N} = ([N], <_1, <_2)$ where $N = \mathrm{g}_t(k)$, such that for every coloring $c\colon [N]^2 \to \Gamma$ there is a substructure $\mathbf{M}'$ of $\mathbf{N}$ isomorphic to $\mathbf{M}$ such that $c(p,q)$ depends only on $\mathrm{ot}_1(p,q)$ and $\mathrm{ot}_2(p,q)$, for all distinct $p,q \in \mathbf{M}'$. More precisely, $c(p,q) = \gamma(\mathrm{ot}_1(p,q), \mathrm{ot}_2(p,q))$ holds for some function $\gamma\colon \{-1,0,1\}^2 \to \Gamma$.*

Lemma 17 translates to the following statement on permutations.

▶ **Lemma 18.** *Fix a finite set of colors Γ. For every $k \geqslant 1$ and permutation $\sigma \in \mathfrak{S}_k$ there is $N \geqslant 1$ and a permutation $\pi \in \mathfrak{S}_N$ such that for every coloring $c\colon [N]^2 \to \Gamma$ there is a set $U \subseteq [N]$ of size k such that σ is the subpermutation of π induced by U, and $c(i,j)$ depends only on $\mathrm{ot}(i,j)$ and $\mathrm{ot}(\pi(i),\pi(j))$, for all $i,j \in U$.*

4 Effective equivalence of bounded twin-width and no large rich division

In this section we show the equivalence between (i) and (ix). As a by-product, we obtain an $f(\mathrm{OPT})$-approximation algorithm for the twin-width of matrices, or ordered graphs. We first show that a large rich division implies large twin-width. This direction is crucial for the algorithm but *not* for the main circuit of implications.

▶ **Lemma 19.** *If M has a $2k(k+1)$-rich division $\mathcal{D}$, then $\mathrm{tww}(M) > k$.*

Proof. We prove the contrapositive. Let M be a matrix of twin-width at most k. In particular, M admits a (k,k)-sequence $\mathcal{P}_1, \ldots, \mathcal{P}_{n+m-1}$. Let $\mathcal{D}$ be any division of M. We want to show that $\mathcal{D}$ is *not* $2k(k+1)$-rich.

Let t be the smallest index such that either a part R_i of $\mathcal{P}^R_t$ intersects three parts of $\mathcal{D}^R$, or a part C_j of $\mathcal{P}^C_t$ intersects three parts of $\mathcal{D}^C$. Without loss of generality we can assume that $C_j \in \mathcal{P}^C_t$ intersects three parts C'_a, C'_b, C'_c of $\mathcal{D}^C$, with $a < b < c$ where the parts $C'_1, \ldots, C'_d$ of the division $\mathcal{D}$ are ordered from left to right. Since $\mathcal{P}^C_t$ is a k-overlapping partition, the subset S, consisting of the parts of $\mathcal{P}^C_t$ intersecting C'_b, has size at most $k+1$. Indeed, S contains C_j plus at most k parts which C_j is in conflict with.

Here a part R'_s of $\mathcal{D}^R$ is called *red* if there exist a part R_i of $\mathcal{P}^R_t$ intersecting R'_s and a part C_z in S such that the submatrix $R_i \cap C_z$ is not constant (see Figure 8). We then say that C_z is a *witness* of R'_s being red. Let $N \subseteq R$ be the subset of rows *not* in a red part of $\mathcal{D}^R$. Note that for every part $C_z \in S$, the submatrix $N \cap C_z$ consists of the same column vector repeated $|C_z|$ times. Therefore $N \cap C'_b$ has at most $k+1$ distinct column vectors.

Besides, the number of red parts witnessed by $C_z \in S$ is at most $2k$. This is because the number of non-constant submatrices $R_i \cap C_z$, with $R_i \in \mathcal{P}^R_t$, is at most k (since $\mathcal{P}_1, \ldots, \mathcal{P}_{n+m-1}$ is a (k,k)-sequence) and because every R_i intersects at most two parts of $\mathcal{D}^R$ (by definition of t). Hence the total number of red parts is at most $2k|S|$, thus at most $2k(k+1)$. Consequently, there is a subset X of at most $2k(k+1)$ parts of $\mathcal{D}^R$, namely the red parts, and a part C'_b of $\mathcal{D}^C$ such that $(R \setminus \cup X) \cap C'_b = N \cap C'_b$ consists of at most $k+1$ distinct column vectors. Thus $\mathcal{D}$ is not a $2k(k+1)$-rich division. ◀

Our main algorithmic result is that approximating the twin-width of matrices (or ordered graphs) is FPT. Let us observe that this remains a challenging open problem for (unordered) graphs. We finally state Theorem 2 in the language of matrices, since rich divisions are only natural in that setting. The previous statement of Theorem 2, for ordered binary structures, readily follows.

Indeed recall that the twin-width of an ordered binary structure is defined as the twin-width of its matrix encoding. Besides a contraction sequence for the matrix can be turned into a contraction sequence for the ordered binary structure. Every time the i-th and j-th, say, column parts are merged, we symmetrically merge the i-th and j-th row parts. Because our matrix encoding is mixed-symmetric, this produces a sequence with the same error value (see [8, Theorem 14]). The now-symmetric sequence can then be interpreted as contracting the vertices of a graph, or more generally, the domain elements of a binary structure.

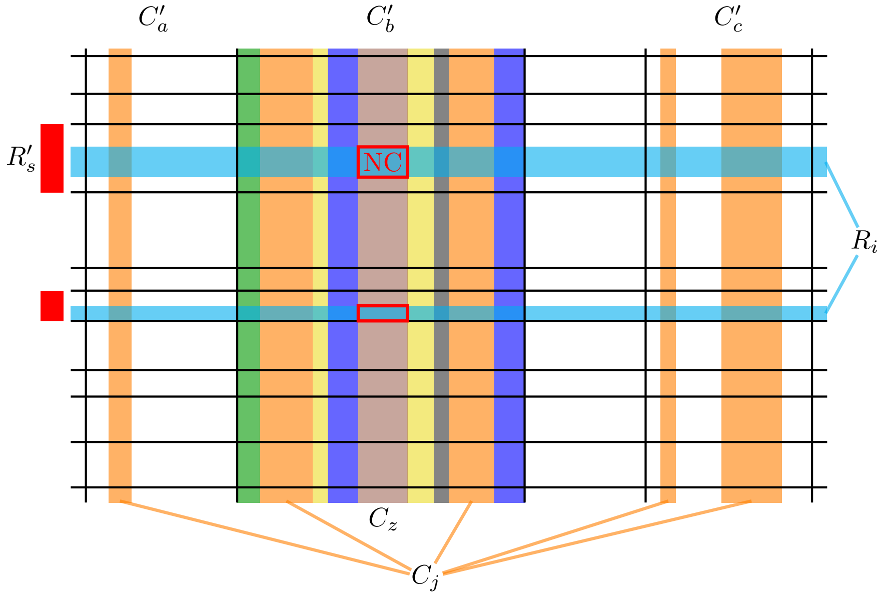

Figure 8 The division D in black. The column part $C_j \in \mathcal{P}_t^C$, first to intersect three division parts, in orange. Two row parts of $\mathcal{D}$ turn red because of the non-constant submatrix $C_z \cap R_i$, with $C_z \in S$ and $R_i \in \mathcal{D}^R$. After removal of the at most $2k|S|$ red parts, $|S| \leqslant k+1$ bounds the number of distinct columns.

▶ **Theorem 2.** *Given as input an $n \times m$ matrix M over a finite alphabet A, and an integer k, there is an $2^{2^{O(k^2 \log k)}}(n+m)^{O(1)}$ time algorithm which returns*

- *either a $2k(k+1)$-rich division of M, certifying that $\text{tww}(M) > k$,*
- *or an $(|A|^{O(k^4)}, |A|^{O(k^4)})$-sequence, certifying that $\text{tww}(M) = |A|^{O(k^4)}$.*

Proof. We try to construct a division sequence $\mathcal{D}_1, \ldots, \mathcal{D}_{n+m-1}$ of M such that every $\mathcal{D}_i$ satisfies the following properties $\mathscr{P}^R$ and $\mathscr{P}^C$. Let r be equal to $4k(k+1)+1$.

- $\mathscr{P}^R$: For every part R_a of $\mathcal{D}_i^R$, there is a set Y of at most r parts of $\mathcal{D}_i^C$, such that the submatrix $R_a \cap (C \setminus \cup Y)$ has at most $r-1$ distinct row vectors.
- $\mathscr{P}^C$: For every part C_b of $\mathcal{D}_i^C$, there is a set X of at most r parts of $\mathcal{D}_i^R$, such that the submatrix $(R \setminus \cup X) \cap C_b$ has at most $r-1$ distinct column vectors.

The algorithm is greedy: Whenever we can merge two consecutive row parts or two consecutive column parts in $\mathcal{D}_i$ so that the above properties are preserved, we do so, and obtain $\mathcal{D}_{i+1}$. We first show that checking properties $\mathscr{P}^R$ and $\mathscr{P}^C$ can be done in fixed-parameter time.

▶ **Lemma 20.** *Whether $\mathscr{P}^R$, or $\mathscr{P}^C$, holds can be decided in time $2^{2^{O(k^2 \log k)}}(n+m)^{O(1)}$.*

Proof. We show the lemma with $\mathscr{P}^R$, since the case of $\mathscr{P}^C$ is symmetric. For every $R_a \in \mathcal{D}_i^R$, we denote by $\mathscr{P}^R(R_a)$ the fact that R_a satisfies the condition $\mathscr{P}^R$ starting at "there is a set Y." If one can check $\mathscr{P}^R(R_a)$ in time T, one can thus check $\mathscr{P}^R$ and $\mathscr{P}^C$ for the current division $\mathcal{D}_i$ in time $(|\mathcal{D}_i^R|+|\mathcal{D}_i^C|)T \leqslant (n+m)T$.

To decide $\mathscr{P}^R(R_a)$, we initialize the set Y with all the column parts $C_b \in \mathcal{D}_i^C$ such that the zone $R_a \cap C_b$ contains more than $r-1$ distinct rows. Indeed these parts *have to* be in Y. At this point, if $R_a \cap (C \setminus \cup Y)$ has more than $(r-1)^{r+1}$ distinct rows, then $\mathscr{P}^R(R_a)$ is false. Indeed, each further removal of a column part divides the number of distinct rows in R_a by at most $r-1$. Thus after at most r further removals, more than $r-1$ distinct rows would remain.

Let us suppose instead that $R_a \cap (C \setminus \cup Y)$ has at most $(r-1)^{r+1}$ distinct rows. We keep one representative for each distinct row. For every $C_b \in \mathcal{D}_i^C \setminus Y$, the number of distinct *columns* in zone $R_a \cap C_b$ is at most $|A|^{r-1}$. In each of these zones, we keep only one representative for every occurring column vector. Now every zone of R_a has dimension at most $(r-1)^{r+1} \times |A|^{r-1}$. Therefore the maximum number of distinct zones is $\exp(\exp(O(r \log r))) = \exp(\exp(O(k^2 \log k)))$.

If a same zone Z is repeated in R_a more than r times, at least one occurrence of the zone will not be included in Y. In that case, putting copies of Z in Y is pointless: it eventually does not decrease the number of distinct rows. Thus if that happens, we keep exactly $r+1$ copies of Z. Now R_a has at most $(r+1) \cdot \exp(\exp(O(k^2 \log k))) = \exp(\exp(O(k^2 \log k)))$ zones. We can try out all $\exp(\exp(O(k^2 \log k)))^r$, that is, $\exp(\exp(O(k^2 \log k)))$ possibilities for the set Y, and conclude whether or not one of them works. ◀

Two cases can arise.

Case 1. The algorithm terminates on some division $\mathcal{D}_i$ and no merge is possible.
Let us assume that $\mathcal{D}_i^R := \{R_1, \dots, R_s\}$ and $\mathcal{D}_i^C := \{C_1, \dots, C_t\}$, where the parts are ordered by increasing vector indices. We consider the division $\mathcal{D}$ of M obtained by merging in $\mathcal{D}_i$ the pairs $\{R_{2a-1}, R_{2a}\}$ and $\{C_{2b-1}, C_{2b}\}$, for every $1 \leqslant a \leqslant \lfloor s/2 \rfloor$ and $1 \leqslant b \leqslant \lfloor t/2 \rfloor$. Let C'_j be any column part of $\mathcal{D}^C$. Since the algorithm has stopped, for every set X of at most $(r-1)/2$ parts of $\mathcal{D}^R$, the matrix $(R \setminus \cup X) \cap C'_j$ has at least r distinct (column) vectors. This is because $(r-1)/2$ parts of $\mathcal{D}^R$ corresponds to at most $r-1$ parts of $\mathcal{D}_i^R$. The same applies to the row parts, so we deduce that $\mathcal{D}$ is $(r-1)/2$-rich, that is, $2k(k+1)$-rich. Therefore, by Lemma 19, M has twin-width greater than k.

Case 2. The algorithm terminates with a full sequence $\mathcal{D}_1, \dots, \mathcal{D}_{n+m-1}$.
Given a division $\mathcal{D}_i$ with $\mathcal{D}_i^R := \{R_1, \dots, R_s\}$ and $\mathcal{D}_i^C := \{C_1, \dots, C_t\}$, we now define a partition $\mathcal{P}_i$ that refines $\mathcal{D}_i$ and has small error value. To do so, we fix a, say, column part C_j and show how to partition it further in $\mathcal{P}_i$.

By assumption on $\mathcal{D}_i$, there exists a subset X of at most r parts of $\mathcal{D}_i^R$ such that $(R \setminus \cup X) \cap C_j$ has less than r distinct column vectors. We now denote by F the set of parts R_a of $\mathcal{D}_i^R$ such that the zone $R_a \cap C_j$ has at least r distinct rows and r distinct columns. Such a zone is called *full*. Observe that $F \subseteq X$. Moreover, for every R_a in $X \setminus F$, the total number of distinct column vectors in $R_a \cap C_j$ is at most $\max(r, |A|^{r-1}) = |A|^{r-1}$, assuming that the alphabet A has at least two letters. Indeed, if the number of distinct columns in $R_a \cap C_j$ is at least r, then the number of distinct rows is at most $r-1$.

In particular, the total number of distinct column vectors in $(R \setminus \cup F) \cap C_j$ is at most $w := r(|A|^{r-1})^r$; a multiplicative factor of $|A|^{r-1}$ for each of the at most r zones $R_a \in X \setminus F$, and a multiplicative factor of r for $(R \setminus \cup X) \cap C_j$. We partition the columns of C_j accordingly to their subvector in $(R \setminus \cup F) \cap C_j$ (by grouping columns with equal subvectors together). The partition $\mathcal{P}_i$ is obtained by refining, as described for C_j, all column parts and all row parts of $\mathcal{D}_i$.

By construction, $\mathcal{P}_i$ is a refinement of $\mathcal{P}_{i+1}$ since every full zone of $\mathcal{D}_i$ remains full in $\mathcal{D}_{i+1}$. Hence if two columns belong to the same part of $\mathcal{P}_i$, they continue belonging to the same part of $\mathcal{P}_{i+1}$. Besides, $\mathcal{P}_i$ is a w-overlapping partition of M, and its error value is at most $r \cdot w$ since non-constant zones can only occur in full zones (at most r per part of $\mathcal{D}_i$), which are further partitioned at most w times in $\mathcal{P}_i$. To finally get a contraction sequence, we greedily merge parts to fill the intermediate partitions between $\mathcal{P}_i$ and $\mathcal{P}_{i+1}$. Note that all intermediate refinements of $\mathcal{P}_{i+1}$ are w-overlapping partitions. Moreover the error value of a column part does not exceed $r \cdot w$. Finally the error value of a row part can increase during the intermediate steps by at most $2w$. All in all, we get a $(w, (r+2) \cdot w)$-sequence. This implies that M has twin-width at most $(r+2) \cdot w = |A|^{O(k^4)}$.

The running time of the overall algorithm follows from Lemma 20. ◀

The approximation ratio, of $2^{O(\mathrm{OPT}^4)}$, can be analyzed more carefully by observing that bounded twin-width implies bounded VC dimension. Then the threshold $|A|^{r-1}$ can be replaced by r^d, where d upperbounds the VC dimension. As a direct corollary of our algorithm, if the matrix M does not admit any large rich division, the only possible outcome is a contraction sequence. Considering the size of A as an absolute constant, we thus obtain the following.

▶ **Theorem 21.** *If M has no r-rich division, then* $\mathrm{tww}(M) = 2^{O(r^2)}$.

This is the direction which is important for the circuit of implications. The algorithm of Theorem 2 further implies that Theorem 21 is effective.

5 Large rich divisions imply large rank divisions

In this section we show how to extract a large rank division from a huge rich division. We remind the reader that a rank-k division is a k-division in which every zone has at least k distinct rows or at least k distinct columns. We recall that mt is the Marcus-Tardos bound of Theorem 6.

▶ **Theorem 22.** *Let A be a finite set, and $K := |A|^{k\,\mathrm{mt}(k^2)}$. Every A-matrix M with a K-rich division $\mathcal{D}$ has a rank-k division.*

Proof. Without loss of generality, we can assume that $\mathcal{D}^C$ has size at least the size of $\mathcal{D}^R$. We color *red* every zone of $\mathcal{D}$ which has at least k distinct rows or at least k distinct columns. We now color *blue* a zone $R_i \cap C_j$ of $\mathcal{D}$ if it contains a row vector r (of length $|C_j|$) which does not appear in any non-red zone $R_{i'} \cap C_j$ with $i' < i$. We then call r a *blue witness* of $R_i \cap C_j$.

Let us now denote by U_j the subset of $\mathcal{D}^R$ such that every zone $R_i \cap C_j$ with $R_i \in U_j$ is *uncolored*, i.e., neither red nor blue. Since the division $\mathcal{D}$ is K-rich, if the number of *colored* (i.e., red or blue) zones $R_i \cap C_j$ is less than K, the matrix $(\cup U_j) \cap C_j$ has at least K distinct column vectors. So $(\cup U_j) \cap C_j$ has at least $\log_{|A|} K = k\,\mathrm{mt}(k^2)$ distinct *row* vectors. By design, every row vector appearing in some uncolored zone $R_i \cap C_j$ must appear in some blue zone $R_{i'} \cap C_j$ with $i' < i$. Therefore at least $k\,\mathrm{mt}(k^2)$ distinct row vectors must appear in some blue zones within column part C_j. Since a blue zone contains less than k distinct row vectors (otherwise it would be a red zone), there are, in that case, at least $k\,\mathrm{mt}(k^2)/k = \mathrm{mt}(k^2)$ blue zones within C_j. Therefore in any case, the number of colored zones $R_i \cap C_j$ is at least $\mathrm{mt}(k^2)$ per C_j.

Thus, by Theorem 6, we can find D' a $k^2 \times k^2$ division of M, coarsening D, with at least one colored zone of D in each *cell* of D'. Now we consider D'' the $k \times k$ division of M, coarsening D', where each *supercell* of D'' corresponds a $k \times k$ square block of cells of D' (see Figure 9). Our goal is to show that every supercell Z of D'' has at least k distinct rows or k distinct columns. If the supercell Z contains a red zone of D, the property immediately holds for Z. If not, each of the $k \times k$ cells of D' within the supercell Z contains at least one blue zone of D. Let $Z_{i,j}$ be the cell in the i-th row block and j-th column block of hypercell Z, for every $i, j \in [k]$. Consider the diagonal cells $Z_{i,i}$ ($i \in [k]$) of D' within the supercell Z. In each of them, there is at least one blue zone witnessed by a row vector, say, $\tilde{r_i}$. Let r_i be the prolongation of $\tilde{r_i}$ up until the two vertical limits of Z. We claim that every r_i (with $i \in [k]$) is distinct. Indeed by definition of a blue witness, if $i < j$, $\tilde{r_j}$ is different from all the row vectors below it, in particular from r_i restricted to these columns. So Z has at least k distinct row vectors. ◀

6 Rank Latin divisions

In this section, we show a Ramsey-like result which establishes that every (hereditary) matrix class with unbounded grid rank can encode all the n-permutations with some of its $2n \times 2n$ matrices. In particular and in light of the previous sections, this proves the small conjecture for ordered graphs.

We recall that a rank-k d-division of a matrix M is a d-by-d division of M whose every zone has rank at least k, and *rank-k division* is a short-hand for *rank-k k-division*. Then a matrix class $\mathcal{M}$ has bounded grid rank if there is an integer k such that no matrix of $\mathcal{M}$ admits a rank-k division.

Henceforth it will be more convenient to only work with $0,1$-matrices. Lemma 23 allows us to do so. It directly implies that for every matrix class $\mathcal{M}$ over a finite alphabet A, if $\mathcal{M}$ unbounded grid rank then there is some $a \in A$ such that the class of $0,1$-matrices $s_a(\mathcal{M})$ obtained from $\mathcal{M}$ by replacing a by 1 and each other letter by 0, has unbounded grid rank, too.

▶ **Lemma 23.** *Let A be a finite alphabet, $k, d \in \mathbb{N}$ and let M be any A-matrix admitting a rank-K D-division for $K = (k-1)^{|A|-1} + 1$ and $D = \mathrm{b}_{|A|}(d)$. Then for some $a \in A$ the $0,1$-matrix $s_a(M)$ admits a rank-k d-division. In particular, if a class $\mathcal{M}$ of A-matrices has unbounded grid rank then there is some $a \in A$ such that $s_a(\mathcal{M})$ has unbounded grid rank.*

Proof. First note that if an A-matrix N has at least K distinct rows or columns then there is some $a \in A$ such that $s_a(N)$ has at least k distinct rows or columns. Indeed, suppose that $s_a(N)$ has at most $k-1$

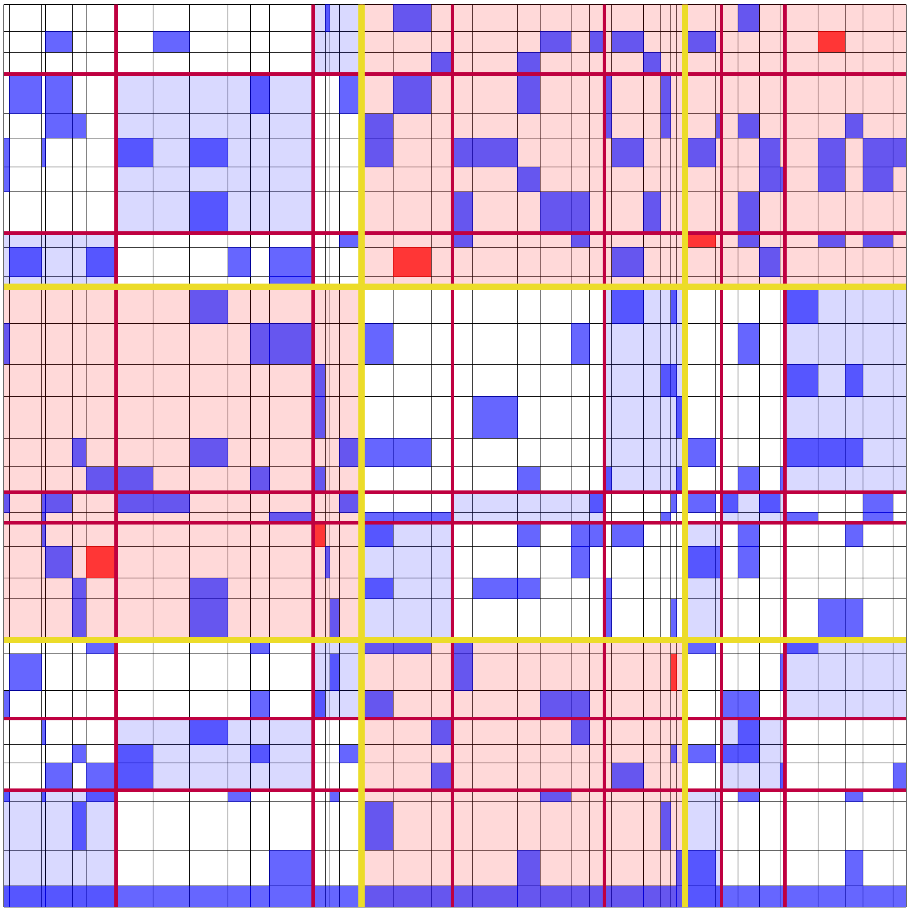

Figure 9 In black (purple, and yellow), the rich division D. In purple (and yellow), the Marcus-Tardos division D' with at least one colored zone of D per cell. In yellow, the rank-k division D''. Each supercell of D'' has large rank, either because it contains a red zone (light red) or because it has a diagonal of cells of D' with a blue zone (light blue).

distinct rows, for each $a \in A$. Let $a_0 \in A$ be any letter. Each row v of N (treated as a matrix with one row) is uniquely determined by the tuple $(s_a(v))_{a \in A \setminus \{a_0\}}$ of all its a-selections with $a \neq a_0$. As $s_a(v)$ is a row of $s_a(N)$, there are $k-1$ possible values for $s_a(v)$, by assumption. Hence, $(s_a(v))_{a \in A \setminus \{a_0\}}$ ranges over a set of size at most $(k-1)^{|A|-1} = K-1$, which yields the conclusion.

We now prove the statement in the lemma. Let $\mathcal{D}$ be a rank-K division of M. By definition in every cell C of $\mathcal{D}$ there are (at least) K distinct row or K distinct column vectors. By the above, there is some $a \in A$ such that $s_a(C)$ has at least k distinct row or k distinct column vectors. We label C by a.

There are only $|A|$ possible labels for the cells and $D = \mathrm{b}_{|A|}(d)$. Thus by Theorem 16, there is a fixed letter, say, $a \in A$ and a d-division $\mathcal{D}'$ coarsening the D-division $\mathcal{D}$ such that each cell of $\mathcal{D}'$ contains a cell of $\mathcal{D}$ labeled by a. By construction, $\mathcal{D}'$ is a rank-k d-division of $s_a(M)$. ◀

Let $\mathbf{I}_k$ be the $k \times k$ identity matrix, and $\mathbf{1}_k$, $\mathbf{0}_k$, $\mathbf{U}_k$, and $\mathbf{L}_k$ be the $k \times k$ 0,1-matrices that are full 1, full 0, full-1 upper triangular, and full-1 lower triangular, respectively. More precisely, the $k \times k$ full-1 upper (resp. lower) triangular matrix is a 0,1-matrix with a 1 at position $(i,j) \in [k]^2$ if and only if $i \leqslant j$ (resp. $i \geqslant j$). Let A^{M} be the vertical mirror of matrix A, that is, its reflection about a vertical line separating the matrix in two equal parts.[5] The following Ramsey-like result states that every $0,1$-matrix with huge rank (or equivalently a huge number of distinct row or column vectors) admits a regular matrix with large rank.

[5] i.e., column $\lceil n/2 \rceil$ if the number n of columns is odd, and between columns $n/2$ and $n/2+1$ if n is even

▶ **Theorem 24.** *There is a function $T : \mathbb{N}^+ \to \mathbb{N}^+$ such that for every natural k, every matrix with at least $T(k)$ rows or at least $T(k)$ columns contains as a submatrix one of the following $k \times k$ matrices:* $\mathbf{I}_k$, $\mathbf{1}_k - \mathbf{I}_k$, $\mathbf{U}_k$, $\mathbf{L}_k$, $\mathbf{I}_k^{\mathrm{M}}$, $(\mathbf{1}_k - \mathbf{I}_k)^{\mathrm{M}}$, $\mathbf{U}_k^{\mathrm{M}}$, $\mathbf{L}_k^{\mathrm{M}}$.

The previous theorem is a folklore result. For instance, it can be readily derived from Gravier et al. [24] or from [16, Corollary 2.4.] combined with the Erdős-Szekeres theorem.

Let $\mathcal{N}_k$ be the set of the eight matrices of Theorem 24. The first four matrices are called *diagonal*, and the last four (those defined by vertical mirror) are called *anti-diagonal*. By Theorem 24, if a matrix class $\mathcal{M}$ has unbounded grid rank, then one can find in $\mathcal{M}$ arbitrarily large divisions with a matrix of $\mathcal{N}_k$ as submatrix in each zone of the division, for arbitrarily large k. We want to acquire more control on the horizontal-vertical interactions between these submatrices of $\mathcal{N}_k$. We will prove that in large rank divisions, one can find so-called *rank Latin divisions*.

An *embedded submatrix* M' of a matrix M is a submatrix of M together with the implicit information of the position of M' in M. In particular, we will denote by $\text{rows}(M')$, respectively $\text{cols}(M')$ the rows of M, respectively columns of M, intersecting precisely at M'. The argument of $\text{rows}(\cdot)$ or $\text{cols}(\cdot)$ is implicitly cast in an embedded submatrix of M. In particular, $\text{rows}(M)$ simply denotes the set of rows of M. A *contiguous* (embedded) submatrix is defined by a *zone*, that is, a set of consecutive rows and a set of consecutive columns. The (i,j)*-cell* of a d-division $\mathcal{D}$, for any $i,j \in [d]$, is the zone formed by the i-th row block and the j-th column block of $\mathcal{D}$. We will often denote that zone by $\mathcal{D}_{i,j}$.

A *rank-k Latin d-division* of a matrix M is a d-division $\mathcal{D}$ of M such that for every $i,j \in [d]$ there is a contiguous embedded submatrix $M_{i,j} \in \mathcal{N}_k$ in the (i,j)-cell of $\mathcal{D}$ satisfying:

- $\{\text{rows}(M_{i,j})\}_{i,j}$ partitions $\text{rows}(M)$, and $\{\text{cols}(M_{i,j})\}_{i,j}$ partitions $\text{cols}(M)$.
- $\text{rows}(M_{i,j}) \cap \text{cols}(M_{i',j'})$ equals $\mathbf{1}_k$ or $\mathbf{0}_k$, whenever $(i,j) \neq (i',j')$.

Note that since the submatrices $M_{i,j}$ are supposed contiguous, the partition is necessarily a 0-overlapping partition, hence a division. A *rank-k pre-Latin d-division* is the same, except that the second item need not be satisfied.

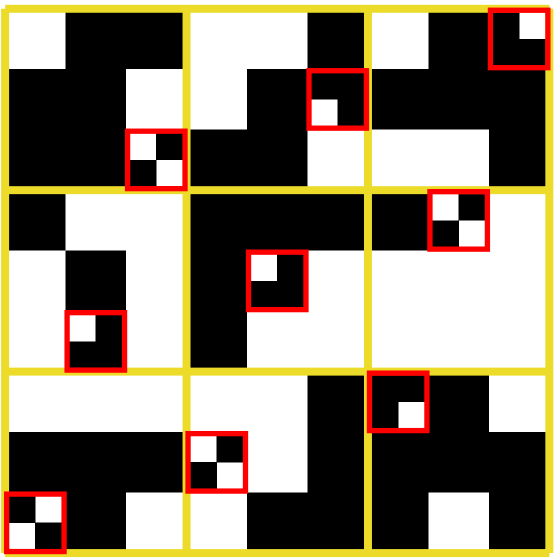

Figure 10 A 18×18 0,1-matrix with a rank-2 Latin 3-division (in yellow) where 1 entries are depicted in black, 0 entries, in white, and every $M_{i,j}$ is highlighted in red.

We can now state our technical lemma.

▶ **Lemma 25.** *For every positive integer k, there is an integer K such that every 0,1-matrix M with a rank-K division has a submatrix with a rank-k Latin division.*

Proof. We start by showing the following claim, a first step in the global cleaning process of Lemma 25. We recall that $T(\cdot)$ is the function of Theorem 24.

▷ Claim 26. Let M be a $0,1$-matrix with a rank-$T(\kappa)$ d^2-division $\mathcal{D}$. There is a $\kappa d^2 \times \kappa d^2$ submatrix $\tilde{M}$ of M with a rank-κ pre-Latin d-division; i.e., a d-division $\mathcal{D}'$, coarsening $\mathcal{D}$, such that the (i,j)-cell of $\mathcal{D}'$ contains $M_{i,j} \in \mathcal{N}_\kappa$ as a contiguous submatrix, $\{\text{rows}(M_{i,j})\}_{i,j\in[d]}$ partitions $\text{rows}(\tilde{M})$, and $\{\text{cols}(M_{i,j})\}_{i,j\in[d]}$, $\text{cols}(\tilde{M})$.

Proof of the claim. Let $\mathcal{D}^R$ be $(R_1, \ldots, R_{d^2})$ and, $\mathcal{D}^C$ be $(C_1, \ldots, C_{d^2})$. Let $\mathcal{D}'$ be the coarsening of $\mathcal{D}$ defined by

$$\mathcal{D}'^R := (\bigcup_{i\in[d]} R_i, \bigcup_{i\in[d+1,2d]} R_i, \ldots, \bigcup_{i\in[(d-1)d+1,d^2]} R_i), \text{ and}$$

$$\mathcal{D}'^C := (\bigcup_{j\in[d]} C_j, \bigcup_{j\in[d+1,2d]} C_j, \ldots, \bigcup_{j\in[(d-1)d+1,d^2]} C_j).$$

By Theorem 24, each cell of $\mathcal{D}$ contains a submatrix in $\mathcal{N}_\kappa$. Thus there are d^2 such submatrices in each cell of $\mathcal{D}'$. For every $i,j \in [d]$, we keep in $\tilde{M}$ the κ rows and κ columns of a single submatrix of $\mathcal{N}_\kappa$ in the (i,j)-cell of $\mathcal{D}'$, and more precisely, one $M_{i,j}$ in the $(j+(i-1)d, i+(j-1)d)$-cell of $\mathcal{D}$. In other words, we keep in the (i,j)-cell of $\mathcal{D}'$, a submatrix of $\mathcal{N}_\kappa$ in the (j,i)-cell of $\mathcal{D}$ restricted to $\mathcal{D}'$.[6] The submatrices $M_{i,j}$ are contiguous in $\tilde{M}$. The set $\{\text{rows}(M_{i,j})\}_{i,j\in[d]}$ partitions $\text{rows}(\tilde{M})$ since $j+(i-1)d$ describes $[d]^2$ when $i \times j$ describes $[d] \times [d]$. Similarly $\{\text{cols}(M_{i,j})\}_{i,j\in[d]}$ partitions $\text{cols}(\tilde{M})$. ◀

We recall that $\text{b}_2(k)$ is the minimum integer b such that every 2-edge coloring of the biclique $K_{b,b}$ contains a monochromatic $K_{k,k}$. We set $\text{b}_2^{(1)}(k) := \text{b}_2(k)$, and for every integer $s \geqslant 2$, we denote by $\text{b}_2^{(s)}(k)$, the minimum integer b such that every 2-edge coloring of $K_{b,b}$ contains a monochromatic $K_{q,q}$ with $q = \text{b}_2^{(s-1)}(k)$. We set $\kappa := \text{b}_2^{(k^4-k^2)}(k)$ and $K := \max(T(\kappa), k^2) = T(\kappa)$, so that applying Claim 26 on a rank-K division (hence in particular a rank-$T(\kappa)$ k^2-division) gives a rank-κ pre-Latin k-division, with the k^2 submatrices of $\mathcal{N}_\kappa$ denoted by $M_{i,j}$ for $i,j \in [k]$.

At this point the zones $\text{rows}(M_{i,j}) \cap \text{cols}(M_{i',j'})$, with $(i,j) \neq (i',j')$, are arbitrary. We now gradually extract a subset of k rows and the k *corresponding columns* (i.e., the columns crossing at the diagonal if $M_{i,j}$ is diagonal, or at the anti-diagonal if $M_{i,j}$ is anti-diagonal) within each $M_{i,j}$, to turn the rank pre-Latin division into a rank Latin division. To keep our notation simple, we still denote by $M_{i,j}$ the initial submatrix $M_{i,j}$ after one or several extractions.

For every (ordered) pair $(M_{i,j}, M_{i',j'})$ with $(i,j) \neq (i',j')$, we perform the following extraction (in any order of these $\binom{k^2}{2}$ pairs). Let s be such that all the $M_{a,b}$'s have size $\text{b}_2^{(s)}(k)$. We find two subsets of size $\text{b}_2^{(s-1)}(k)$, one in $\text{rows}(M_{i,j})$ and one in $\text{cols}(M_{i',j'})$, intersecting at a constant $\text{b}_2^{(s-1)}(k) \times \text{b}_2^{(s-1)}(k)$ submatrix. In $M_{i,j}$ we keep only those rows and the corresponding columns, while in $M_{i',j'}$ we keep only those columns and the corresponding rows. In every other $M_{a,b}$, we keep only the first $\text{b}_2^{(s-1)}(k)$ rows and corresponding columns.

After this extraction performed on the $k^4 - k^2$ zones $\text{rows}(M_{i,j}) \cap \text{cols}(M_{i',j'})$ (with $(i,j) \neq (i',j')$), we obtain the desired rank-k Latin division (on a submatrix of M). ◀

A simple consequence of Lemma 25 is that every class $\mathcal{M}$ with unbounded grid rank satisfies $|\mathcal{M}_n| \geqslant \lfloor \frac{n}{2} \rfloor!$. Indeed there is a simple injection from n-permutations to $2n \times 2n$ submatrices of any rank-2 Latin n-division. This is enough to show that classes of unbounded grid rank are not small. We will need some more work to establish the sharper lower bound of $n!$.

7 Matrix classes with unbounded grid rank

In this section, we prove our main result concerning matrix classes, Theorem 5. The plan is to refine the cleaning of rank Latin divisions, and prove the following.

[6] Or for readers familiar with the game ultimate tic-tac-toe, at positions of moves forcing the next move in the symmetric cell about the diagonal.

▶ **Theorem 27.** *Let $\mathcal{M}$ be a matrix class over a finite alphabet with unbounded grid rank. Then some a-selection $s_a(\mathcal{M})$ of $\mathcal{M}$ includes $\mathcal{F}_{\mathrm{s}}$ for some $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$.*

We refer the reader to Section 3.9 for the definition of $\mathcal{F}_{\mathrm{s}}$ and a-selections. Theorem 27 will later simplify our task when we move to the growth of ordered graphs. Moreover, it easily yields Theorem 5, as we will prove in Theorem 36 that each of the classes $\mathcal{F}_{\mathrm{s}}$ is independent and intractable. In addition, we show that the six classes in Theorem 27 constitute a minimal family, in the sense that none of the classes is contained in another.

7.1 Finding $k!$ distinct $k \times k$ matrices when the grid rank is unbounded

In this section, we prove that matrix classes of unbounded grid rank have at least factorial growth. Apart from that, we prove the following, weaker variant of Theorem 27.

▶ **Theorem 28.** *Let $\mathcal{M}$ be a $0,1$-matrix class with unbounded grid rank. Then there exists $\eta : \{-1,1\} \times \{-1,1\} \to \{0,1\}$ such that $\mathcal{F}_\eta \subseteq \mathcal{M}$.*

The sixteen classes $\mathcal{F}_\eta$ are defined below, and include the classes $\mathcal{F}_{\mathrm{s}}$ for $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$. In the following section, we will reduce those sixteen classes down to six.

Recall that the *order type* $\mathrm{ot}(x,y)$ of a pair (x,y) of elements in a totally ordered set is equal to -1 if $x > y$, 0 if $x = y$, and 1 if $x < y$.

▶ **Definition 29.** *Let $k \geqslant 1$ be an integer and $\eta : \{-1,1\} \times \{-1,1\} \to \{0,1\}$. For every $\sigma \in \mathfrak{S}_k$ we define the $k \times k$ matrix $F_\eta(\sigma) = (f_{i,j})_{1 \leqslant i,j \leqslant k}$ by setting for every $i, j \in [k]$,*

$$f_{i,j} := \begin{cases} \eta(\mathrm{ot}(\sigma^{-1}(j), i), \mathrm{ot}(j, \sigma(i))) & \text{if } \sigma(i) \neq j \\ 1 - \eta(1,1) & \text{if } \sigma(i) = j. \end{cases}$$

Finally $\mathcal{F}_\eta$ is the submatrix closure of $\{F_\eta(\sigma) \mid \sigma \in \mathfrak{S}_n, n \geqslant 1\}$.

These matrices generalize reorderings of matrices in $\mathcal{N}_k$. For example, we find exactly the permutation matrices (reorderings of $\mathbf{I}_k$) when η is constant equal to 0 and their complement when η is constant equal to 1. See Figure 11 for more interesting examples of such matrices.

The six classes $\mathcal{F}_{\mathrm{s}}$'s with $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$ correspond to six of the sixteen possible encodings η. More specifically, $\mathcal{F}_{\mathrm{s}} = \mathcal{F}_\eta$ where η is defined as follows, depending on s:
($=$) $\eta(x,y) = 0$ for every $(x,y) \neq (1,1)$;
($\neq$) $\eta(x,y) = 1$ for every $(x,y) \neq (1,1)$;
($\leqslant_R$) $\eta(1,1) = \eta(-1,1) = 1$ and $\eta(-1,-1) = \eta(1,-1) = 0$;
($\geqslant_R$) $\eta(1,1) = \eta(-1,1) = 0$ and $\eta(-1,-1) = \eta(1,-1) = 1$;
($\leqslant_C$) $\eta(-1,-1) = \eta(-1,1) = 1$ and $\eta(1,1) = \eta(1,-1) = 0$;
($\geqslant_C$) $\eta(-1,-1) = \eta(-1,1) = 0$ and $\eta(1,1) = \eta(1,-1) = 1$.
A careful reader might notice that the entries at positions $(i, \sigma(i))$ differ between the encodings of, say, $\mathcal{F}_{\leqslant_R}$ and $\mathcal{F}_\eta$ with $\eta(1,1) = \eta(-1,1) = 1$ and $\eta(-1,-1) = \eta(1,-1) = 0$. The latter class could rather be denoted by $\mathcal{F}_{<_R}$. As the $\mathcal{F}_{\mathrm{s}}$'s are closed under taking submatrices, we already observed that $\mathcal{F}_{\geqslant R} = \mathcal{F}_{>R}$ holds.

With the next lemma, we get even cleaner universal patterns out of a large rank Latin division. We use the notation of Lemma 17.

▶ **Lemma 30.** *Let $k \geqslant 1$ be an integer. Let M be a $0,1$-matrix with a rank-k Latin N-division with $N := \mathrm{g}_2(k)$. Then there exists $\eta : \{-1,1\} \times \{-1,1\} \to \{0,1\}$ such that the submatrix closure of M contains the set $\{F_\eta(\sigma) \mid \sigma \in \mathfrak{S}_k\}$.*

Proof. Let $(\mathcal{R}, \mathcal{C})$ be the rank-k Latin N-division, with $\mathcal{R} := \{R_1, \ldots, R_N\}$ and $\mathcal{C} := \{C_1, \ldots, C_N\}$, so that every row of R_i (resp. column of C_i) is smaller than every row of R_j (resp. column of C_j) whenever $i < j$. Let $M_{i,j}$ be the *chosen* contiguous submatrix of $\mathcal{N}_k$ in $R_i \cap C_j$ for every $i, j \in [N]$. Recall that,

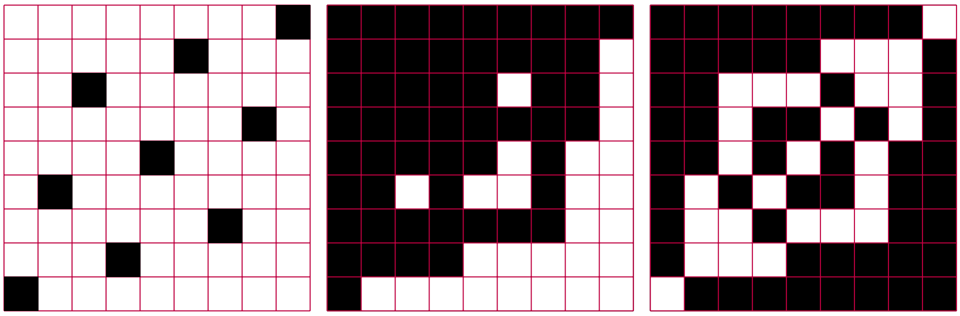

Figure 11 Left: 9×9 permutation matrix M_σ. Center: The matrix $F_{\eta_1}(\sigma)$ with $\eta_1(1,1) = 0$ and $\eta_1(-1,-1) = \eta_1(-1,1) = \eta_1(1,-1) = 1$. Right: The matrix $F_{\eta_2}(\sigma)$ with $\eta_2(1,1) = \eta_2(-1,-1) = 1$ and $\eta_2(-1,1) = \eta_2(1,-1) = 0$.

by definition of a rank Latin division, $\{\mathrm{rows}(M_{i,j})\}_{i,j\in[N]}$ partitions $\mathrm{rows}(M)$ (resp. $\{\mathrm{cols}(M_{i,j})\}_{i,j\in[N]}$ partitions $\mathrm{cols}(M)$) into intervals.

Let $\mathbf{N} := ([N]^2, <_1, <_2)$, where $<_1$ is the lexicographic order on $[N]^2$ which first orders according to the first coordinate and then the second one, while $<_2$ is the lexicographic order on $[N]^2$ which first orders according to the second coordinate and then the second one.

We now define a coloring $c : [N]^4 \to \{0,1\}$ as follows: for every $(i,j) \neq (i',j') \in [N]^2$, we let $c((i,j),(i',j')) \in \{0,1\}$ be the value of the constant entries in $\mathrm{rows}(M_{i,j}) \cap \mathrm{cols}(M_{i',j'})$. We choose arbitrary colors when $(i,j) = (i',j')$. By Lemma 17, there are two sets $R, C \in \binom{[N]}{k}$ such that for every $i,i',j,j' \in [R]$, the value $c((i,j),(i',j'))$ only depends on $\mathrm{ot}_1((i,j),(i',j'))$ and $\mathrm{ot}_2((i,j),(i',j'))$. In particular when $i \neq i'$ and $j \neq j'$, this value only depends on $\mathrm{ot}(i,i')$ and $\mathrm{ot}(j,j')$.

Let $\eta : \{-1,1\} \times \{-1,1\} \to \{0,1\}$ be such that for every $i \neq i' \in [N]$ and $j \neq j' \in [N]$ we have:

$$c((i,j),(i',j')) = \eta(\mathrm{ot}(j,j'), \mathrm{ot}(i,i')).$$

In terms of the rank Latin division, this means that for every $i < i' \in R$ and $j < j' \in C$,

- $\mathrm{cols}(M_{i,j}) \cap \mathrm{rows}(M_{i',j'})$ has constant value $\eta(-1,-1)$,
- $\mathrm{rows}(M_{i,j}) \cap \mathrm{cols}(M_{i',j'})$ has constant value $\eta(1,1)$,
- $\mathrm{cols}(M_{i',j}) \cap \mathrm{rows}(M_{i,j'})$ has constant value $\eta(-1,1)$, and
- $\mathrm{rows}(M_{i',j}) \cap \mathrm{cols}(M_{i,j'})$ has constant value $\eta(1,-1)$.

Figure 12 How zones are determined by η, $\mathrm{ot}(i,i')$, and $\mathrm{ot}(j,j')$.

In other words, $\mathrm{rows}(M_{i,j}) \cap \mathrm{cols}(M_{i',j'})$ is entirely determined by η, $\mathrm{ot}(i,i')$, and $\mathrm{ot}(j,j')$ (see Figure 12).

Let $\sigma \in \mathfrak{S}_k$. We now show how to find $F_\eta(\sigma) = (f_{i,j})_{1\leqslant i,j\leqslant k}$ as a submatrix of M. For every $i \in [k]$, we choose a row $r_i \in \mathrm{rows}(M_{i,\sigma(i)})$ and a column $c_{\sigma(i)} \in \mathrm{cols}(M_{i,\sigma(i)})$ such that the entry of M at the intersection of r_i and $c_{\sigma(i)}$ has value $f_{i,\sigma(i)}$. This is possible since the submatrices $M_{i,j}$ are in $\mathcal{N}_k$ and have disjoint row and column supports. We consider the $k \times k$ submatrix M' of M with rows $\{r_i \mid i \in [k]\}$ and columns $\{c_i \mid i \in [k]\}$.

By design $M' = F_\eta(\sigma)$ holds. Let us write $M' := (m_{i,j})_{1\leqslant i,j\leqslant k}$ and show for example that if $\mathrm{ot}(\sigma^{-1}(j), i) = -1$ and $\mathrm{ot}(j, \sigma(i)) = 1$ for some $i,j \in [k]$, then we have $m_{i,j} = \eta(-1,1) = f_{i,j}$. The other cases are obtained in a similar way. Let $i' := \sigma^{-1}(j) > i$ and $j' := \sigma(i) > j$. In M', $m_{i,j}$ is obtained by taking the entry of M associated to the row r_i of the matrix $M_{i,\sigma(i)} = M_{i,j'}$ and the column c_j of

$M_{\sigma^{-1}(j),j} = M_{i',j}$. The entry $m_{i,j}$ lied in M in the zone $\text{rows}(M_{i,j'}) \cap \text{cols}(M_{i',j})$ with constant value $\eta(-1,1)$. ◀

We now check that $\sigma \in \mathfrak{S}_k \mapsto F_\eta(\sigma)$ is indeed injective.

▶ **Lemma 31.** *For every $k \geqslant 1$ and $\eta : \{-1,1\} \times \{-1,1\} \to \{0,1\}$:*

$$|\{F_\eta(\sigma) \mid \sigma \in \mathfrak{S}_k\}| = k!$$

Proof. We let $k \geqslant 1$ and $\eta : \{-1,1\} \times \{-1,1\} \to \{0,1\}$. The inequality $|\{F_\eta(\sigma) \mid \sigma \in \mathfrak{S}_k\}| \leqslant k!$ simply holds. We thus focus on the converse inequality.

When we read out the first row (bottom one) of $F_\eta(\sigma) = (f_{i,j})_{1 \leqslant i,j \leqslant k}$ by increasing column indices (left to right), we get a possibly empty list of values $\eta(-1,1)$, one occurrence of $1 - \eta(1,1)$ at position $(1, \sigma(1))$, and a possibly empty list of values $\eta(1,1)$. The last index j such that $f_{1,j} \neq f_{1,j+1}$, or $j = k$ if no such index exists, thus corresponds to $\sigma(1)$. We remove the first row and the j-th column and iterate the process on the rest of the matrix. ◀

By piecing Lemmas 25, 30, and 31 together, we get:

▶ **Theorem 32.** *Every $0,1$-matrix class $\mathcal{M}$ with unbounded grid rank satisfies $|\mathcal{M}_k| \geqslant k!$, for every integer k.*

Proof. We fix

$$k \geqslant 1,\ n := \mathrm{R}_{16}(k),\ N := \mathrm{R}_{16}^{(\binom{n}{2}+1)}(k).$$

Now we let $K := K(N)$ be the integer of Lemma 25 sufficient to get a rank-N Latin division. As $\mathcal{M}$ has unbounded grid rank, it contains a matrix M with grid rank at least K. By Lemma 25, a submatrix $\tilde{M} \in \mathcal{M}$ of M admits a rank-N Latin division, from which we can extract a rank-k Latin N-division (since $k \leqslant N$). By Lemma 30 applied to $\tilde{M}$, there exists η such that $\{F_\eta(\sigma) \mid \sigma \in \mathfrak{S}_k\} \subseteq \mathcal{M}_k$. By Lemma 31, this implies that $|\mathcal{M}_k| \geqslant k!$. ◀

Proof of Theorem 28. We just showed that for every matrix class $\mathcal{M}$ of unbounded grid rank, for every integer k, there is an $\eta(k) : \{-1,1\} \times \{-1,1\} \to \{0,1\}$ such that $\left\{F_{\eta(k)}(\sigma) \mid \sigma \in \mathfrak{S}_k\right\} \subseteq \mathcal{M}_k \subseteq \mathcal{M}$. As there are only 16 possible *functions* η, the sequence $\eta(1), \eta(2), \ldots$ contains at least one function η infinitely often. Besides for every $k' < k$, $\{F_\eta(\sigma) \mid \sigma \in \mathfrak{S}_{k'}\}$ is included in the submatrix closure of $\{F_\eta(\sigma) \mid \sigma \in \mathfrak{S}_k\}$. This proves $\mathcal{M} \supseteq \mathcal{F}_\eta$. ◀

7.2 Minimal family of six unavoidable classes

Theorem 28 shows that each matrix class with unbounded twin-width contains one of the sixteen classes $\mathcal{F}_\eta$. We will now see that some of these classes are contained in some others. We say that a mapping $\eta\colon \{-1,1\} \times \{-1,1\} \to \{0,1\}$ *depends only on one coordinate* if either $\eta(x,y) = \eta(x',y)$ for all $x, x', y \in \{-1,1\}$, or $\eta(x,y) = \eta(x,y')$ for all $x, y, y' \in \{-1,1\}$. Note that there are only six such functions η. Indeed, once we fix the coordinate (first or second) η depends on, there are four mappings from $\{-1,1\}$ to $\{0,1\}$. This adds up to eight mappings but the constant-0 and constant-1 mappings are each counted twice, hence a total of six functions.

These six mappings η correspond to the six classes $\mathcal{F}_=$, $\mathcal{F}_\neq$, $\mathcal{F}_{\leqslant R}$, $\mathcal{F}_{\geqslant R}$, $\mathcal{F}_{\leqslant C}$, and $\mathcal{F}_{\geqslant C}$ defined in Section 3.9, where we recall that the matrix encoding a given permutation σ respectively follows the Iverson brackets

$$[\sigma(i) = j], \quad [\sigma(i) \neq j], \quad [i \leqslant \sigma^{-1}(j)], \quad [i \geqslant \sigma^{-1}(j)], \quad [j \leqslant \sigma(i)], \text{ and } \quad [j \geqslant \sigma(i)]$$

for its entry at position (i,j). Thus to establish the refined milestone of the section, Theorem 27, we shall prove the following.

▶ **Lemma 33.** *Let $\eta\colon \{-1,1\} \times \{-1,1\} \to \{0,1\}$. Then there is a $\gamma\colon \{-1,1\} \times \{-1,1\} \to \{0,1\}$ depending only on one coordinate, such that $\mathcal{F}_\eta \supseteq \mathcal{F}_\gamma$.*

Proof. If η depends only on one coordinate, we are done. In particular, we can suppose that η is not constant. Then there are $x, y, x', y' \in \{-1, 1\}$ with $x = x'$ or $y = y'$ such that $\eta(x, y) \neq \eta(x', y')$.

We will assume that $\eta(1, 1) \neq \eta(-1, 1)$. The three other cases ($\eta(1, 1) \neq \eta(1, -1)$, $\eta(-1, -1) \neq \eta(1, -1)$, and $\eta(-1, -1) \neq \eta(-1, 1)$) are similar and correspond to rotating Figure 13 by a 90, 180, 270-degree angle, respectively. We choose $\gamma \colon \{-1, 1\} \times \{-1, 1\} \to \{0, 1\}$ by setting $\gamma(x, y) = \eta(x, 1)$ for $x, y \in \{-1, 1\}$. By construction γ depends only on the first coordinate. We are left with proving that $\mathcal{F}_\eta \supseteq \mathcal{F}_\gamma$.

Let σ be any k-permutation. We build a $2k$-permutation τ as a perfect shuffle of σ and the identity k-permutation. More precisely, τ has its 1 entries at positions $(i, 2\sigma(i) - 1)$ and $(k + i, 2i)$, for every $i \in [k]$. See Figure 13 for an illustration on a particular 5-permutation.

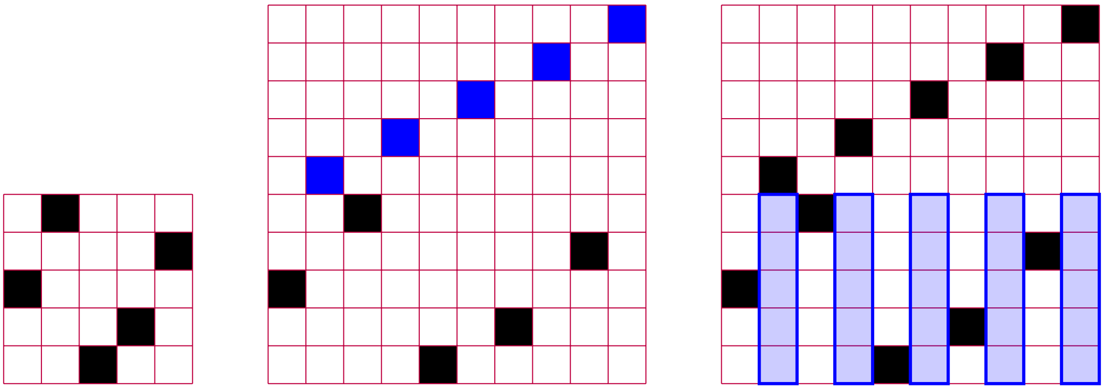

Figure 13 Left: An example of a k-permutation σ with $k = 5$. Center: The matrix of the corresponding $2k$-permutation τ, where the "initial" 1 entries are still in black, and those coming from the identity are in blue. Right: The submatrix where $F_\eta(\tau)$ is only populated by the (distinct) values $\eta(1, 1)$ and $\eta(-1, 1)$.

We claim that $F_\gamma(\sigma)$ appears as the submatrix N of $F_\eta(\tau)$ obtained by keeping the even-indexed columns and the first k rows (see shaded blue area in the right matrix of Figure 13). Indeed as the 1 entry of τ in each kept column is above the kept rows, N depends only on $\eta(1, 1)$ and $\eta(-1, 1)$ (recall Figure 12). Specifically $N_{i,j} = \eta(1, 1)$ if $j < \sigma(i)$ and $N_{i,j} = \eta(-1, 1)$ otherwise. As $\gamma(x, y) = \eta(x, 1)$, the encoding γ follows Figure 12 where $\eta(1, -1)$ is replaced by $\eta(1, 1)$, and $\eta(-1, -1)$ is replaced by $\eta(-1, 1)$. Thus it also holds that $F_\gamma(\sigma)_{i,j} = \eta(1, 1)$ if $j < \sigma(i)$ and $\eta(-1, 1)$ otherwise. Hence $N = F_\gamma(\sigma)$. This proves that $F_\gamma(\sigma) \in \mathcal{F}_\eta$ since $\mathcal{F}_\eta$ is closed under submatrices. And we conclude that $\mathcal{F}_\gamma \subseteq \mathcal{F}_\eta$. ◀

This proves Theorem 27.

Proof of Theorem 27. Follows from Lemmas 23 and 33 and Theorem 28. ◀

We end this section showing that the set of six encoding functions γ's depending only on one coordinate is minimal, in the sense of the following lemma.

▶ **Lemma 34.** *Let $\gamma, \gamma' : \{-1, 1\} \times \{-1, 1\} \to \{0, 1\}$ be two distinct functions depending only on one coordinate. Then $\mathcal{F}_\gamma$ and $\mathcal{F}_{\gamma'}$ are incomparable for $\subseteq$, i.e., neither $\mathcal{F}_\gamma \subseteq \mathcal{F}_{\gamma'}$ nor $\mathcal{F}_{\gamma'} \subseteq \mathcal{F}_\gamma$ hold.*

Proof. We consider the permutation product of two transpositions $\tau := (135)(24)$ over [5]. Let $\gamma :$ $\{-1, 1\} \times \{-1, 1\} \to \{0, 1\}$ be an encoding function that only depends on one coordinate and $N := F_\gamma(\tau)$.

Up to a symmetric argument, we assume that $\gamma(x, y) = \gamma(x, y')$ for every $x, y, y' \in \{0, 1\}$, that is, γ depends on the first coordinate. Let $\eta : \{-1, 1\} \times \{-1, 1\} \to \{0, 1\}$ only depending on one coordinate be such that $N \in \mathcal{F}_\eta$. We will show that $\gamma = \eta$, which implies the desired result.

As $N \in \mathcal{F}_\eta$, there exists some permutation σ such that $M := F_\eta(\sigma)$ contains N as a submatrix.

Case 1. η depends only on the first coordinate, i.e., $\eta(x, y) = \eta(x, y')$ for every $x, y, y' \in \{0, 1\}$.

In this case, observe that every row of M consists of a sequence of consecutive entries with value $\eta(-1, -1) = \eta(-1, 1)$, then an entry with value $1 - \eta(1, 1)$, and then a sequence of consecutive entries with value $\eta(1, 1) = \eta(1, -1)$. Moreover the first row (the bottommost in Figure 14) of N has exactly the values $\gamma(-1,-1), \gamma(-1,-1), 1-\gamma(1,1), \gamma(1,1), \gamma(1,1)$, thus we must have $\eta(-1,-1) = \gamma(-1,-1)$ and $\eta(1,1) = \gamma(1,1)$. Thus we conclude that $\eta = \gamma$.

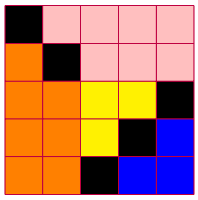

Figure 14 The matrix N in the proof of Lemma 34. The entries with value $\eta(1,1), \eta(-1,-1), \eta(1,-1), \eta(-1,1), 1-\eta(1,1)$ are respectively associated to the colors blue, yellow, pink, orange, and black.

Case 2. η depends only on the second coordinate, i.e., $\eta(x,y) = \eta(x',y)$ for every $x, x', y \in \{0,1\}$.

In that case, observe that every column of M consists of a sequence of consecutive entries with value $\eta(1,1) = \eta(-1,1)$, then an entry with value $1-\eta(1,1)$, and then a sequence of consecutive entries with value $\eta(-1,-1) = \eta(1,-1)$. Moreover the last column of N has exactly the values $\gamma(1,1), \gamma(1,1), 1-\gamma(1,1), \gamma(-1,-1), \gamma(-1,-1)$, thus we must have $\eta(-1,-1) = \gamma(-1,-1)$ and $\eta(1,1) = \gamma(1,1)$. Now observe that the fourth column of N has exactly the values $\gamma(1,1), 1-\gamma(1,1), \gamma(-1,-1), \gamma(1,-1), \gamma(1,-1)$. As $\gamma(1,1) = \eta(1,1)$, we must have $\gamma(1,-1) = \gamma(-1,-1)$. Thus γ is constant, and so is η. In particular we have $\gamma = \eta$ and we are done. ◀

▶ **Corollary 35.** *The classes $\mathcal{F}_{\mathrm{s}}$ for $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$ are precisely all the matrix classes $\mathcal{M}$ with growth $2^{\omega(n)}$ such that every proper subclass of $\mathcal{M}$ has growth $2^{O(n)}$.*

Proof. Each of the classes $\mathcal{F}_{\mathrm{s}}$ has growth $n! = 2^{\omega(n)}$. Suppose $\mathcal{M}$ is a proper subclass of $\mathcal{F}_{\mathrm{s}}$ with growth $2^{\omega(n)}$. Then $\mathcal{M}$ has unbounded twin-width, and hence contains one of the classes $\mathcal{F}_{\mathrm{s}'}$ for some $\mathrm{s}' \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$. But then $\mathcal{F}_{\mathrm{s}'}$ is a proper subclass of $\mathcal{F}_{\mathrm{s}}$, contradicting Lemma 34.

On the other hand, if $\mathcal{M}$ is some class with growth $2^{\omega(n)}$, then $\mathcal{M}$ contains some class $\mathcal{F}_{\mathrm{s}}$, which also has growth $2^{\omega(n)}$. ◀

7.3 Matrix classes of unbounded twin-width are independent

The goal of this section is to prove the following.

▶ **Theorem 36.** *Let $\mathcal{M}$ be a (hereditary) class of matrices over a finite alphabet. If $\mathcal{M}$ has unbounded twin-width then $\mathcal{M}$ efficiently interprets the class of all graphs. In particular, $\mathcal{M}$ is independent and* FO *model checking is* AW[∗]*-hard on $\mathcal{M}$.*

Recall (see Lemma 10) that the class $\mathscr{M}$ of all ordered matchings efficiently interprets the class of all graphs. To prove Theorem 36 it remains to show that each of the classes $\mathcal{F}_{\mathrm{s}}$ efficiently interprets $\mathscr{M}$.

▶ **Lemma 37.** *For each $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$, the class $\mathcal{F}_{\mathrm{s}}$ efficiently interprets the class $\mathscr{M}$ of all ordered matchings.*

Proof. Fix $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$. We construct an interpretation I such that $\mathsf{I}(\mathcal{F}_{\mathrm{s}})$ contains all the ordered matchings of $\mathscr{M}$. Given an ordered matching H with vertices $u_1 < \ldots < u_n < v_1 < \ldots < v_n$, let $\sigma \in \mathfrak{S}_n$ be the permutation such that u_i is adjacent to $v_{\sigma(i)}$, for $1 \leqslant i \leqslant n$. Let M be the 0,1-matrix $F_{\mathrm{s}}(\sigma)$. Clearly, M can be constructed from H in polynomial time. We now describe an interpretation I such that $\mathsf{I}(M)$ is isomorphic to H. The interpretation I will not depend on H, but only on s.

Recall that the matrix M is viewed as a structure with two unary predicates R, C indicating the rows and columns, respectively, a total order $<$ on $R \cup C$, and a binary relation $E \subseteq R \times C$ defining the non-zero entries in M.

If s is '$=$' then the interpretation I simply forgets the unary predicates R and C. If s is '$\neq$' then the interpretation I replaces edges with non-edges. In either case, $\mathsf{I}(M)$ is isomorphic to H.

Suppose s is '$\leqslant_C$', the other cases being symmetric. Recall that in the matrix $M = F_{\mathrm{s}}(\sigma)$ the entries of M are defined by the Iverson bracket $[j \leqslant \sigma(i)]$. Hence, given i, $\sigma(i)$ is the largest value j such that M has entry 1 at position (i, j). This can be defined by a first-order formula:

$$\varphi(x, y) := R(x) \wedge C(y) \wedge E(x, y) \wedge \forall z.(C(z) \wedge y < z) \to \neg E(x, z).$$

The interpretation I, given $M = F_{\mathrm{s}}(\sigma)$ with relations $R, C, <$, and E, outputs the matching with the same domain $R \cup C$, order $<$, and edge relation defined by the formula $\varphi(x, y) \vee \varphi(y, x)$. Then $\mathsf{I}(M)$ is isomorphic to the ordered matching H. ◀

Theorem 36 now follows:

Proof of Theorem 36. By Lemmas 10 and 37 and transitivity, each of the classes $\mathcal{F}_{\mathrm{s}}$ efficiently interprets the class of all graphs. By Theorem 27, every matrix class $\mathcal{M}$ with unbounded twin-width contains one of the classes $\mathcal{F}_{\mathrm{s}}$. Hence $\mathcal{M}$ efficiently interprets the class of all graphs. It follows that $\mathcal{M}$ is independent and FO model checking is AW[$*$]-hard on $\mathcal{M}$ (see Corollary 9). ◀

7.4 Proof of Theorem 5

We can now conclude with the proof of Theorem 5, which we restate below for arbitrary finite alphabets, with all the conditions negated to ease the reasoning.

▶ **Theorem 5.** *Given a class $\mathcal{M}$ of matrices over a finite alphabet A, the following are equivalent.*

- *(¬i) $\mathcal{M}$ has unbounded twin-width.*
- *(¬ii) $\mathcal{M}$ has unbounded grid rank.*
- *(¬iii) $\mathcal{M}$ is not pattern-avoiding.*
- *(¬iv) $\mathcal{M}$ interprets the class of all graphs.*
- *(¬v) $\mathcal{M}$ transduces the class of all graphs.*
- *(¬vi) $\mathcal{M}$ has growth at least $n!$.*
- *(¬vii) $\mathcal{M}$ has growth at least $2^{\omega(n)}$.*
- *(¬viii) FO model checking is not FPT on $\mathcal{M}$. (The implication to (¬viii) holds if FPT≠AW[$*$].)*
- *(¬ix) For all $r \in \mathbb{N}$ there is a matrix $M \in \mathcal{M}$ which admits an r-rich division.*

Proof. The implication (¬*i*)→(¬*ix*) is by Theorem 21, and (¬*ix*)→(¬*ii*) is by Theorem 22. By Lemma 23 if $\mathcal{M}$ has unbounded grid rank then some a-selection $s_a(\mathcal{M})$ has unbounded grid rank. Then (¬*ii*)→(¬*iii*) follows from Theorem 27 applied to $s_a(\mathcal{M})$. The implication (¬*iii*)→(¬*vi*) is clear, as each of the classes $\mathcal{F}_{\mathrm{s}}$ has factorial growth. The implications (¬*iii*)→(¬*iv*), and (¬*iii*)→(¬*viii*) (assuming FPT≠AW[$*$]) are by Lemma 37. The implications (¬*iv*)→(¬*v*) and (¬*vi*)→(¬*vii*) are immediate. The implication (¬*v*)→(¬*i*) is by [8] (¬*vii*)→(¬*i*) is by [7], whereas (¬*viii*)→(¬*i*) follows from [8] and Theorem 2. ◀

8 Classes of ordered graphs with unbounded twin-width

We now move to the world of hereditary classes of ordered graphs. In this language, we will refine the lower bound on the growth of classes of ordered graphs, in order to match the conjecture of Balogh, Bollobás, and Morris [4]. We will also establish that bounded twin-width, NIP, monadically NIP, and tractable (provided that FPT≠AW[$*$]) are all equivalent. This will prove our main result, Theorem 1.

8.1 Twenty-five unavoidable classes of ordered graphs

In the context of graph classes, we rename the parameters $\{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$ to $\{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$, since we interpret the rows/columns as left/right vertices, respectively.

In this section we prove the following two theorems.

▶ **Theorem 38.** *There exist 25 classes of ordered graphs with unbounded twin-width, namely the classes $\mathscr{P}$ and $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ for $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ and $\lambda, \rho \in \{0, 1\}$, such that every hereditary class of ordered graphs with unbounded twin-width includes at least one of these classes.*

▶ **Theorem 39.** *Let $\mathscr{C}$ be a class of ordered graphs with unbounded twin-width. Then $\mathscr{C}$ efficiently interprets the class of all graphs. In particular,* FO *model checking is* AW[$*$]*-hard and $\mathscr{C}$ is independent.*

Let $\sigma \in \mathfrak{S}_n$ be a permutation. For a parameter $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$, an (s, σ)*-matching* is an ordered graph G with vertices $a_1 < \ldots < a_n < b_1 < \ldots < b_n$ such that a_i and b_j are adjacent in G if and only if there is a 1 on position (i, j) in the matrix $F_{\mathrm{s}}(\sigma)$, where now s is treated as an element of $\{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$. In other words, the adjacency between a_i and b_j is defined by the Iverson bracket

$$[\sigma(i) = j], \quad [\sigma(i) \neq j], \quad [i \leqslant \sigma^{-1}(j)], \quad [i \geqslant \sigma^{-1}(j)], \quad [j \leqslant \sigma(i)], \text{ or } \quad [j \geqslant \sigma(i)]$$

depending on the parameter $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$. The vertices $a_1, \ldots, a_n$ are called the *left* vertices, while the vertices $b_1, \ldots, b_n$ are the *right* vertices in the (s, σ)-matching G.

As a first step, we use Theorem 27 to find arbitrary (s, σ)-matchings in graph classes of unbounded twin-width.

▶ **Lemma 40.** *Let $\mathscr{C}$ be a hereditary class of ordered graphs with unbounded twin-width. Then there exists $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ such that for every $n \geqslant 1$ and permutation $\sigma \in \mathfrak{S}_n$, the class $\mathscr{C}$ contains an (s, σ)-matching.*

Proof. Let $\mathcal{M}$ be the submatrix closure of the set of adjacency matrices of graphs in $\mathscr{C}$, along their respective orders. $\mathcal{M}$ has unbounded twin-width (see last paragraph of Section 3.3), and hence unbounded grid rank. By Theorem 27, there exists some $\mathrm{s} \in \{=, \neq, \leqslant_R, \geqslant_R, \leqslant_C, \geqslant_C\}$ such that $\mathcal{F}_{\mathrm{s}} \subseteq \mathcal{M}$.

Let $\sigma \in \mathfrak{S}_n$ be a permutation. We construct a (s, σ)-matching $G \in \mathscr{C}$. Consider its associated *matching permutation* $\widetilde{\sigma} \in \mathfrak{S}_{2n}$ defined by

$$\widetilde{\sigma}(i) := \begin{cases} \sigma(i) + n & \text{if } i \leqslant n \\ \sigma^{-1}(i - n) & \text{if } n + 1 \leqslant i \leqslant 2n. \end{cases}$$

In other words $M_{\widetilde{\sigma}}$ consists of the two blocks M_σ and $M_{\sigma^{-1}}$ on its anti-diagonal. We have $F_{\mathrm{s}}(\widetilde{\sigma}) \in \mathcal{M}$, so there exists a graph $H \in \mathcal{C}$ such that $F_{\mathrm{s}}(\widetilde{\sigma})$ is a submatrix of its adjacency matrix. Denote by U_1, U_2 the (disjoint) ordered sets of vertices corresponding to the rows indexed respectively by $\{1, \ldots, n\}$ and $\{n + 1, \ldots, 2n\}$, such that $\max(U_1) < \min(U_2)$. Take similarly V_1, V_2 associated to the column indices. If $\max(U_1) < \min(V_2)$ we let $A = U_1$ and $B = V_2$; otherwise, $\min(U_2) > \max(U_1) \geqslant \min(V_2) > \max(V_1)$ and we let $A = V_1$ and $B = U_2$. Then, if $a_1 < \cdots < a_n$ are the elements of A and $b_1 < \cdots < b_n$ are the elements of B, we have $a_n < b_1$ and $a_i b_j \in E(H)$ if and only if $F_{\mathrm{s}}(\sigma)$ has a 1 on position (i, j). Hence $G = H[A \cup B]$ is an (s, σ)-matching in $\mathscr{C}$. ◀

As a second step, we make the (s, σ)-matchings more organized, by controlling the edges among the left and right parts. Let $f, g \colon \{-1, 1\} \to \{0, 1\}$ be two functions. An (s, σ)-matching G with vertices $a_1 < \ldots < a_n < b_1 < \ldots < b_n$ is (f, g)*-regular* if the following hold for $1 \leqslant i < j \leqslant n$:

$$[E(a_i, a_j)] = f(\mathrm{ot}(\sigma(i), \sigma(j))) \qquad \textit{and} \qquad [E(b_i, b_j)] = g(\mathrm{ot}(\sigma^{-1}(i), \sigma^{-1}(j))).$$

Let $\mathscr{R}_{\mathrm{s},f,g}$ denote the hereditary closure of the class of all (f, g)-regular (s, σ)-matchings, for all permutations σ.

We now further improve the statement of Lemma 40 to obtain one of the 96 classes $\mathscr{R}_{\mathrm{s},f,g}$. To this end, from a huge (s, π)-matching we will extract a large (f, g)-regular (s, σ)-matching, for some f and g.

▶ **Lemma 41.** *Let $\mathscr{C}$ be a hereditary class of ordered graphs with unbounded twin-width. Then there exist $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ and $f, g \colon \{-1, 1\} \to \{0, 1\}$ such that $\mathscr{C} \supseteq \mathscr{R}_{\mathrm{s},f,g}$.*

Proof. Let $\mathrm{s} \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ be given by Lemma 40. We first show that for every k and permutation $\sigma \in \mathfrak{S}_k$ there is some regular (s, σ)-matching $H \in \mathscr{C}$. Fix a permutation $\sigma \in \mathfrak{S}_k$. Let N and

$\pi \in \mathfrak{S}_N$ be given by Lemma 18. By Lemma 40 some (s, π)-matching G belongs to $\mathscr{C}$. Denote its vertices $a_1 < \ldots < a_N < b_1 < \ldots < b_N$.

Let $c\colon [N]^2 \to \{0,1\}^2$ be such that

$$c(i,j) = ([E(a_i,a_j)], [E(b_{\pi(i)}, b_{\pi(j)}]) \qquad \text{for } i,j \in [N],$$

where $[E(u,v)] = 1$ if the vertices u,v are adjacent in G and 0 otherwise.

Apply Lemma 18 to c. There is a subset $U \subseteq [N]$ such that the subpermutation of π induced by U is isomorphic to σ and such that $c(i,j)$ depends only on $\mathrm{ot}(i,j) = \mathrm{ot}(a_i,a_j)$ and on $\mathrm{ot}(\pi(i),\pi(j)) = \mathrm{ot}(b_{\pi(i)}, b_{\pi(j)})$ for $u,v \in U$. Hence, the subgraph H of G induced by $\{a_i \mid i \in U\} \cup \{b_{\pi(i)} \mid i \in U\}$ is an (f,g)-regular (s,σ)-matching in $\mathscr{C}$, for some $f,g\colon \{-1,1\} \to \{0,1\}$.

Let $\sigma_1, \sigma_2, \sigma_3, \ldots$ be a sequence of permutations such that σ_n is a subpermutation of σ_{n+1}, for all $n \geqslant 1$, and such that for every k and permutation $\sigma \in \mathfrak{S}_k$ there is some n such that σ is a subpermutation of σ_n. For each $n \geqslant 1$ let $f_n, g_n\colon \{-1,1\} \to \{0,1\}$ be such that there is an (f_n,g_n)-regular (s,σ_n)-matching in $\mathscr{C}$. Since there are finitely many possible pairs (f_n,g_n), by taking a subsequence we may assume that $f_n = f$ and $g_n = g$ for some fixed $f,g\colon \{-1,1\} \to \{0,1\}$. Note that if σ is a subpermutation of σ' then the (f,g)-regular (s,σ)-matching is an induced subgraph of the (f,g)-regular (s,σ')-matching. It follows that $\mathscr{C}$ contains all (f,g)-regular (s,σ)-matchings. ◀

▶ **Lemma 42.** *If one of the functions $f,g\colon \{-1,1\} \to \{0,1\}$ is non-constant then $\mathscr{R}_{\mathrm{s},f,g} \supseteq \mathscr{P}$. Otherwise, $\mathscr{R}_{\mathrm{s},f,g} = \mathscr{M}_{\mathrm{s},\lambda,\rho}$, where $\lambda,\rho \in \{0,1\}$ are such that $\lambda = f(-1) = f(1)$ and $\rho = g(-1) = g(1)$.*

Proof. Consider an (f,g)-regular (s,σ)-matching G and let $G[L]$ and $G[R]$ be its ordered subgraphs induced by the left and right vertices, respectively.

If $f(1) = 0$ and $f(-1) = 1$ then by definition, $G[L]$ is isomorphic to G_σ. Similarly, if $g(1) = 0$ and $g(-1) = 1$ then $G[R]$ is isomorphic to $G_{\sigma^{-1}}$. It follows that if $f(1) < f(-1)$ or $g(1) < g(-1)$ then $\mathscr{R}_{\mathrm{s},f,g}$ contains $\mathscr{P}$.

On the other hand, if $f(1) > f(-1)$ then $G[L]$ is isomorphic to the edge complement of G_σ, and if $g(1) > g(-1)$ then $G[R]$ is isomorphic to the edge complement of $G_{\sigma^{-1}}$. Therefore, if $f(1) > f(-1)$ or $g(1) > g(-1)$ then $\mathscr{R}_{\mathrm{s},f,g}$ contains the class of edge complements of ordered graphs in $\mathscr{P}$, but this class is again $\mathscr{P}$ (see Lemma 14).

If f and g are constantly equal to λ and ρ, respectively, then $\mathscr{R}_{\mathrm{s},f,g} = \mathscr{M}_{\mathrm{s},\lambda,\rho}$ by definition. ◀

To prove Theorem 39, it suffices to show that each of the classes $\mathscr{R}_{\mathrm{s},\lambda,\rho}$ efficiently interprets the class of all graphs.

▶ **Lemma 43.** *For every $\mathrm{s} \in \{=,\neq,\leqslant_l,\geqslant_l,\leqslant_r,\geqslant_r\}$ and $f,g\colon \{-1,1\} \to \{0,1\}$, the class $\mathscr{R}_{\mathrm{s},f,g}$ efficiently interprets the class $\mathscr{M}$ of all ordered matchings, via an interpretation I which maps n-element structures to n-element structures. In particular, $\mathscr{R}_{\mathrm{s},f,g}$ has growth at least $\lfloor \frac{n}{2} \rfloor !$.*

Proof. Suppose s is '$\leqslant_r$', the other cases being similar. Let H be an ordered matching with left vertices L and right vertices R, and let $\sigma\colon L \to R$ be the bijection mapping each $v \in L$ to its neighbor in R.

Let G be the (f,g)-regular (s,σ)-matching, with vertices L and R. As G is regular, it is fully determined by H, s, λ,ρ, and can be constructed from H in polynomial time.

We now show how to define the edges of H from G, using an FO formula not depending on H.

The following formula with parameter z and two variables x,y expresses that $x \leqslant z$ and y is the least neighbor of E larger than z:

$$\mu(x,y;z) \equiv (x \leqslant z) \wedge (y > z) \wedge E(x,y) \wedge \forall y'.(z < y' < y) \to \neg E(x,y').$$

Let $\rho(z)$ be the formula expressing that the set of pairs x,y satisfying $\mu(x,y;z)$ defines the graph of a bijection between $\{x \mid x \leqslant z\}$ and $\{x \mid x > z\}$. There is a unique vertex in G satisfying $\rho(z)$, namely the largest element of L, since $|L| = |R|$. Consider the formula

$$\varphi(x,y) \equiv \exists z.\rho(z) \wedge (\mu(x,y;z) \vee \mu(y,x;z)).$$

Then $\varphi(x,y)$ defines in G the matching H, directed from L to R. Hence, $\varphi(x,y) \vee \varphi(y,x)$ defines the edges of the matching H in G, whereas $x<y$ defines the order of H. Together those formulas form an interpretation I such that $\mathsf{I}(G)=H$.

As the interpretation I from Lemma 10 preserves the domains of the structures, and there are $\lfloor \frac{n}{2} \rfloor !$ ordered matchings with n elements, it follows that the class $\mathscr{R}_{\mathrm{s},\lambda,\rho}$ has growth at least $\lfloor \frac{n}{2} \rfloor !$. ◀

Theorems 38 and 39 now follow.

Proof of Theorem 38. By Lemmas 41–43, every class with unbounded twin-width contains one of the classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ or $\mathscr{P}$. The former have growth at least $\lfloor \frac{n}{2} \rfloor !$ by Lemma 43, while the latter has growth $n!$ by Lemma 14. In particular, they have unbounded twin-width by Theorem 7. ◀

Proof of Theorem 39. Follows by Lemmas 10, 41, and 43. ◀

In the next section, we improve the lower bound on the growth of the classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ further, to match the growth of the hereditary closure of the class of matchings.

8.2 Lowerbounding the growth of $\mathscr{M}_{\mathrm{s},\lambda,\rho}$

There is still a bit of work to get the exact value of $\sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k!$ conjectured in [4] as a lower bound of the growth of hereditary classes of ordered graphs with superexponential growth. We show how to derive this bound for $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ in each case of s, λ, ρ.

▶ **Lemma 44.** *For every* $\mathrm{s} \in \{=,\neq,\leqslant_l,\geqslant_l,\leqslant_r,\geqslant_r\}$ *and* $\lambda,\rho \in \{0,1\}$ *and every positive integer* n *we have*

$$|(\mathscr{M}_{\mathrm{s},\lambda,\rho})_n| \geqslant \sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k!, \tag{1}$$

with equality in the case when s *is '='* *and* $\lambda=\rho=0$.

Proof. We first observe that symmetries allow to reduce the number of classes to consider.

▷ Claim 45. It is sufficient to consider the following 7 classes: $\mathscr{M}_{\mathrm{s},0,\rho}$ with $\mathrm{s} \in \{=,\leqslant_r,\geqslant_r\}$ and $\rho \in \{0,1\}$, and $\mathscr{M}_{=,1,1}$.

Proof. Considering the complements of the graph and/or a reverse linear order we can first restrict our study to the case where $\mathrm{s} \in \{=,\leqslant_r\}$. If $\lambda=0$ these classes are in the claimed list, so assume $\lambda=1$. If s is '=' we consider the class with reversed linear order, which has type $\mathscr{M}_{=,\lambda',1}$ for some $\lambda' \in \{0,1\}$, and is therefore in the claimed list. Finally, if the class is $\mathscr{M}_{\leqslant_r,1,\rho}$ we consider the class of its edge complements, which is equal[7] to $\mathscr{M}_{\geqslant_r,0,\rho'}$ where $\rho'=1-\rho$. ◁

Let $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ be one of the 7 classes above. We now prove that (1) holds for any given $n \geqslant 1$. In the case of $\mathscr{M}_{=,1,1}$ we assume $n>3$ and verify the claim for the values $n=1,2,3$ by separately. Indeed, the first 3 values of the growth function are 1, 2, and 6, so (1) holds.

In the general case, we obtain the claimed number of distinct n-element ordered graphs G in $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ by the following process.

Pick a number k with $0 \leqslant 2k \leqslant n$ and $2k$ elements $a_1<\ldots<a_k<b_1<\ldots<b_k$ in $[n]$. Pick a matching E between the a_i's and b_j's in one of $k!$ possible ways. If $k=0$ let $m=n$, otherwise, let $m=b_1-1$. Denote $L=[1,m]$ and $R=[m+1,n]$. Then E is an 'ordered partial matching' with vertices $L \cup R$ (see Figure 15). Construct the ordered graph G with vertices $[n]=L \cup R$ and edges:

- E if s is '='

[7] To be more explicit, the class of edge complements of graphs in $\mathscr{M}_{\leqslant_r,1,\rho}$ directly corresponds to a class which could be denoted $\mathscr{M}_{>_r,0,\rho'}$, which is however equal to $\mathscr{M}_{\geqslant_r,0,\rho'}$ as the classes are hereditary; see analogous discussion following Definition 29.

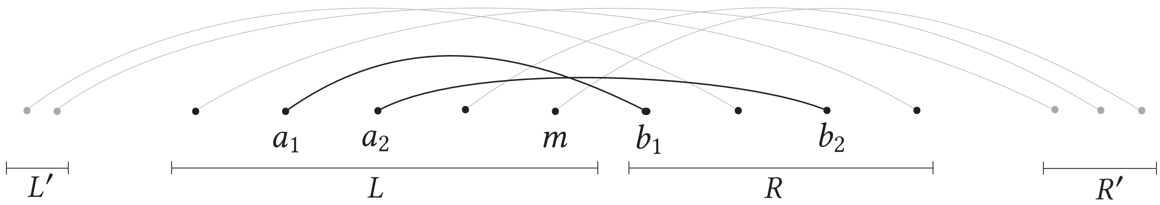

Figure 15 The construction from Lemma 44. The process described in the proof starts with an ordered partial matching, marked with black vertices and edges above. This is then completed to an ordered matching, by matching the isolated vertices with newly added vertices. The vertices R' are inserted after $\max R$ if s is '=' or '$\geqslant_r$', and between $\max L$ and $\min R$ if s is '$\leqslant_r$'.

- $\{uv' \mid uv \in E, u < v \leqslant v', v' \in R\}$ if s is '$\geqslant_r$',
- $\{uv' \mid uv \in E, u < v' \leqslant v, v' \in R\}$ if s is '$\leqslant_r$',

and additionally form in G a clique on L if $\lambda = 1$ and a clique on R if $\rho = 1$.

Clearly, the number of possible outcomes of the above process is at most $\sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k!$. Observe that every ordered graph $G \in \mathscr{M}_{=,0,0}$ can be obtained as a result of the above process, so $|(\mathscr{M}_{=,0,0})_n| \leqslant \sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k!$.

We now argue that different choices made above lead to non-isomorphic outcomes G. Hence, the number of possible outcomes is exactly $\sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k!$. To this end, we show how to recover E from G.

Suppose first that $\lambda = 0$. Then m is the largest vertex in $[n]$ such that $[1, m]$ forms an independent set, since $m + 1 = b_1$ is adjacent to $a_{\sigma^{-1}(1)} < b_1$ (unless $k = 0$, but then $[1, n]$ is still an independent set). Also, for $1 \leqslant i \leqslant k$, the vertex $b_{\sigma(i)}$ is the unique, resp. smallest, resp. largest, neighbor of a_i which is larger than m, depending on $\mathrm{s} \in \{=, \geqslant_r, \leqslant_r\}$. Moreover, all remaining vertices $u \in L \setminus \{a_1, \ldots, a_k\}$ are isolated in G. This allows us to recover $a_1, \ldots, a_k, b_1, \ldots, b_k$ as well as the matching E.

Suppose now that $\lambda, \rho = 1$ and s is '='. Then m is the largest vertex v such that $[1, v]$ forms a clique in G, unless $m = 1, k = 1, a_1 = 1, b_1 = 2$, so that G is the ordered graph G_n obtained from the clique on $[2, n]$ by adding the edge joining 1 with 2. So unless $G = G_n$, we can recover m as described above. If $G = G_n$, then m can be still determined, since $m + 1$ is the smallest vertex v' such that $[v', n]$ forms a clique, unless $m = n - 1, k = 2, a_1 = n - 1, b_1 = n$ so that G is the ordered graph G'_n obtained from the clique on $[1, n - 1]$ by adding the edge joining $(n - 1)$ with n. However, if $G = G_n = G'_n$ implies $n = 3$, but we have assumed that $n > 3$. Hence, we can determine m, given G. Then E is recovered from G by removing all the edges uv with $u < v \leqslant m$ or $m < u < v$.

It remains to verify that $G \in \mathscr{M}_{\mathrm{s},\lambda,\rho}$. Let $A' = L \setminus \{a_1, \ldots, a_k\}$ and $B' = R \setminus \{b_1, \ldots, b_k\}$. Extend the ordered partial matching E to an ordered matching H (see Figure 15) as follows. Create a set L' of $|B'|$ new vertices matched with B' arbitrarily, such that $\max L' < \min L = 1$. Also, create a set R' of $|A'|$ new vertices matched with A' arbitrarily, such that $\min R' > \max R = n$ if s is '=' or '$\geqslant_r$', and $m = \max L < \min R' < \max R' < \min R = m+1$ if s is '$\leqslant_r$'. Then H is an ordered matching with vertices $(L \cup L') \cup (R \cup R')$. Construct the ordered graph G' with the same vertices and order, and with edges:

- $E(H)$ if s is '='
- $\{uv' \mid uv \in E(H), v \leqslant v', v' \in R \cup R'\}$ if s is '$\geqslant_r$',
- $\{uv' \mid uv \in E(H), v' \leqslant v, v' \in R \cup R'\}$ if s is '$\leqslant_r$',

and additionally form in G' a clique on $L \cup L'$ if $\lambda = 1$ and a clique on $R \cup R'$ if $\rho = 1$. Then $G' = H[\mathrm{s}, \lambda, \rho]$ in the terminology of Section 3.9, so $G' \in \mathscr{M}_{\mathrm{s},\lambda,\rho}$. As G is an induced subgraph of G', this shows $G \in \mathscr{M}_{\mathrm{s},\lambda,\rho}$. Altogether, this proves (1). ◀

8.3 Proof of Theorem 1

Finally, we can piece together the proof of our main result, Theorem 1, which we restate below in full generality for arbitrary classes of ordered, binary structures (see Section 3.7 for the definition of an atomic type τ and the interpretation I_τ).

▶ **Theorem 1.** *Let $\mathscr{C}$ be a hereditary class of finite ordered binary structures. Then either $\mathscr{C}$ satisfies conditions (i)-(v), or $\mathscr{C}$ satisfies conditions (i')-(v') below:*

(i) $\mathscr{C}$ has bounded twin-width
(ii) $\mathscr{C}$ has bounded grid rank
(iii) $\mathscr{C}$ has growth $2^{O(n)}$
(iv) $\mathscr{C}$ does not transduce the class of all graphs
(v) FO *model checking is* FPT *on $\mathscr{C}$*

(i') $\mathscr{C}$ has unbounded twin-width
(ii') there is some atomic type $\tau(x, y)$ such that $\mathsf{I}_\tau(\mathscr{C})$ contains $\mathscr{P}$ or one of the 24 classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$
(iii') $\mathscr{C}$ has growth at least $\sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k! \geqslant \lfloor \frac{n}{2} \rfloor !$
(iv') $\mathscr{C}$ interprets the class of all graphs
(v') FO *model checking is* AW[∗]*-hard on $\mathscr{C}$.*

Proof. If $\mathscr{C}$ has bounded twin-width then (*iii*) holds by [7], (*iv*) holds by [8], and (*v*) holds by [8] and Theorem 2, and finally, (*ii*) holds by Theorem 5 (*ii*),(*i*).

Conversely, if $\mathscr{C}$ has unbounded twin-width then the class $\mathcal{M} = \{M(\mathbf{A}) \mid \mathbf{A} \in \mathscr{C}\}$ of adjacency matrices of $\mathscr{C}$ has unbounded twin-width. Recall that $M(\mathbf{A})$ is the adjacency matrix of $\mathbf{A}$ over the alphabet A_Σ consisting of atomic types of pairs of elements (see definition preceding Lemma 12). By Theorem 5 (*i*) (*viii*), model checking is AW[∗]-hard on $\mathcal{M}$. By Lemma 12, this yields AW[∗]-hardness for $\mathscr{C}$, proving (*v'*). By Theorem 5 (*iii*),(*i*) there is some letter (atomic type) $\tau \in A_\Sigma$ such that the selection $s_\tau(\mathcal{M})$ contains one of the classes $\mathcal{F}_\mathrm{s}$, so in particular, $s_\tau(\mathcal{M})$ has unbounded twin-width.

▷ Claim 46. The class $\mathsf{I}_\tau(\mathscr{C})$ has unbounded twin-width.

Proof. Let I'_τ be the interpretation which, given $\mathbf{A} \in \mathscr{C}$, produces the ordered directed graph with the same domain and order as $\mathbf{A}$, such that there is an edge from u to v (possibly $u = v$) if and only if $\tau(u, v)$ holds in $\mathbf{A}$. Note that for $\mathbf{A} \in \mathscr{C}$, the matrices $s_\tau(M(\mathbf{A}))$ and $M(\mathsf{I}'_\tau(\mathbf{A}))$ are equal, as both matrices have entry equal to 1 if $\tau(a, b)$ holds in $\mathbf{A}$ and 0 otherwise. Hence, $s_\tau(\mathcal{M})$ is equal to the class of adjacency matrices of the ordered graphs in $\mathsf{I}'_\tau(\mathscr{C})$. As $\mathcal{M}$ has unbounded twin-width, also $\mathsf{I}'_\tau(\mathscr{C})$ has unbounded twin-width. As the class $\mathsf{I}_\tau(\mathscr{C})$ interprets the class $\mathsf{I}'_\tau(\mathscr{C})$, it follows that $\mathsf{I}_\tau(\mathscr{C})$ has unbounded twin-width, too. ◀

Since $\mathsf{I}_\tau(\mathscr{C})$ is a class of ordered graphs, we may apply the results of the preceding sections. In particular, $\mathsf{I}_\tau(\mathscr{C})$ contains one of the 25 classes $\mathscr{P}$ and $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ by Theorem 38, proving (*ii'*). By Lemma 44, $\mathsf{I}_\tau(\mathscr{C})$, and so also $\mathscr{C}$, has growth at least $\sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} k!$, proving (*iii'*). By Theorem 39, $\mathsf{I}_\tau(\mathscr{C})$, and so also $\mathscr{C}$, interprets the class of all graphs. This proves (*iv'*), and concludes the theorem. ◀

8.4 Minimality

▶ **Lemma 47.** *None of the classes $\mathscr{P}$ and $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ is included in another.*

Proof. We first prove that the class $\mathscr{P}$ is incomparable with the classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$: The class $\mathscr{P}$ contains the ordered graph 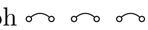 (as witnessed by the permutation (214365)) but none of the classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ does (as all the ordered graphs in these classes admits a vertex partition into two intervals L and R with $\max L < \min R$, each inducing a complete or an edgeless graph). On the other hand, none of 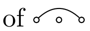 and are ordered permutation graphs, but each class $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ contains one of these. Indeed, as this set of two ordered graphs is closed by complement and order reversal it is sufficient to consider the next 6 classes: $\mathscr{M}_{\mathrm{s},0,\rho}$ with $s \in \{=, \leqslant_r, \geqslant_r\}$ and $\mathscr{M}_{=,1,1}$ (see Claim 45). For $\mathscr{M}_{=,1,1}$ we immediately check that $\in \mathscr{M}_{=,1,1}$. It is easily checked that 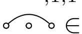 $\in \mathscr{M}_{\mathrm{s},0,0}$ for $s \in \{=, \leqslant_r, \geqslant_r\}$ and that $\in \mathscr{M}_{\mathrm{s},\lambda,\rho}$ for $s \in \{=, \leqslant_r, \geqslant_r\}$ and $(\lambda, \rho) \neq (0, 0)$. Hence $\mathscr{P}$ is incomparable with classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$.

Thus we can restrict our attention to the comparison of classes $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ and $\mathscr{M}_{\mathrm{s}',\lambda',\rho'}$. We consider the ordered matching $M =$  and let $G = M[\mathrm{s}, \lambda, \rho]$ be the graph in $\mathscr{M}_{\mathrm{s},\lambda,\rho}$ originating from M. Assume G is an induced subgraph of a graph $G' \in \mathscr{M}'_{\mathrm{s}',\lambda',\rho'}$, originating from an ordered matching M', so that $G' = M'[\mathrm{s}', \lambda', \rho']$, and that M' has a minimal number of vertices. Note that there is exactly one vertex-partition of G' into intervals L, R with $\max L < \min R$ such that both L and R are either cliques or independent sets. This readily implies $\lambda' = \lambda$ and $\rho' = \rho$. By deleting the internal edges of the parts L and R we reduce to the case where $\lambda = \rho = 0$. Every ordered graph in $\mathscr{M}_{\mathrm{s}',0,0}$ with parts L and R with

$|L| = |R| = 4$ has the following properties: if s' is '$=$' then the maximum degree is at most 1, if s' is '$\neq$' then the minimum degree is at least 3, if s' is '$\geqslant_r$' then the last vertex has maximum degree in R, if s' is '$\leqslant_r$' then $\min R$ has maximum degree in R, if s' is '$\leqslant_l$' then the first vertex has maximum degree in L, and if s' is '$\leqslant_r$' then $\min L$ has maximum degree in L. When s is $=, \neq, \leqslant_r, \geqslant_r, \leqslant_l, \geqslant_l$ the degrees of the vertices of G are $(1,1,1,1,1,1,1,1)$, $(3,3,3,3,3,3,3,3)$, $(3,4,1,2,4,3,2,1)$, $(2,1,4,3,1,2,3,4)$, $(4,3,2,1,3,4,1,2)$, and $(1,2,3,4,2,1,4,3)$. Thus it is easily checked that the only possible choice is $s' = s$. ◀

By the same argument as in Corollary 35, we get:

▶ **Corollary 48.** *The 25 classes $\mathscr{P}$ and $\mathscr{M}_{s,\lambda,\rho}$ for $s \in \{=, \neq, \leqslant_l, \geqslant_l, \leqslant_r, \geqslant_r\}$ and $\lambda, \rho \in \{0,1\}$ are precisely all the hereditary classes of ordered graphs $\mathscr{C}$ with growth $2^{\omega(n)}$ such that every proper subclass of $\mathscr{C}$ has growth $2^{O(n)}$.*

9 Model-theoretic characterizations

In this section, we present further model-theoretic characterizations of classes of bounded twin-width, as well as prove more general results concerning arbitrary classes of structures, over an arbitrary signature. In particular, we generalize the implications $(ii) \to (iii) \to (v)$ from Theorem 5 to arbitrary classes of structures, by proving Theorem 3. Namely, we show that every monadically dependent class of structures excludes certain grid-like patterns, and every class of structures which excludes such grid-like patterns satisfies a property generalizing bounded grid rank.

We start with defining the notion of a restrained class, generalizing the notion of bounded grid rank for matrices. First, we introduce a notion generalizing the concept of the number of distinct rows in a zone of a matrix, in arbitrary structures.

In this section, whenever $\mathbf{S}$ is a structure then we identify $\mathbf{S}$ with its domain, when writing e.g. $a \in \mathbf{S}$ or $A \subseteq \mathbf{S}$. We also write $\mathbf{S}^{\bar{x}}$ for the set of all *valuations* $\bar{a}$ of a set of variables $\bar{x}$ in $\mathbf{S}$, where a valuation is a function $\bar{a} \colon \bar{x} \to \mathbf{S}$.

Let $\Delta(\bar{u};\bar{v})$ be a finite set of formulas $\theta(\bar{u};\bar{v})$ with free variables contained in $\bar{u}$ and $\bar{v}$. For a structure $\mathbf{S}$, tuple $\bar{a} \in \mathbf{S}^{\bar{u}}$ and a set $B \subseteq \mathbf{S}$ define the *Δ-type of $\bar{a}$ over B* as:

$$\mathrm{tp}^{\Delta}(\bar{a}/B) = \{(\theta, \bar{b}) \in \Delta \times B^{\bar{v}} \mid \mathbf{S} \models \theta(\bar{a};\bar{b})\}.$$

For a set $A \subseteq \mathbf{S}$, denote

$$\mathrm{Types}^{\Delta}(A/B) := \{\mathrm{tp}^{\Delta}(\bar{a}/B) \mid \bar{a} \in A^{\bar{u}}\}.$$

▶ **Example 49.** Let M be a 0-1 matrix, viewed as an (ordered) binary structure with the unary predicate $R \subseteq M$ indicating the rows and the binary relation E defining the entries of the matrix. Let $\Delta(u,v) = \{E(u,v)\}$. Let $B \subseteq M \setminus R$ be a set of columns of M. Then, for a row $a \in R$, the set $\mathrm{tp}^{\Delta}(a/B)$ corresponds to the set of those columns $b \in B$ with a non-zero entry in row a. For a set of rows $A \subseteq R$, $|\mathrm{Types}^{\Delta}(A/B)|$ is the number of distinct rows in the submatrix of M with rows A and columns B.

Let $\varphi(x;\bar{y})$ be a formula and $\mathbf{S}$ a structure. A *φ-definable disjoint family* is a family $\mathcal{R}$ of pairwise disjoint of subsets of $\mathbf{S}$, where for each $R \in \mathcal{R}$ there is $\bar{b} \in \mathbf{S}^{\bar{y}}$ with $R = \{a \in \mathbf{S} \mid \mathbf{S} \models \varphi(a;\bar{b})\}$. For example, if $\mathbf{S}$ is a finite ordered structure and $\mathcal{R}$ is a partition of $\mathbf{S}$ into convex sets, then $\mathcal{R}$ is a φ-definable family of pairwise disjoint sets, for $\varphi(x;y_1,y_2) = y_1 \leqslant x \leqslant y_2$.

▶ **Definition 50** (Restrained class)**.** *A class $\mathscr{C}$ of structures is* restrained *if the following condition holds. Let $\varphi(x;\bar{y})$, $\psi(x;\bar{z})$ and be formulas over the signature of $\mathscr{C}$, and let $\Delta(\bar{u};\bar{v})$ be a finite set of formulas. Then there are natural numbers t and k such that for any $\mathbf{S} \in \mathscr{C}$ and any φ-definable disjoint family $\mathcal{R}$ and ψ-definable disjoint family $\mathcal{L}$ with $|\mathcal{R}| = |\mathcal{L}| \geqslant t$ there are $R \in \mathcal{R}$ and $L \in \mathcal{L}$ with $|\mathrm{Types}^{\Delta}(R/L)| < k$.*

The following proposition is an analogue of the statement that matrices of bounded grid rank have bounded twin-width.

▶ **Proposition 51.** *Let $\mathscr{C}$ be a class of finite, ordered binary structures. If $\mathscr{C}$ is restrained then $\mathscr{C}$ has bounded twin-width.*

Proof. Consider the formula $\varphi(x; y_1, y_2) \equiv y_1 \leqslant x \leqslant y_2$.. Then every partition of a structure $\mathbf{S} \in \mathscr{C}$ into convex sets is a φ-definable disjoint family.

Let $\Delta(u, v)$ be the set of atomic formulas over the signature of $\mathscr{C}$ with free variables u and v. Let k, t be the numbers as in the definition of a restrained class, applied to the formulas $y_1 \leqslant x \leqslant y_2$ and $z_1 \leqslant x \leqslant z_2$. Thus, for every two convex partitions $\mathcal{R}$ and $\mathcal{L}$ of $\mathbf{S} \in \mathscr{C}$ into at least t parts each there are $R \in \mathcal{R}$ and $L \in \mathcal{L}$ with

$$|\mathrm{Types}^{\Delta}(R/L)| < k \tag{2}$$

Let M be the adjacency matrix of $\mathbf{S}$, that is, the $\mathbf{S} \times \mathbf{S}$ matrix whose entry at position (a, b) is the atomic type of (a, b) in $\mathbf{S}$ (up to a certain encoding as described in Section 3.3, which is however irrelevant here). We show that M has grid rank less than $p := 2^{k2^{|\Delta|}}$.

Suppose that $\mathbf{S}$ has grid rank at least p. Then there are divisions $\mathcal{L}$ of the rows of M and $\mathcal{R}$ of the columns of M, into $p > k$ parts each, such that each zone has at least p different rows or at least p different columns. As the size of the alphabet of M is at most $2^{|\Delta|}$, this implies that each zone has at least k different rows. Hence, for all $L \in \mathcal{L}$ and $R \in \mathcal{R}$ we have that $|\mathrm{Types}^{\Delta}(L/R) \geqslant k|$, which contradicts (2). ◀

We now define a notion which generalizes the notion of avoiding certain patterns. In this case, rather than defining patterns which encode all permutations, it is more convenient to define patterns which encode grids, in the following way.

Fix any signature Σ and a first-order formula $\varphi(\bar{x}, \bar{y}, z)$, where $\bar{x}$ and $\bar{y}$ are sets of variables and z is a single variable. An *$m \times n$ grid* defined by φ in a structure $\mathbf{S}$ is a triple of sets $A \subseteq \mathbf{S}^{\bar{x}}$, $B \subseteq \mathbf{S}^{\bar{y}}$ and $C \subseteq \mathbf{S}$ with $|A| = m$, $|B| = n$ and $|C| = m \times n$, such that the relation

$$\{(\bar{a}, \bar{b}, c) \in A \times B \times C \mid \mathbf{S} \models \varphi(\bar{a}, \bar{b}, c)\}$$

is the graph of a bijection from $A \times B$ to C. More explicitly, for each $c \in C$ there is a unique pair $(\bar{a}, \bar{b}) \in A \times B$ such that $\mathbf{S} \models \varphi(\bar{a}, \bar{b}, c)$, and conversely, for each $(\bar{a}, \bar{b}) \in A \times B$ there is a unique $c \in C$ such that $\mathbf{S} \models \varphi(\bar{a}, \bar{b}, c)$.

▶ **Definition 52** (Defining large grids). *A class of structures $\mathscr{C}$* defines large grids *if there is a formula $\varphi(\bar{x}, \bar{y}, z)$ such that for all $n \in \mathbb{N}$, φ defines an $n \times n$ grid in some structure $\mathbf{S} \in \mathscr{C}$.*

Intuitively, if $\mathscr{C}$ defines large grids then the product of two sets $A \times B$ can be represented by a set of single elements C in some structure $\mathbf{S} \in \mathscr{C}$. Hence an arbitrary relation $R \subseteq A \times B$ can be represented by some subset of C, so $\mathscr{C}$ is monadically independent. This is expressed by the following.

▶ **Lemma 53.** *If $\mathscr{C}$ defines large grids then $\mathscr{C}$ is monadically independent.*

Proof. Suppose $\mathscr{C}$ defines large grids and let $R \subseteq A \times B$ be a binary relation between two finite sets A, B. Then there is a structure $\mathbf{S} \in \mathscr{C}$ and sets A', B', C such that $\varphi(\bar{x}, \bar{y}, z)$ defines a grid (A', B', C) in $\mathbf{S}$, and $|A| = |A'|$ and $|B| = |B'|$. Fix an arbitrary bijections $f \colon A \to A'$ and $g \colon B \to B'$. Let $U \subseteq C$ be such that for all $c \in C$,

$$c \in U \qquad \Leftrightarrow \qquad \mathbf{S} \models \varphi(f(a), g(b), c) \text{ for some } a \in A \text{ and } b \in B \text{ with } R(a, b).$$

As φ defines a bijection between $A \times B$ and C, it follows that for all $a \in A$ and $b \in B$,

$$R(a, b) \quad \Leftrightarrow \quad \mathbf{S} \models \exists z. \varphi(f(a), g(b), c) \wedge U(z).$$

Consider the monadic lift $\mathscr{C}^+$ of $\mathscr{C}$ which consists of all structures $\mathbf{S} \in \mathscr{C}$ expanded with a unary predicate U which is interpreted as a finite set. The above shows that the formula $\exists z. \varphi(\bar{x}, \bar{y}, z) \wedge U(z)$ is independent over $\mathscr{C}^+$. Hence, $\mathscr{C}$ is not monadically dependent. ◀

The following result generalizes the implications (*ii*)→ (*iii*)→ (*v*) from Theorem 5 to arbitrary classes of structures – finite or infinite, ordered or unordered, and over an arbitrary signature. It also involves a notion which we call *1-dimensionality*, which is a model-theoretic notion originating from Shelah (it is called *finite satisfiability dichotomy* in [11]). This is a central tool in the study of monadically dependent classes. It is defined in the following section, in terms of a variant of forking independence – a key concept in stability theory, generalizing e.g. independence in vector spaces or algebraic independence.

▶ **Theorem 54.** *Let $\mathscr{C}$ be a class of structures over a fixed signature.*
Then the implications (1)↔(2)↔(3)↔(4)→(5) hold among the following conditions.

(1) $\mathscr{C}$ does not transduce the class of all graphs,
(2) $\mathscr{C}$ is monadically dependent,
(3) $\mathscr{C}$ does not define large grids,
(4) $\mathscr{C}$ is 1*-dimensional (cf. Def. 60),*
(5) $\mathscr{C}$ is restrained.

The equivalence (*1*) ↔ (*2*) is Theorem 13, due to Baldwin and Shelah. The implication (*2*) → (*3*) is Lemma 53. The implication (*3*) → (*4*) is due to Shelah (cf. Prop. 63). The implication (*4*)→(*2*) has been recently proved by Braunfeld and Laskowski [11]. Our contribution is the implication (*4*)→(*5*).

We believe this result may be of independent interest, and possibly of broader applicability than just in the context of ordered structures. For example, by Theorem 54, all graph classes of bounded twin-width (without an order) and all interpretations of nowhere-dense classes are restrained. Conversely, every class of structures which is not restrained defines large grids.

Theorem 54 allows us to provide further, model theoretic characterizations of hereditary classes of finite, ordered, binary structures of bounded twin-width:

▶ **Theorem 55.** *Let $\mathscr{C}$ be a hereditary class of finite, ordered, binary structures. Then the following conditions are equivalent:*

(1) $\mathscr{C}$ does not transduce the class of all finite graphs,
(2) $\mathscr{C}$ is monadically dependent,
(3) $\mathscr{C}$ does not define large grids,
(4) $\mathscr{C}$ is 1*-dimensional,*
(5) $\mathscr{C}$ is a restrained class,
(6) $\mathscr{C}$ has bounded twin-width,
(7) $\mathscr{C}$ is dependent.

Proof. The implications (*1*) → (*2*) → (*3*) → (*4*) → (*5*) follow from Theorem 54. The implication (*5*) →(*6*) is proved in Proposition 51. The implication (*6*)→(*1*) is by [8] (cf. Theorem 11). This proves the equivalence of the first six items.

The implication (*2*)→(*7*) is immediate, whereas the implication (*7*)→(*6*) is by Theorem 1,(*iv'*),(*i*). ◀

To prove Theorem 54, it remains to prove that every 1-dimensional class is restrained. First, we need to define 1-dimensionality.

9.1 1-dimensionality

We now introduce a wholly model-theoretic notion which can be used to characterize bounded twin-width, but also arbitrary monadically dependent classes of structures. For this, we first recall some basic notions from model theory.

By a *model* we mean a structure which is typically infinite, as opposed to the structures considered earlier, which were typically finite. We give a brief account of basic notions from model theory in Appendix C, although they are not needed to follow the main text below.

The *elementary closure* of a class of structures $\mathscr{C}$ is the class of all models $\mathbf{M}$ that satisfy every sentence φ that holds in all structures $\mathbf{S} \in \mathscr{C}$. In particular, if $\mathscr{C}$ does not define large grids, then neither does its elementary closure. This is because for any fixed $n \in \mathbb{N}$ the existence of an $n \times n$-grid defined by a fixed formula $\varphi(\bar{x}, \bar{y}, z)$ can be expressed by a first-order sentence φ' which existentially quantifies $(|\bar{x}| + |\bar{y}|) \cdot n + n^2$ variables, corresponding to sets A, B, C of $\bar{x}$-tuples, $\bar{y}$-tuples and single vertices, and then checks that φ defines a bijection between $A \times B$ and C.

By the compactness theorem (cf. Thm. 68), if $\mathscr{C}$ defines large grids, then its elementary closure contains a structure that defines a grid (A, B, C) with A and B of arbitrarily large infinite cardinalities.

▶ **Definition 56** (Elementary extension). *Let* $\mathbf{M}, \mathbf{N}$ *be two models. Then* $\mathbf{N}$ *is an* elementary extension *of* $\mathbf{M}$*, written* $\mathbf{M} \prec \mathbf{N}$*, if the domain of* $\mathbf{M}$ *is contained in the domain of* $\mathbf{N}$*, and for every formula* $\varphi(\bar{x})$ *and tuple* $\bar{a} \in \mathbf{M}^{\bar{x}}$ *of elements of* $\mathbf{M}$*,*

$$\mathbf{M} \models \varphi(\bar{a}) \quad \textit{if and only if} \quad \mathbf{N} \models \varphi(\bar{a}).$$

In other words, it doesn't matter if we evaluate formulas in $\mathbf{M}$ or in $\mathbf{N}$. In particular, $\mathbf{M}$ and $\mathbf{N}$ satisfy the same sentences.

A *formula* $\varphi(\bar{x})$ *with parameters* from $C \subseteq \mathbf{N}$ is a formula using constant symbols denoting elements from C. Such a formula can be evaluated in $\mathbf{N}$ on a tuple $\bar{a} \in \mathbf{N}^{\bar{x}}$, as expected. Note that if $\mathbf{M} \prec \mathbf{N}$ and $\varphi(\bar{x})$ is a formula with parameters from $\mathbf{N}$ and $\bar{a} \in \mathbf{M}^{\bar{x}}$ then it is not necessarily the case that $\mathbf{M} \models \varphi(\bar{a})$ if and only if $\mathbf{N} \models \varphi(\bar{a})$, although this does hold for formulas with parameters from $\mathbf{M}$.

▶ **Definition 57** (Independence). *Let* $\mathbf{M}$ *be a model and* $\mathbf{N}$ *its elementary extension. For a tuple* $\bar{a} \in \mathbf{N}^{\bar{x}}$ *and a set* $B \subseteq \mathbf{N}$ *say that* $\bar{a}$ *is* independent *from* B *over* $\mathbf{M}$*, denoted* $\bar{a} \mathop{|}\nolimits_{\mathbf{M}} B$*, if for every formula* $\varphi(\bar{x})$ *with parameters from* $B \cup \mathbf{M}$ *such that* $\mathbf{N} \models \varphi(\bar{a})$ *there is some* $\bar{c} \in \mathbf{M}^{\bar{x}}$ *such that* $\mathbf{N} \models \varphi(\bar{c})$*.*

Abusing notation, if B is enumerated by a tuple $\bar{b}$, then we may write $\bar{a} \mathop{|}\nolimits_{\mathbf{M}} \bar{b}$. We write $\not\mathop{|}\nolimits_{\mathbf{M}}$ for the negation of the relation $\mathop{|}\nolimits_{\mathbf{M}}$.

▶ **Example 58.** Let $\mathbf{N}$ be $(\mathbb{R}, \leqslant)$ and let $\mathbf{M}$ be the union of the open intervals $]0, 1[$ and $]8, 9[$, equipped with the relation $\leqslant$. Then $\mathbf{M} \prec \mathbf{N}$. This is easy to derive from the fact that $(\mathbb{R}, \leqslant)$ has quantifier elimination, that is, every formula $\varphi(\bar{x})$ is equivalent to a quantifier-free formula. Figure 16 illustrates

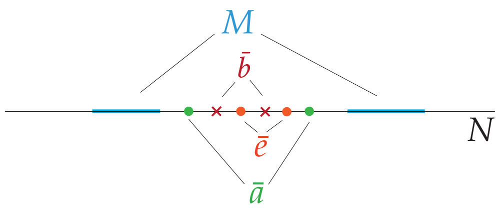

Figure 16 The structures $\mathbf{M} \prec \mathbf{N}$, a set B and two tuples, $\bar{a}, \bar{e}$, with $\bar{a} \mathop{|}\nolimits_{\mathbf{M}} B$ and $\bar{e} \not\mathop{|}\nolimits_{\mathbf{M}} B$.

independence over $\mathbf{M}$.

▶ **Example 59.** Let $\mathbf{M} \prec \mathbf{N}$. Then $\bar{a} \mathop{|}\nolimits_{\mathbf{M}} \mathbf{M}$ for every $\bar{a} \in \mathbf{N}^{\bar{x}}$ (cf. Lemma 70).

▶ **Definition 60** (1-dimensionality). *A model* $\mathbf{M}$ *is* 1-dimensional *if for every* $\mathbf{M} \prec \mathbf{N}$*, tuples* $\bar{a}, \bar{b}$ *of elements of* $\mathbf{N}$ *and* $c \in \mathbf{N}$ *a single element, if* $\bar{a} \mathop{|}\nolimits_{\mathbf{M}} \bar{b}$ *then* $\bar{a}c \mathop{|}\nolimits_{\mathbf{M}} \bar{b}$ *or* $\bar{a} \mathop{|}\nolimits_{\mathbf{M}} \bar{b}c$*. A class* $\mathscr{C}$ *of structures is* 1-dimensional *if every model in the elementary closure of* $\mathscr{C}$ *is 1-dimensional.*

▶ **Example 61.** Any total order $(X, \leqslant)$ is 1-dimensional. As an illustration, in the situation in Fig. 16, consider the tuples $\bar{a}, \bar{b}$ marked therein. Then $\bar{a} \mathop{|}\nolimits_{\mathbf{M}} \bar{b}$. Let $c \in \mathbf{N}$. If c belongs to the interval $]b_1, b_2[$ then $\bar{a} \mathop{|}\nolimits_{\mathbf{M}} \bar{b}c$. Otherwise, $\bar{a}c \mathop{|}\nolimits_{\mathbf{M}} \bar{b}$.

▶ **Example 62.** Let $\mathbf{N} = (\mathbb{R} \times \mathbb{R}, \sim_1, \sim_2)$ where for $i = 1, 2$, the relation $\sim_i$ denotes equality of the ith coordinates. Let $\mathbf{M}$ be the induced substructure of $\mathbf{N}$ with domain $I \times I$ for some infinite subset $I \subseteq \mathbf{N}$. Then $\mathbf{M} \prec \mathbf{N}$. In the situation depicted in Fig. 17, $a \mathop{|}\nolimits_{\mathbf{M}} b$ but both $ac \not\mathop{|}\nolimits_{\mathbf{M}} b$ and $a \not\mathop{|}\nolimits_{\mathbf{M}} bc$. So $\mathbf{M}$ is not 1-dimensional.

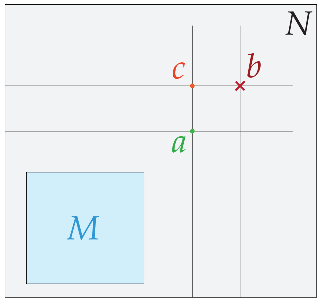

Figure 17 The structures $\mathbf{M} \prec \mathbf{N}$ and elements $a, b, c \in \mathbf{N}$ with $a \mathop{\vert}_{\mathbf{M}} b$, $ac \mathop{\not\vert}_{\mathbf{M}} b$ and $a \mathop{\not\vert}_{\mathbf{M}} bc$.

The following result is essentially [38, Lemma 2.2] (see also [11]).

▶ **Proposition 63.** *If a model* $\mathbf{M}$ *does not define large grids then it is 1-dimensional.*

In particular, every class $\mathscr{C}$ that does not define large grids is 1-dimensional, proving the implication *(3)*→*(4)* in Theorem 54.

9.2 Proof of Theorem 54

In this section we prove that every 1-dimensional class is restrained, proving the remaining implication *(4)*→*(5)* in Theorem 54.

Let $\mathscr{C}$ be a class which is not restrained. We show that $\mathscr{C}$ is not 1-dimensional. We first construct a structure $\mathbf{M}$ in the elementary closure of $\mathscr{C}$ that witnesses that $\mathscr{C}$ is not restrained in a convenient way.

▶ **Lemma 64.** *Suppose $\mathscr{C}$ is not restrained. Then there exist:*

- *formulas* $\varphi(z, \bar{x}), \psi(z, \bar{y}), \theta(\bar{u}, \bar{v})$,
- *a structure* $\mathbf{M}$ *in the elementary closure of* $\mathscr{C}$,
- *an elementary extension* $\mathbf{N}$ *of* $\mathbf{M}$,
- *tuples* $\bar{a}_0, \bar{a}_1 \in \mathbf{N}^{\bar{x}}$ *and* $\bar{b}_0 \in \mathbf{N}^{\bar{y}}$,

such that the following properties hold:

1. $\mathrm{tp}(\bar{a}_0/\mathbf{M}) = \mathrm{tp}(\bar{a}_1/\mathbf{M})$,
2. *the set* $\mathrm{Types}^{\theta}(A/B)$ *is infinite, where* $A = \varphi(\mathbf{N}, \bar{a}_1)$ *and* $B = \psi(\mathbf{N}, \bar{b}_0)$,
3. $\bar{a}_1 \mathop{\vert}_{\mathbf{M}} \bar{a}_0\bar{b}_0$,
4. $\varphi(z; \bar{a}_0) \wedge \varphi(z, \bar{a}_1)$ *is not satisfiable in* $\mathbf{N}$,
5. $\psi(z; \bar{b}_0)$ *is not satisfiable in* $\mathbf{M}$.

The proof of the lemma is a standard application of basic tools from model theory: compactness, (mutually) indiscernible sequences and Morley sequences, which are recalled in Appendix C. The proof of Lemma 64 is in Appendix D. Using the lemma, we now show that $\mathscr{C}$ is not 1-dimensional.

Through the rest of Section 9.2, we fix $\mathbf{M}$ and $\mathbf{N}$ as in Lemma 64 and use the notation from the lemma.

▷ Claim 65. There is an elementary extension $\mathbf{N}'$ of $\mathbf{N}$, and a tuple $\bar{a}$ of elements in $\varphi(\mathbf{N}'; \bar{a}_1)$, a tuple $\bar{b}$ of elements in $\psi(\mathbf{N}'; \bar{b}_0)$ such that $\bar{a} \mathop{\not\vert}_{\mathbf{M}} \bar{b}$.

Proof. Let $\hat{\theta}(\bar{v}; \bar{u}) = \theta(\bar{u}; \bar{v})$. As $\mathrm{Types}^{\theta}(A/B)$ is infinite by *(2)*, so is $\mathrm{Types}^{\hat{\theta}}(B/A)$.

Let $E_\theta(\bar{u}; \bar{u}')$ be the formula defining the equivalence relation on $\psi(\mathbf{N}; \bar{b}_0)^{\bar{u}}$ such that

$$E_\theta(\bar{b}, \bar{b}') \quad \Leftrightarrow \quad \mathrm{tp}^{\hat{\theta}}(\bar{b}/\varphi(\mathbf{N}; \bar{a}_1)) = \mathrm{tp}^{\hat{\theta}}(\bar{b}'/\varphi(\mathbf{N}; \bar{a}_1)) \qquad \text{for } \bar{b}, \bar{b}' \in \psi(\mathbf{N}; \bar{b}_0)^{\bar{u}}.$$

More precisely,

$$E_\theta(\bar{v}; \bar{v}') \quad \equiv \quad \psi(\bar{v}; \bar{b}_0) \wedge \psi(\bar{v}'; \bar{b}_0) \wedge \forall \bar{u}. \varphi(\bar{u}; \bar{a}_1) \implies (\theta(\bar{u}; \bar{v}) \iff \theta(\bar{u}'; \bar{v})).$$

Hence the formula E_θ defines infinitely many classes in $\mathbf{N}$. By compactness there is an elementary extension $\mathbf{N}'$ of $\mathbf{N}$ such that E_θ induces more than $2^{|\mathbf{M}|}$ equivalence classes in $\mathbf{N}'$ (cf. Lemma 69). As there are at most $2^{|\mathbf{M}|}$ distinct $\hat{\theta}$-types in $\mathbf{M}$, there exist $\bar{b}', \bar{b}'' \in \psi(\mathbf{N}'; \bar{b}_0)$ such that $\mathrm{tp}^{\hat{\theta}}(\bar{b}'/\mathbf{M}) = \mathrm{tp}^{\hat{\theta}}(\bar{b}''/\mathbf{M})$ and $\neg E_\theta(\bar{b}', \bar{b}'')$ holds in $\mathbf{N}'$. Hence $\theta(\bar{a}; \bar{b}') \triangle \theta(\bar{a};\ \bar{b}'')$ holds in $\mathbf{N}'$ for some $\bar{a} \in \varphi(\mathbf{N}'; \bar{a}_1)^{\bar{u}}$. The claim follows for $\bar{b} = \bar{b}'\bar{b}''$. ◀

▷ Claim 66. Let $\mathbf{N}'$ be an elementary extension of $\mathbf{N}$ and $a \in \varphi(\mathbf{N}'; \bar{a}_1)$. Then $\bar{a}_1 \not\mathop{\big|}_{\mathbf{M}} \bar{a}_0 a$, as witnessed by the formula

$$\zeta(\bar{y}; a, \bar{a}_0) := \varphi(a; \bar{y}) \wedge \neg\exists x. \varphi(x; \bar{y}) \wedge \varphi(x; \bar{a}_0). \tag{$*$}$$

Proof. We have that $\zeta(\bar{a}_1; a, \bar{a}_0)$ holds since $\varphi(x; \bar{a}_1)$ and $\varphi(x; \bar{a}_0)$ are inconsistent by *(4)*. Assume that there is some $\bar{a}' \in \mathbf{M}^{\bar{y}}$ such that $\zeta(\bar{a}'; a, \bar{a}_0)$ holds. Then

$$\exists x. \varphi(x; \bar{a}') \wedge \varphi(x; \bar{a}_1)$$

holds in $\mathbf{N}'$, as witnessed by $x = a$. By property *(1)* and as $\bar{a}'$ is in $\mathbf{M}$, this implies that

$$\exists x. \varphi(x; \bar{a}') \wedge \varphi(x; \bar{a}_0)$$

holds in $\mathbf{N}'$, contradicting $\zeta(\bar{a}'; a, \bar{a}_0)$. Thus $\zeta(\bar{y}; a, \bar{a}_0)$ is not satisfiable in $\mathbf{M}$. In particular, $\bar{a}_1 \not\mathop{\big|}_{\mathbf{M}} \bar{a}_0 a$, proving the claim. ◀

Fix tuples $\bar{a}, \bar{b}$ as in Claim 65. We show that if $\mathbf{M}$ is 1-dimensional then $\bar{a}_1\bar{a} \mathop{\big|}_{\mathbf{M}} \bar{a}_0\bar{b}_0\bar{b}$, implying $\bar{a} \mathop{\big|}_{\mathbf{M}} \bar{b}$, contrary to Claim 65.

▷ Claim 67. Suppose $\mathbf{M}$ is 1-dimensional, $\mathbf{N}'$ is an elementary extension of $\mathbf{N}$, and let $\bar{a}$ be a tuple in $\varphi(\mathbf{N}'; \bar{a}_1)$ and $\bar{b}$ a tuple in $\psi(\mathbf{N}'; \bar{b}_0)$. Then

$$\bar{a}_1\bar{a} \mathop{\Big|}_{\mathbf{M}} \bar{a}_0\bar{b}_0\bar{b}.$$

In particular, $\bar{a} \mathop{\big|}_{\mathbf{M}} \bar{b}$.

Proof. We show the statement by induction on the length of $\bar{a}$ and $\bar{b}$. The base case where $\bar{a}$ and $\bar{b}$ are empty is given by property *(3)* in Lemma 64.

Assume we know the result for $\bar{a}, \bar{b}$ and we want to add an element $b \in \psi(\mathbf{N}'; \bar{b}_0)$ to $\bar{b}$. By 1-dimensionality, one of the two cases holds:

$$\bar{a}_1\bar{a}\, b \mathop{\Big|}_{\mathbf{M}} \bar{a}_0\bar{b}_0\bar{b} \qquad \text{or} \qquad \bar{a}_1\bar{a} \mathop{\Big|}_{\mathbf{M}} \bar{a}_0\bar{b}_0\bar{b}\, b.$$

Note that property *(5)* implies $b \not\mathop{\big|}_{\mathbf{M}} \bar{b}_0$, excluding the first case, so the second case must hold, as required.

Now assume we want to add $a \in \varphi(\mathbf{N}'; \bar{a}_1)$ to $\bar{a}$. By 1-dimensionality,

$$\bar{a}_1\bar{a}\, a \mathop{\Big|}_{\mathbf{M}} \bar{a}_0\bar{b}_0\bar{b} \qquad \text{or} \qquad \bar{a}_1\bar{a} \mathop{\Big|}_{\mathbf{M}} \bar{a}_0\bar{b}_0\bar{b}a,$$

but the second possibility is excluded by Claim 66, and the first one concludes the inductive step. This proves Claim 67. ◀

Since there are $\bar{a} \in \varphi(\mathbf{N}'; \bar{a}_1)$ and $\bar{b} \in \psi(\mathbf{N}'; \bar{b}_0)$ with $\bar{a} \not\mathop{\big|}_{\mathbf{M}} \bar{b}$ by Claim 65, Claim 67 implies $\mathbf{M}$ is not 1-dimensional. This finishes the proof of Theorem 54.

A Ordered matchings interpret all graphs

Here we give a proof of the following.

▶ **Lemma 10.** *The class $\mathscr{M}$ of ordered matchings efficiently interprets the class of all graphs.*

Proof. We prove that the class of all graphs interprets in $\mathscr{M}$. Before describing the interpretation of graphs in ordered matchings, we show how the ordered matching M_G corresponding to an ordered graph G is constructed, in polynomial time.

Let G be an ordered graph with vertices $v_1 < \cdots < v_n$ and edges $e_1, \dots, e_m$. For $i \in [n]$ and $1 \leqslant j \leqslant \mathrm{d}(v_i)$ we define $\epsilon_{i,j}$ as the index of the jth edge incident to v_i. The left vertices of M_G will be (in order) $v_1, \dots, v_n, x, {e'_1}^-, e'_1, {e'_1}^+, \dots, {e'_m}^-, e'_m, {e'_m}^+$, and y'. The right vertices of M_G will be (in order) x', $\epsilon_{n,1}, \dots, \epsilon_{n,\mathrm{d}(v_n)}, v'_n, \dots, \epsilon_{1,1}, \dots, \epsilon_{1,\mathrm{d}(v_n)}, v'_1, y, e_m, \dots, e_1$. The matching M_G matches v_i and v'_i, x and x', y and y', e'_i and e_i, and finally $\epsilon_{i,j}$ either with ${e'_{\epsilon_{i,j}}}^-$ or ${e'_{\epsilon_{i,j}}}^+$, depending on whether v_i is the smallest or biggest incidence of $e_{\epsilon_{i,j}}$ (see Figure 18).

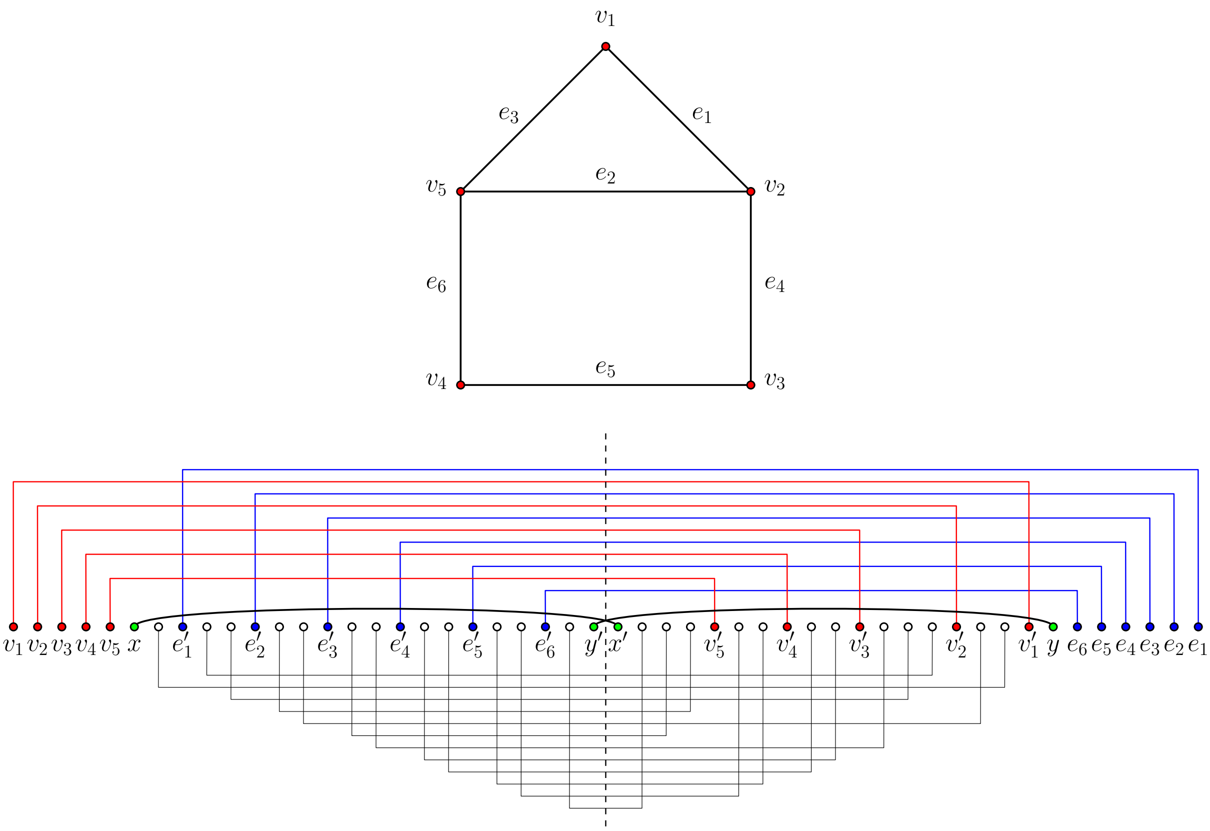

Figure 18 Encoding of a graph in a matching.

We now prove that there is a simple interpretation G, which reconstructs G from M_G. First note that x' is definable as the minimum vertex adjacent to a smaller vertex, and y' is definable as the maximum vertex adjacent to a bigger vertex. Also, x is definable from x' and y is definable from y'. Now we can define $v_1, \dots, v_n$ to be the vertices smaller than x, ordered with the order of M_G. Two vertices $v_i < v_j < x$ are adjacent in the interpretation if there exists an element $e_k > y$ adjacent to a vertex e'_k preceded in the order by an element ${e'_k}^-$ and followed in the order by an element ${e'_k}^+$ with the following properties: ${e'_k}^-$ is adjacent to a vertex z^- strictly between the neighbor v'_i of v_i and the neighbor of the successor of v_i in the order and, similarly, ${e'_k}^+$ is adjacent to a vertex z^+ strictly between the neighbor v'_j of v_j and the neighbor of the successor of v_j in the order. ◀

B Reducing the model checking problem for matrices to structures

In this appendix, we prove:

▶ **Lemma 12.** *Let $\mathscr{C}$ be a class of ordered binary structures and let $\mathcal{M} = \{M(\mathbf{A}) \mid \mathbf{A} \in \mathscr{C}\}$ be the class of adjacency matrices of structures in $\mathscr{C}$. Then there is an* FPT *reduction of the* FO *model checking problem for $\mathcal{M}$ to the* FO *model checking problem for $\mathscr{C}$.*

Proof. Let φ be a sentence in the signature of $\mathcal{M}$. Hence, φ may use the unary relations R and C denoting the rows/columns respectively, the total order $<$, as well as binary relations $E_\tau(x, y)$, for each atomic type $\tau \in A_\Sigma$. Without loss of generality, by rewriting the sentence φ to an equivalent one if necessary, we may assume that the variables are partitioned into *row variables* or a *column variables*, and for every subformula of φ of the form $\exists x.\psi$, the subformula ψ is of the form $R(x) \wedge \psi'(x)$ if x is a row variable, and ψ is of the form $C(x) \wedge \psi'(x)$ if x is a column variable.

In this case, construct from φ a formula φ' in the signature of $\mathscr{C}$ as follows:

- replace each atom $R(x)$ or $C(x)$ by $\top$ (denoting true),
- if x is a row variable and y is a column variable then replace each atom $x = y$ or $y = x$ by $\perp$ (denoting false), each atom $E_\tau(x, y)$ or $E_\tau(y, x)$ by $\tau(x, y)$,
- if x, y are both row variables or both column variables, then replace each atom $E_\tau(x, y)$ by $\perp$.

Then φ holds in $M(\mathbf{A})$ if and only if φ' holds in $\mathbf{A}$. This yields the reduction. ◀

C Model theoretic preliminaries

C.1 Basic notions from model theory

Models and theories. In model theory, structures are called *models*, and we will therefore denote them $\mathbf{M}, \mathbf{N}$, etc. They will typically be infinite.

A (first-order) *theory* is a set T of sentences over a fixed signature. A *model of a theory* T is a model $\mathbf{M}$ (finite or not) which satisfies all the sentences in T, which is denoted $\mathbf{M} \models T$. We say that T *has a model* if there is some model $\mathbf{M}$ of T.

The theory of a class of structures $\mathscr{C}$ is the set T of all sentences φ such that $\mathbf{S} \models \varphi$ for all $\mathbf{S} \in \mathscr{C}$. Trivially, every structure in $\mathscr{C}$ is a model of T, but typically, T has also other models. Those can be constructed using the compactness theorem:

▶ **Theorem 68** (Compactness of first-order logic)**.** *Let T be a theory such that every finite subset $T' \subseteq T$ has a model. Then T has a model.*

For example, let $\mathscr{C}$ be a class of structures over a signature Σ, and assume that $\mathscr{C}$ contains structures of arbitrarily large finite size. Then the models of the theory of $\mathscr{C}$ also include infinite models of arbitrarily large cardinality. To see this, consider the theory T of $\mathscr{C}$ and let Σ' extend the signature of $\mathscr{C}$ by an arbitrary set C of constant symbols. For $c, d \in C$, let φ_{cd} be the Σ'-sentence $c{\neq}d$. Then $T \cup \{\varphi_{cd} \mid c, d \in C, c{\neq}d\}$ satisfies the assumption of the compactness theorem, so it has a model $\mathbf{M}$, and this model has at least the cardinality of C.

Elementary extensions. Let $\mathbf{M}, \mathbf{N}$ be two models such that the domain of $\mathbf{M}$ is contained in the domain of $\mathbf{N}$. Then $\mathbf{N}$ is an *elementary extension* of $\mathbf{M}$, written $\mathbf{M} \prec \mathbf{N}$, if for every formula $\varphi(\bar{x})$ and tuple $\bar{a} \in \mathbf{M}^{\bar{x}}$ of elements of $\mathbf{M}$,

$$\mathbf{M} \models \varphi(\bar{a}) \quad \text{if and only if} \quad \mathbf{N} \models \varphi(\bar{a}).$$

In other words, it doesn't matter if we evaluate formulas in $\mathbf{M}$ or in $\mathbf{N}$.

A typical way of constructing an elementary extension of $\mathbf{M}$ is by considering the following theory, called the *elementary diagram* of $\mathbf{M}$. Let Σ be the signature of $\mathbf{M}$, and let $\Sigma' = \Sigma \cup \mathbf{M}$, where the elements of $\mathbf{M}$ are viewed as constant symbols.

For a Σ-formula $\varphi(\bar{y})$ and tuple $\bar{a} \in \mathbf{M}^{\bar{y}}$ write $\varphi(\bar{a})$ for the Σ'-sentence obtained by replacing the variables in $\bar{y}$ by constants in $\mathbf{M}$, according to $\bar{a}$. Let T be the Σ'-theory consisting of all sentences $\varphi(\bar{a})$, for all Σ-formulas $\varphi(\bar{x})$ and tuples $\bar{a}$ such that $\mathbf{M} \models \varphi(\bar{a})$.

Pick a model $\mathbf{N}'$ of T, and let $\mathbf{N}$ denote the Σ-structure obtained from $\mathbf{N}'$ by forgetting the constants in $\mathbf{M} \subseteq \Sigma'$. The interpretation of the constants $m \in \mathbf{M}$ of Σ' in $\mathbf{N}'$ yields a function $i \colon \mathbf{M} \to \mathbf{N}$. By the definition of T, for any formula $\varphi(\bar{y})$ and tuple $\bar{a} \in \mathbf{M}^{\bar{y}}$,

$$\mathbf{M} \models \varphi(\bar{a}) \text{ if and only if } \mathbf{N} \models \varphi(i(\bar{a})).$$

Therefore, we may view (identyfing each $m \in \mathbf{M}$ with $i(m) \in \mathbf{N}$) the Σ-structure $\mathbf{N}$ as an elementary extension of $\mathbf{M}$.

Reassuming, models of the elementary diagram of $\mathbf{M}$ correspond precisely to elementary extensions of $\mathbf{M}$. In particular, by extending the elementary diagram of $\mathbf{M}$ by an arbitrary set of constants, from compactness we get that $\mathbf{M}$ has elementary extensions of arbitrarily large cardinality (unless $\mathbf{M}$ is finite). More generally, we have the following.

▶ **Lemma 69.** *Let $\mathbf{M}$ be a model and let $\alpha(\bar{x}, \bar{x}')$ be a formula with $|\bar{x}| = |\bar{x}'|$ defining an equivalence relation in $\mathbf{M}$ with infinitely many classes. Then for every cardinality $\mathfrak{n}$ there is an elementary extension $\mathbf{N} \succ \mathbf{M}$ in which α defines an equivalence relation with at least $\mathfrak{n}$ equivalence classes.*

Proof. To simplify notation, assume that $|\bar{x}| = |\bar{x}'| = 1$. The case of $|\bar{x}| = |\bar{x}'| = k > 1$ proceeds similarly, or can be deduced from the case $k = 1$ by extending the domain of $\mathbf{M}$ by $\mathbf{M}^k$ and the k projection functions.

Let $\alpha(x, x')$ be formula defining an equivalence relation $\sim$ in $\mathbf{M}$ with infinitely many classes. Let Σ be the signature of $\mathbf{M}$. Fix any set of constants C and let Σ' extend Σ by $C \cup \mathbf{M}$, where all the added elements are constant symbols. For any $c, d \in C$ consider the Σ'-sentence $\varphi_{cd} = \neg\alpha(c, d)$. Let T be the Σ'-theory consisting of:

- the sentences φ_{cd}, for all $c{\neq}d$ in C,
- the elementary diagram of $\mathbf{M}$.

We show that every $T' \subseteq T$ containing finitely many sentences of the form φ_{cd} has a model. Let $C' \subseteq C$ be the finite set of constants appearing in the sentences $\varphi_{cd} \in T'$. Let $\mathbf{M}'$ be the model $\mathbf{M}$ together with each constant c in $\mathbf{M} \subseteq \Sigma'$ interpreted as the corresponding element $c \in \mathbf{M}$, and constants in C' interpreted as pairwise $\sim$-inequivalent elements of $\mathbf{M}$, and constants in $C \setminus C'$ interpreted as arbitrary elements of $\mathbf{M}$. This can be done, since there are infinitely many pairwise $\sim$-inequivalent elements in $\mathbf{M}$. This shows that T' has a model.

By compactness, T has a model $\mathbf{N}'$. This model can be seen as an elementary extension of T together with a set of $|C|$ elements which are pairwise inequivalent with respect to the equivalence relation defined by α in $\mathbf{N}$. Since C was taken arbitrary, this proves the lemma. ◀

Parameters. Let $\mathbf{M}$ be a model over a signature Σ and let $A \subseteq \mathbf{M}$ be a set of elements. We may view $\mathbf{M}$ as a model over a signature $\Sigma \cup A$, where the elements of A are seen as constant symbols, interpreted in $\mathbf{M}$ in the expected way: a constant $a \in A$ is interpreted as the element $a \in \mathbf{M}$. We call the elements of A *parameters* in this context. A Σ-formula *with parameters* from A is a formula over the signature $\Sigma \cup A$.

Types. A *type* with variables $\bar{x}$ and parameters from A, or a *type over* A is a set p of formulas $\varphi(\bar{x})$ with parameters from A. We may write $p(\bar{x})$ to indicate that p has variables $\bar{x}$.

If $p(\bar{x})$ is a type over A and $B \subseteq A$ then $p|B$ denotes the subset of p consisting of all formulas with parameters from B. If $\bar{b} \in \mathbf{M}^{\bar{x}}$ is a tuple of elements of $\mathbf{M}$ then *the type* of $\bar{b}$ *over* A in $\mathbf{M}$ is the set of formulas $\varphi(\bar{x})$ with parameters from A that are satisfied by $\bar{b}$ in $\mathbf{M}$. This type is denoted $\mathrm{tp}(\bar{b}/A)$ or $\mathrm{tp}_{\bar{x}}(\bar{b}/A)$. Note that $\mathrm{tp}(\bar{b}/A)$ is related to the notion of θ-types as follows, for every formula $\theta(\bar{x}; \bar{y})$ and tuple $\bar{a} \in A^{\bar{y}}$:

$$\theta(\bar{x}; \bar{a}) \in \mathrm{tp}(\bar{b}/A) \iff \bar{a} \in \mathrm{tp}^{\theta}(\bar{b}/A).$$

In particular, $\mathrm{tp}(\bar{b}/A)$ is uniquely determined by $\{\mathrm{tp}^\theta(\bar{b}/A) \mid \theta(\bar{x};\bar{y}) \text{ is a formula}\}$.

A type $p(\bar{x})$ is *satisfiable in a set* C if there is some tuple $\bar{c} \in C^{\bar{x}}$ which satisfies all the formulas in p. A type $p(\bar{x})$ with parameters from $A \subseteq \mathbf{M}$ is *satisfiable* if it is satisfiable in some elementary extension $\mathbf{N}$ of $\mathbf{M}$. By compactness, this is equivalent to saying that for any finite conjunction $\varphi(\bar{x})$ of formulas in $p(\bar{x})$ we have $\mathbf{M} \models \exists\bar{x}.\varphi(\bar{x})$.

A type $p(\bar{x})$ with parameters from A is *complete* if it is satisfiable and for every formula $\varphi(\bar{x})$ with parameters from A, either φ or $\neg\varphi$ belongs to p. Equivalently, $p(\bar{x})$ is the type over A of some tuple $\bar{b} \in \mathbf{N}^{\bar{x}}$, for some elementary extension $\mathbf{N}$ of $\mathbf{M}$. We sometimes say that a type is *partial* to emphasise that it may not be complete. We denote the set of complete types with variables $\bar{x}$ and parameters from A by $\mathrm{S}^{\bar{x}}(A)$ or simply $\mathrm{S}(A)$, if $\bar{x}$ is understood from the context. Note that we have ommitted the model $\mathbf{M}$ from the notation. Indeed, if $\mathbf{M}'$ is a model containing the parameters A and satisfying the same sentences with parameters from A as $\mathbf{M}$, then $\mathbf{M}$ and $\mathbf{M}'$ have identical sets of complete types $p(\bar{x})$ with parameters from A. Hence, $\mathrm{S}^{\bar{x}}(A)$ does not depend on $\mathbf{M}$, but only on the set of sentences satisfied by A in $\mathbf{M}$.

C.2 Finite satisfiability

A (partial) type $p(\bar{x})$ with parameters from A is *finitely satisfiable* in C if every finite subset $p' \subseteq p$ is satisfiable in C. Note that $\bar{a} \mathop{\big|}_{\mathbf{M}} B$ (cf. Def. 57) if and only if $\mathrm{tp}(\bar{a}/\mathbf{M}B)$ is finitely satisfiable in $\mathbf{M}$.

▶ **Lemma 70.** *A type $p(\bar{x})$ with parameters from $\mathbf{M}$ is finitely satisfiable in $\mathbf{M}$ if and only if it is satisfiable. Consequently, $\bar{a} \mathop{\big|}_{\mathbf{M}} \mathbf{M}$ for all $\bar{a}$ in an elementary extension of $\mathbf{M}$.*

Proof. For the right-to-left implication, assume that p is satisfied by some tuple $\bar{c} \in \mathbf{N}^{\bar{x}}$ for some elementary extension $\mathbf{N}$ of $\mathbf{M}$. Pick a finite $p' \subseteq p$, and suppose $p' = \{\varphi_1, \ldots, \varphi_k\}$. Consider the formula $\psi := \varphi_1 \wedge \cdots \wedge \varphi_k$. Note that ψ may use some parameters from $\mathbf{M}$. So we may write ψ as $\psi = \psi'(\bar{x}, \bar{a})$ where $\psi'(\bar{x}, \bar{z})$ is a formula and $\bar{a} \in \mathbf{M}^{\mathbf{z}}$.

The formula $\exists_{\bar{x}}\psi'(\bar{x}, \bar{a})$ holds in $\mathbf{N}$, as witnessed by $\bar{c}$. As $\mathbf{N}$ is an elementary extension of $\mathbf{M}$, this formula also holds in $\mathbf{M}$. So there is some $\bar{m} \in \mathbf{M}$ satisfying $\psi'(\bar{b}, \bar{a})$. Therefore, p' is satisfied by $\bar{m}$ in $\mathbf{M}$, proving that p is finitely satisfiable in $\mathbf{M}$.

The left-to-right implicaiton is a basic application of the compactness theorem.

Consider the signature $\Sigma' = \Sigma \cup \mathbf{M} \cup \bar{x}$ extending Σ by constant symbols for each element of $\mathbf{M}$ and each variable in $\bar{x}$. Let T be the theory over Σ' consisting of:

- For every formula $\varphi(\bar{x}) \in p$, the Σ'-sentence obtained from $\varphi(\bar{x})$ by viewing each parameter $a \in \mathbf{M}$ as the constant $a \in \mathbf{M} \subseteq \Sigma'$, and each variable $x \in \bar{x}$ as the constant $x \in \bar{x} \subseteq \Sigma'$.
- the elementary diagram of $\mathbf{M}$.

Then every finite subset T' of T has a model. Indeed, let p' be the set of formulas $\varphi(\bar{x})$ which occur (as Σ'-sentences) in T'. Since $p(\bar{x})$ is finitely satisfiable in $\mathbf{M}$, $p'(\bar{x})$ is satisfied by some tuple $\bar{m} \in \mathbf{M}^{\bar{x}}$. The pair $(\mathbf{M}, \bar{m})$ may be seen as a Σ'-structure, where a constant $m \in \mathbf{M}$ is interpreted by the corresponding element of $\mathbf{M}$, and a constant $x \in \bar{x}$ is interpreted as $\bar{m}(x)$. Then $(\mathbf{M}, \bar{m})$ is a model of T'.

By compactness, T has a model $\mathbf{N}'$. This model can be seen as an elementary extension $\mathbf{N}$ of $\mathbf{M}$ together with a tuple $\bar{c} \in \mathbf{N}^{\bar{x}}$ of elements (obtained by the interpretation of the constants $\bar{x}$ in $\mathbf{N}'$), such that $\mathbf{N} \models \varphi(\bar{c})$ for every formula $\varphi(\bar{x}) \in p$. Hence, $\bar{c}$ satisfies $p(\bar{x})$ in $\mathbf{N}$. ◀

Finite satisfiability and filters. Recall that a *filter* on a set U is a nonempty set $F \subseteq P(U)$ that is closed under taking supersets (if $A \subseteq B$ then $A \in F$ implies $B \in F$), under binary intersections, and does not contain the empty set. A filter is an *ultrafilter* if for every $A \subseteq U$, either $A \in F$ or $U \setminus A \in F$. Every filter is contained in some ultrafilter, by the Kuratowski-Zorn lemma.

Let $\mathbf{N}$ be a model, $A \subseteq \mathbf{N}$ be a set and $\bar{x}$ be a set of variables. Fix a filter F on $A^{\bar{x}}$. The *average (partial) type* over $\mathbf{N}$ is the partial type denoted $\mathrm{Av}_F(\bar{x})$ such that for every formula $\varphi(\bar{x})$ with parameters from $\mathbf{N}$,

$$\varphi(\bar{x}) \in \mathrm{Av}_F(\bar{x}) \iff \{\bar{a} \in A^{\bar{x}} : \mathbf{N} \models \varphi(\bar{a})\} \in F.$$

This is a consistent partial type: if say $\varphi_1(\bar{x}), \dots, \varphi_n(\bar{x}) \in \pi(\bar{x})$, then since any finitely many elements of F have non-empty intersection, there is $\bar{a} \in A^{\bar{x}}$ which satisfies the conjunction $\varphi_1(\bar{x}) \wedge \cdots \wedge \varphi_n(\bar{x})$. Hence this conjunction is consistent, indeed we have shown that $\mathrm{Av}_F(\bar{x})$ is finitely satisfiable in A.

If F is an ultrafilter on $A^{\bar{x}}$, then $\mathrm{Av}_F(\bar{x})$ is a complete type: for every formula $\varphi(\bar{x})$, either $\varphi(\bar{x}) \in \mathrm{Av}_F(\bar{x})$ or $\neg\varphi(\bar{x}) \in \mathrm{Av}_F(\bar{x})$.

▶ **Lemma 71.** *Let $\pi(\bar{x})$ be a partial type, then $\pi(\bar{x})$ is finitely satisfiable in A if and only if there is a filter F on $A^{\bar{x}}$ such that $\pi(\bar{x}) \subseteq \mathrm{Av}_F(\bar{x})$.*

Proof. We have already observed that $\mathrm{Av}_F(\bar{x})$ is finitely satisfiable in A. Conversely, assume that $\pi(\bar{x})$ is finitely satisfiable in A, then define $F_0 \subseteq P(A^{\bar{x}})$ by: $F_0 = \{\varphi(A) : \varphi(\bar{x}) \in \pi(\bar{x})\}$. The fact that $\pi(\bar{x})$ is finitely satisfiable in A implies that any finitely many elements of F_0 have non-empty intersection. Let F be the filter generated by F_0. Then we have $\pi(\bar{x}) \subseteq \mathrm{Av}_F(\bar{x})$. ◀

▶ **Lemma 72.** *Let $p(\bar{x}) \in \mathrm{S}(\mathbf{N})$ be a complete type, then $p(\bar{x})$ is finitely satisfiable in A if and only if there is an ultrafilter F on $A^{\bar{x}}$ such that $p(\bar{x}) = \mathrm{Av}_F(\bar{x})$.*

Proof. We have already seen that if F is an ultrafilter on $A^{\bar{x}}$, then $\mathrm{Av}_F(\bar{x})$ is a complete type over $\mathbf{N}$, which is finitely satisfiable in A. Conversely, if $p(\bar{x}) \in \mathrm{S}(\mathbf{N})$ is finitely satisfiable in A, then by the previous lemma, there is a filter F_0 on $A^{\bar{x}}$ such that $p(\bar{x}) \subseteq \mathrm{Av}_{F_0}(\bar{x})$. Extend F_0 to an ultrafilter F on $A^{\bar{x}}$. Then $p(\bar{x}) \subseteq \mathrm{Av}_{F_0}(\bar{x}) \subseteq \mathrm{Av}_F(\bar{x})$. But since $p(\bar{x})$ is a complete type, one cannot add any formulas to it without making it inconsistent. Since $\mathrm{Av}_F(\bar{x})$ is consistent, we must have $p(\bar{x}) = \mathrm{Av}_F(\bar{x})$. ◀

▶ **Lemma 73.** *Let $\pi(\bar{x})$ be a partial type finitely satisfiable in A. Then there is a complete type $p(\bar{x}) \in \mathrm{S}(\mathbf{N})$ finitely satisfiable in A which extends $\pi(\bar{x})$.*

Proof. Let F be a filter on $A^{\bar{x}}$ such that $\pi(\bar{x}) \subseteq \mathrm{Av}_F(\bar{x})$. Let F' be an ultrafilter extending F and let $p(\bar{x}) = \mathrm{Av}_{F'}(\bar{x})$. Then p is finitely satisfiable in A and extends π. ◀

▶ **Lemma 74.** *Let $p(\bar{x}) \in \mathrm{S}(\mathbf{N})$ be finitely satisfiable in A. Then p is A-*invariant*, that is: for any formula $\varphi(\bar{x}; \bar{y})$ and tuples $\bar{b}, \bar{b}' \in \mathbf{N}^{\bar{y}}$, we have:*

$$\mathrm{tp}(\bar{b}/A) = \mathrm{tp}(\bar{b}'/A) \Longrightarrow \varphi(\bar{x}; \bar{b}) \in p \leftrightarrow \varphi(\bar{x}; \bar{b}') \in p.$$

Proof. If $\mathrm{tp}(\bar{b}/A) = \mathrm{tp}(\bar{b}'/A)$, then the formula $\varphi(\bar{x}; \bar{b}) \triangle \varphi(\bar{x}; \bar{b}')$ has no solution in A. Since p is finitely satisfiable in A that formula cannot be in p. Hence as p is a complete type, the formula $\varphi(\bar{x}; \bar{b}) \leftrightarrow \varphi(\bar{x}; \bar{b}')$ is in p as required. ◀

C.3 Indiscernible sequences

▶ **Definition 75.** *Let $\mathbf{M}$ be a structure and $A \subseteq \mathbf{M}$. Let I be a linear order. A sequence $(\bar{a}_i : i \in I)$ of tuples of $\mathbf{M}$ is* indiscernible over A *if for any $n < \omega$ and indices*

$$i_1 < \cdots < i_n \qquad i'_1 < \cdots < i'_n$$

in I, we have

$$\mathrm{tp}(a_{i_1}, \dots, a_{i_n}/A) = \mathrm{tp}(a_{i'_1}, \dots, a_{i'_n}/A).$$

Another way to state this is that the sequence $(\bar{a}_i : i \in I)$ is indiscernible over A if for any $n < \omega$, indices

$$i_1 < \cdots < i_n \qquad i'_1 < \cdots < i'_n$$

in I and formula $\theta(\bar{x}_1, \dots, \bar{x}_n)$ with parameters in A, we have

$$\text{(1)} \qquad \mathbf{M} \models \theta(\bar{a}_{i_1}, \dots, \bar{a}_{i_n}) \leftrightarrow \theta(\bar{a}_{i'_1}, \dots, \bar{a}_{i'_n}).$$

If Δ is a set of formulas with parameters in A, we will say that the sequence $(\bar{a}_i : i \in I)$ is Δ-*indiscernible* if (1) holds for each θ in Δ. If Δ and I are both finite, then this is expressible by a single first order formula.

▶ **Definition 76.** *Two sequences* $(\bar{a}_i : i \in I)$ *and* $(\bar{b}_j : j \in J)$ *are* mutually indiscernible *over* A *if* $(\bar{a}_i : i \in I)$ *is indiscernible over* $A \cup \{\bar{b}_j : j \in J\}$ *and* $(\bar{b}_j : j \in J)$ *is indiscernible over* $A \cup \{\bar{a}_i : i \in I\}$*.*

An equivalent definition is that the sequences $(\bar{a}_i : i \in I)$ and $(\bar{b}_j : j \in J)$ are *mutually indiscernible* over A if for any $n < \omega$, indices

$$i_1 < \cdots < i_n \qquad i'_1 < \cdots < i'_n$$

and

$$j_1 < \cdots < j_n \qquad j'_1 < \cdots < j'_n$$

in I and any formula $\theta(\bar{x}_1, \dots, \bar{x}_n; \bar{y}_1, \dots, \bar{y}_n)$ with parameters in A, we have

$$(2) \qquad \mathbf{M} \models \theta(\bar{a}_{i_1}, \dots, \bar{a}_{i_n}; \bar{b}_{j_1}, \dots, \bar{b}_{j_n}) \leftrightarrow \theta(\bar{a}_{i'_1}, \dots, \bar{a}_{i'_n}; \bar{b}_{j'_1}, \dots, \bar{b}_{j'_n}).$$

If Δ is a set of formulas with parameters in A, we will say that the sequences $(\bar{a}_i : i \in I)$ and $(\bar{b}_j : j < \omega)$ are *mutually Δ-indiscernible* if (2) holds for each θ in Δ. If Δ, I and J are finite, then this is again expressible by a single first-order formula.

In the following lemma, we use the notation $\mathrm{Av}_F|C$ to mean the restriction of the type Av_F to C. We also use the notation $\bar{a}_{<i}$ to mean $\{\bar{a}_j : j < i\}$.

▶ **Lemma 77.** *Let* $A \subseteq B \subseteq \mathbf{M}$*. Let* F *be an ultrafilter on* $A^{\bar{x}}$*. Let* I *be a linear order and let* $(\bar{a}_i : i \in I)$ *be a sequence of tuples of* $\mathbf{M}$ *such that:*

$$\bar{a}_i \models \mathrm{Av}_F|B\bar{a}_{<i}.$$

Then the sequence $(\bar{a}_i : i \in I)$ *is indiscernible over* B*.*

Proof. Write $p = \mathrm{Av}_F$. Note that p is finitely satisfiable in A and a fortiori finitely satisfiable in B.

We prove by induction on n that if $n < \omega$ and $i_1 < \cdots < i_n$, $j_1 < \cdots < j_n$ are in I, then $\mathrm{tp}(\bar{a}_{i_1}, \dots, \bar{a}_{i_n}/A) = \mathrm{tp}(\bar{a}_{j_1}, \dots, \bar{a}_{j_n}/A)$. For $n = 1$ this follows from the fact that all $\bar{a}_i$ realize $\mathrm{Av}_F|B$, which is a complete type over B. Assume we know it for n and take $i_1 < \cdots < i_n < i_{n+1}$, $j_1 < \cdots < j_n < j_{n+1}$ in I. By induction hypothesis, we have

$$\mathrm{tp}(\bar{a}_{i_1}, \dots, \bar{a}_{i_n}/B) = \mathrm{tp}(\bar{a}_{j_1}, \dots, \bar{a}_{j_n}/B).$$

By Lemma 74, for any formula $\theta(\bar{x}; \bar{y}_1, \dots, \bar{y}_n)$ with parameters in B, we have:

$$\theta(\bar{x}; \bar{a}_{i_1}, \dots, \bar{a}_{i_n}) \in p \iff \theta(\bar{x}; \bar{a}_{j_1}, \dots, \bar{a}_{j_n}) \in p.$$

Now since $\bar{a}_{i_{n+1}} \models p|Ba_{i_1} \dots a_{i_n}$, we have

$$\theta(\bar{x}; \bar{a}_{i_1}, \dots, \bar{a}_{i_n}) \in p \iff \mathbf{N} \models \theta(\bar{a}_{i_{n+1}}; \bar{a}_{i_1}, \dots, \bar{a}_{i_n}),$$

and similarly since $\bar{a}_{j_{n+1}} \models p|Ba_{j_1} \dots a_{j_n}$, we have

$$\theta(\bar{x}; \bar{a}_{j_1}, \dots, \bar{a}_{j_n}) \in p \iff \mathbf{N} \models \theta(\bar{a}_{j_{n+1}}; \bar{a}_{j_1}, \dots, \bar{a}_{j_n}).$$

Putting all of this together, we get

$$\mathbf{N} \models \theta(\bar{a}_{i_{n+1}}; \bar{a}_{i_1}, \dots, \bar{a}_{i_n}) \iff \mathbf{N} \models \theta(\bar{a}_{j_{n+1}}; \bar{a}_{j_1}, \dots, \bar{a}_{j_n}).$$

Since the formula θ was an arbitrary formula with parameters in B, we deduce

$$\mathrm{tp}(\bar{a}_{i_1}, \dots, \bar{a}_{i_{n+1}}/B) = \mathrm{tp}(\bar{a}_{j_1}, \dots, \bar{a}_{j_{n+1}}/B)$$

as required. ◀

▶ **Definition 78.** *Let the type* $p(\bar{x}) \in \mathrm{S}(\mathbf{M})$ *be finitely satisfiable in* $A \subseteq \mathbf{M}$ *and let* $A \subseteq B \subseteq \mathbf{M}$*. A sequence* $(\bar{a}_i : i \in I)$ *of tuples in* $\mathbf{M}^{\bar{x}}$ *such that* $\bar{a}_i \models p|B\bar{a}_{<i}$ *is called a* Morley sequence *of* p *over* B*.*

By the previous lemma, a Morley sequence of p over B is indiscernible over B.

C.4 Building indiscernible sequences

Indiscernible sequences are easy to find thanks to Ramsey's theorem.

▶ **Definition 79.** *Let $(\bar{a}_i : i < \omega)$ be a sequence of tuples in some structure $\mathbf{M}$. A family $(\bar{b}_i : i \in I)$ indexed by a linear order I is* based on $(\bar{a}_i)_{i<\omega}$ *if for any formula $\theta(x_1, \ldots, x_n) \in L$ and $i_1 < \ldots < i_n$ in I, if $\mathbf{M} \models \theta(b_{i_1}, \ldots, \bar{b}_{i_n})$ then there are $j_1 < \ldots < j_n < \omega$ such that $\mathbf{M} \models \theta(\bar{a}_{j_1}, \ldots, \bar{a}_{j_n})$.*

Note that if $(\bar{a}_i : i < \omega)$ is indiscernible and $(\bar{b}_i : i \in I)$ is based on it, then it is also indiscernible: indeed for any $i_1 < \ldots < i_n$ in I and any $j_1 < \ldots < j_n < \omega$, we have

$$\text{tp}(\bar{b}_{i_1}, \ldots, \bar{b}_{i_n}) = \text{tp}(\bar{a}_{j_1}, \ldots, \bar{a}_{j_n}).$$

▶ **Proposition 80.** *Let $(\bar{a}_i : i < \omega)$ be a sequence of tuples in some structure $\mathbf{M}$ and let I be any linearly ordered set. There is an elementary extension $\mathbf{M} \prec \mathbf{N}$ and a sequence $(\bar{b}_i : i \in I)$ of tuples of $\mathbf{N}$ that is based on $(\bar{a}_i)_{i<\omega}$.*

Proof. Follows from Ramsey and compactness. ◀

We have analogues for two sequences.

▶ **Definition 81.** *Let $(\bar{a}_i : i < \omega)$ and $(\bar{a}'_i : i < \omega)$ be two sequences of tuples in $\mathbf{M}$. Two families $(\bar{b}_i : i \in I)$, $(\bar{b}'_j : j \in J)$ indexed by linear orders I and J are* based on $(\bar{a}_i)_{i<\omega}$ and $(\bar{a}'_i)_{i<\omega}$ *if for any formula $\theta(\bar{x}_1, \ldots, \bar{x}_n; \bar{y}_1, \ldots, \bar{y}_m) \in L$ and $i_1 < \ldots < i_n$ in I and $j_1 < \ldots < j_m$ in J, if $\mathbf{M} \models \theta(\bar{b}_{i_1}, \ldots, \bar{b}_{i_n}; \bar{b}'_{j_1}, \ldots, \bar{b}'_{j_n})$ then there are $k_1 < \ldots < k_n < \omega$ and $k'_1 < \ldots < k'_m < \omega$ such that $\mathbf{M} \models \theta(\bar{a}_{k_1}, \ldots, \bar{a}_{k_n}; \bar{a}_{k'_1}, \ldots, \bar{a}_{k'_m})$.*

Here is a finitary version of Proposition 80 for two sequences.

▶ **Lemma 82.** *Let Δ be a finite set of formulas and let $m, d < \omega$. Then there is some $m_* < \omega$ such that if $(\bar{a}_i : i < m_*)$ and $(\bar{b}_i : i < m_*)$ are two sequences of d-tuples of a structure $\mathbf{M}$, then there are $(\bar{a}'_i : i < m)$ and $(\bar{b}'_i : i < m)$ subsequences of $(\bar{a}_i)_{i<m_*}$ and $(\bar{b}_i)_{i<m_*}$ respectively such that the sequences $(\bar{a}'_i : i < m)$ and $(\bar{b}'_i : i < m)$ are mutually Δ-indiscernible.*

▶ **Proposition 83.** *Let $(\bar{a}_i : i < \omega)$ and $(\bar{a}'_i : i < \omega)$ be two sequences of tuples in $\mathbf{M}$ and let I, J be two linearly ordered sets. There is an elementary extension $\mathbf{M} \prec \mathbf{N}$ and sequences $(\bar{b}_i : i \in I)$ and $(\bar{b}'_j : j \in J)$ of tuples of $\mathbf{N}$ which are based on $(\bar{a}_i)_{i<\omega}$ and $(\bar{a}'_i)_{i<\omega}$.*

Proof. Follows from Lemma 82 and compactness. ◀

▶ **Lemma 84.** *Let $\mathbf{N}$ be a model and $I = (\bar{a}_i : i < \omega)$ an indiscernible sequence of tuples of $\mathbf{N}$. There is an elementary extension $\mathbf{N} \prec \mathbf{N}'$, a submodel $\mathbf{M} \prec \mathbf{N}'$ and an ultrafilter F on $\mathbf{M}^{\bar{x}}$ such that I is a Morley sequence of Av_F over $\mathbf{M}$.*

Proof. In an elementary extension of $\mathbf{N}$, we can increase the sequence to $I + J$, where $J = (\bar{b}_i : i \in \mathbb{Z})$ so that the sequence $I + J$ is indiscernible. Let F_0 be an ultrafilter on $\{\bar{b}_i \mid i \in \mathbb{Z}\}$ that contains all subsets of the form $\{\bar{b}_i \mid i < n\}$ for $n \in \mathbb{Z}$. It follows from indiscernibility that the sequence I is a Morley sequence of Av_{F_0} over $\{\bar{b}_i \mid i \in \mathbb{Z}\}$. Possibly up to passing to a further elementary extension, we can find an elementary submodel $\mathbf{M}$ such that I is a Morley sequence of Av_{F_0} over $\mathbf{M}$. One can see that by compactness, or alternatively, take $\mathbf{M}_0$ any model containing $\{\bar{a}'_i \mid i \in \mathbb{Z}\}$, let $I' = (\bar{a}'_i : i < \omega)$ be a Morley sequence of Av_{F_0} over M_0. Now I has the same type as I' over J, so passing to an elementary extension, there is an automorphism σ fixing $\{\bar{b}_i \mid i \in \mathbb{Z}\}$ pointwise and sending I' to I. Then take $M = \sigma(M_0)$.

Finally, define F to be the unique ultrafilter on M extending F_0 (so a set A is in F if and only if it contains a set in F_0). Then I is a Morley sequence of Av_F over M. ◀

▶ **Lemma 85.** *Let $\mathbf{N}$ be a structure and let $I = (\bar{a}_i : i < \omega)$ and $J = (\bar{b}_j : j < \omega)$ two mutually indiscernible sequences of tuples of $\mathbf{N}$. There is an elementary extension $\mathbf{N} \prec \mathbf{N}'$, a submodel $\mathbf{M} \prec \mathbf{N}'$ two ultrafilters F and F' on $\mathbf{M}^{\bar{x}}$ such that I is a Morley sequence of Av_F over $\mathbf{M} \cup \{\bar{a}_i : i < \omega\}$ and J is a Morley sequence of $\text{Av}_{F'}$ over $\mathbf{M} \cup \{\bar{b}_j : j < \omega\}$.*

Proof. The proof is very similar to the previous one. First, in an elementary extension, construct sequences $I' = (\bar{a}'_i : i \in \mathbb{Z})$ and $J' = (\bar{b}'_j : j \in \mathbb{Z})$ so that the two sequences $I + I'$ and $J + J'$ are mutually indiscernible. This is possible by compactness. Let F be an ultrafilter on $\{\bar{a}'_i \mid i \in \mathbb{Z}\}$ containing all initial segments as in the previous proof and similarly for F' on $\{\bar{b}'_j \mid j \in \mathbb{Z}\}$. Then I is a Morley sequence of Av_F over $\{\bar{a}'_i \mid i \in \mathbb{Z}\} \cup \{\bar{b}_j \mid j < \omega\} \cup \{\bar{b}'_j \mid j \in \mathbb{Z}\}$ and J is a Morley sequence of $\mathrm{Av}_{F'}$ over $\{\bar{a}_i \mid i < \omega\} \cup \{\bar{a}'_i \mid i \in \mathbb{Z}\} \cup \{\bar{b}'_j \mid j \in \mathbb{Z}\}$. One can then construct the model M as above. ◀

D Proof of Lemma 64

A family $(\varphi_i(x))_{i \in I}$ of formulas with parameters from $\mathbf{N}$ is *pairwise inconsistent* if for any distinct $i, j \in I$, the formula $\varphi_i(x) \wedge \varphi_j(x)$ has no solution in $\mathbf{N}$. For a sequence $\{\bar{a}_i \mid i < \omega\}$ and for $i \leqslant \omega$ by $\bar{a}_{<i}$ denote the set of elements in all the tuples $\bar{a}_j$ with $j < i$.

We prove a stronger variant of Lemma 64.

▶ **Lemma 86.** *Suppose $\mathscr{C}$ is not restrained. Then there exist:*

- *Formulas $\varphi(x, \bar{y}), \psi(x, \bar{z}), \theta(\bar{u}, \bar{v})$,*
- *a structure $\mathbf{M}$ in the elementary closure of $\mathscr{C}$,*
- *an elementary extension $\mathbf{N}$ of $\mathbf{M}$,*
- *a sequence $(\bar{a}_i : i < \omega)$ of tuples in $\mathbf{N}^{\bar{y}}$ and a sequence $(\bar{b}_j : j < \omega)$ of tuples in $\mathbf{N}^{\bar{z}}$,*

such that the following properties hold:

1. *the tuples $\bar{a}_0, \bar{a}_1, \dots$ have equal types over $\mathbf{M}$, and the tuples $\bar{b}_0, \bar{b}_1, \dots$ have equal types over $\mathbf{M}$,*
2. *for all $0 < i, j < \omega$, the set $\mathrm{Types}^{\theta}(A/B)$ is infinite, where $A = \varphi(\mathbf{N}, \bar{a}_i)$ and $B = \psi(\mathbf{N}, \bar{b}_i)$,*
3. *$\bar{a}_i \mathop{\big|}_{\mathbf{M}} \bar{a}_{<i}\bar{b}_{<\omega}$ for $i < \omega$,*
4. *the formulas $\{\varphi(x; \bar{a}_i) \mid i < \omega\}$ are pairwise inconsistent,*
5. *the formulas $\{\psi(x; \bar{b}_j) \mid j < \omega\}$ are pairwise inconsistent.*

It is clear that each of the properties (1)-(4) implies the corresponding property stated in Lemma 64. Properties (5) and (1) together imply that $\psi(\mathbf{M}, \bar{b}_j) = \emptyset$ yielding property (5) in Lemma 64. We thus prove the statement above.

Proof. Assume $\mathscr{C}$ is not restrained. Then there are formulas $\varphi(x, \bar{y})$ and $\psi(y, \bar{z})$ and a finite set of formulas $\Delta(\bar{u}, \bar{v})$ such that for every $n \in \mathbb{N}$ there is a structure $\mathbf{S} \in \mathscr{C}$, two disjoint families $\mathcal{R}, \mathcal{L}$ of size k with $\mathrm{Types}^{\Delta}(L/R) \geqslant k$ for all $L \in \mathcal{L}$ and $R \in \mathcal{R}$, where $\mathcal{L}$ is φ-definable and $\mathcal{R}$ is ψ-definable. We proceed in two steps.

Step 1. There is a model $\mathbf{M}$ in the elementary closure of $\mathscr{C}$, indiscernible sequences $(\bar{a}_i : i < \omega)$ in $\mathbf{M}^{\bar{y}}$ and $(\bar{b}_j : j < \omega)$ in $\mathbf{M}^{\bar{z}}$, formulas $\varphi(x, \bar{y})$, $\psi(x, \bar{z})$ and $\theta(\bar{u}, \bar{v}) \in \Delta$ such that:

- the families $\{\varphi(x, \bar{a}_i) \mid i < \omega\}$ and $\{\psi(x, \bar{b}_j) \mid j < \omega\}$ are both pairwise inconsistent
- for each $0 \leqslant i, j < \omega$ the set $\mathrm{Types}^{\Delta}(A/B)$ is infinite, where $A = \varphi(\mathbf{M}, \bar{a}_i)$ and $B = \psi(\mathbf{M}, \bar{b}_j)$.

As $\mathscr{C}$ is not restrained, for every natural number m, we can find a structure $\mathbf{M}_m \in \mathscr{C}$ sequences $(\bar{a}^m_i : i < m)$ and $(\bar{b}^m_j : j < m)$ of tuples of $\mathbf{M}_m$ such that:

• the two families $\{\varphi(\bar{x}; \bar{a}^m_i) : i < m\}$ and $\{\psi(x; \bar{b}^m_j) : j < m\}$ are pairwise disjoint;

• for every $i, j < m$, the set $\mathrm{Types}^{\Delta}(A/B)$ has size at least m, where $A = \varphi(\mathbf{M}_m; \bar{a}^m_i)$ and $B = \psi(\mathbf{M}_m; \bar{b}^m_j)$.

Add constants to the signature to name two sequences $(\bar{a}_i : i < \omega)$ and $(\bar{b}_j : j < \omega)$. Consider the theory T' in the extended language consisting of the following for every $m < \omega$:

$\bullet_0$ all sentences which hold in all structures in $\mathscr{C}$;

$\bullet_{1,m}$ for every $i < j < m$, the two sets $\varphi(\bar{x}; \bar{a}_i)$ and $\varphi(\bar{x}; \bar{a}_j)$ are disjoint and the two sets $\psi(x; \bar{b}_i)$ and $\psi(x; \bar{b}_j)$ are disjoint;

$\bullet_{2,m}$ the two sequences $(\bar{a}_i : i < m)$ and $(\bar{b}_j : j < m)$ are indiscernible;

$\bullet_{3,m}$ for every $i,j < m$ the set $\mathrm{Types}^{\Delta}(A/B)$ has size at least m, where $A = \varphi(\mathbf{M};\bar{a}^m)$ and $B = \psi(\mathbf{M};\bar{b}^m)$ and $\mathbf{M}$ is the considered model.

Note that all those conditions are expressible by first order formulas (infinitely many in the case of $\bullet_0$ and $\bullet_{2,m}$).

We claim that T' is consistent. Let $T_0 \subseteq T'$ be finite. Then there is $m < \omega$ such that T_0 only contains formulas from T along with formulas $\bullet_{1,m'}$, $\bullet_{2,m'}$ and $\bullet_{3,m'}$ for $m' \leqslant m$. Furthermore, there is a finite set Γ of formulas such that the formulas from $\bullet_2$ appearing in T_0 say at most that $(\bar{a}_i : i < m)$ and $(\bar{b}_j : j < m)$ are Γ-indiscernible.

By Lemma 82, for $m_* < \omega$ is large enough, we can find a subsequences $(\bar{a}'_i : i < m)$ of $(\bar{a}^{m_*}_i : i < m)$ and a subsequence $(\bar{b}'_i : i < m)$ of $(\bar{b}^{m_*}_i : i < m_*)$ that are Γ-indiscernible. But then M_{m_*} where we interpret the constants so as to name the two sequences $(\bar{a}'_i : i < m)$ of $(\bar{a}^{m_*}_i : i < m)$ is a model of T_0. Hence T_0 is consistent. As T_0 was an arbitrary finite subset of T', we conclude by compactness that T' is consistent.

Let $\mathbf{M}$ be a model of T' and set $I = (\bar{a}_i : i < \omega)$ and $J = (\bar{b}_j : j < \omega)$ as interpreted in $\mathbf{M}$. This yields the structure $\mathbf{M}$ as described in Step 1.

Step 2. Apply Lemma 85 to get an elementary extension $\mathbf{N}$ of $\mathbf{M}$, an elementary substructure $\mathbf{M}'$ of $\mathbf{N}$, such that $(\bar{b}_j : j < \omega)$ is a Morley sequence over $\mathbf{M}'\bar{a}_{<\omega}$ and $(\bar{a}_i : j < \omega)$ is a Morley sequence over $\mathbf{M}'\bar{b}_{<\omega}$. In particular:

1. $(\bar{a}_i : i > \omega)$ and $(\bar{b}_j : j > \omega)$ are both indiscernible over $\mathbf{M}'$,
2. the families $\{\varphi(x,\bar{a}_i) \mid i < \omega\}$ and $\{\psi(x,\bar{b}_j) \mid j < \omega\}$ are both pairwise inconsistent,
3. for each $0 \leqslant i,j < \omega$ the set $\mathrm{Types}^{\Delta}(A/B)$ is infinite, where $A = \varphi(\mathbf{N},\bar{a}_i)$ and $B = \psi(\mathbf{N},\bar{b}_j)$,
4. $\bar{a}_i \mathop{\big|}_{\mathbf{M}'} \bar{a}_{<i}\bar{b}_{<\omega}$.

Let $A = \varphi(\mathbf{N},\bar{a}_1)$ and $B = \psi(\mathbf{N},\bar{b}_j)$. Since $\mathrm{Types}^{\Delta}(A/B)$ is infinite and Δ is finite, there is some $\theta(\bar{u},\bar{v}) \in \Delta$ such that $\mathrm{Types}^{\theta}(A/B)$ is infinite. This finishes the proof of Lemma 64. ◀

References

1 Stefan Arnborg, Jens Lagergren, and Detlef Seese. Problems easy for tree-decomposable graphs (extended abstract). In Timo Lepistö and Arto Salomaa, editors, *Automata, Languages and Programming*, pages 38–51, Berlin, Heidelberg, 1988. Springer Berlin Heidelberg.

2 John T. Baldwin and Saharon Shelah. Second-order quantifiers and the complexity of theories. *Notre Dame Journal of Formal Logic*, 26(3):229–303, 1985.

3 József Balogh, Béla Bollobás, and Robert Morris. Hereditary properties of ordered graphs. In *Topics in discrete mathematics*, pages 179–213. Springer, 2006.

4 József Balogh, Béla Bollobás, and Robert Morris. Hereditary properties of partitions, ordered graphs and ordered hypergraphs. *Eur. J. Comb.*, 27(8):1263–1281, 2006. doi:10.1016/j.ejc.2006.05.004.

5 Manuel Bodirsky. *Ramsey classes: examples and constructions*, pages 1–48. London Mathematical Society Lecture Note Series. Cambridge University Press, 2015. doi:10.1017/CBO9781316106853.002.

6 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: Max Independent Set and Coloring. *CoRR*, abs/2007.14161, 2020. URL: https://arxiv.org/abs/2007.14161, arXiv:2007.14161.

7 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1977–1996, 2021. doi:10.1137/1.9781611976465.118.

8 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 601–612. IEEE, 2020. doi:10.1109/FOCS46700.2020.00062.

9 Édouard Bonnet, Jaroslav Nesetril, Patrice Ossona de Mendez, Sebastian Siebertz, and Stéphan Thomassé. Twin-width and permutations. *CoRR*, abs/2102.06880, 2021. URL: https://arxiv.org/abs/2102.06880, arXiv:2102.06880.

10 Simone Bova and Barnaby Martin. First-order queries on finite abelian groups. In *24th EACSL Annual Conference on Computer Science Logic (CSL 2015)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.

11 Samuel Braunfeld and Michael C. Laskowski. Characterizations of monadic NIP, 2021. arXiv:2104.12989.

12 Yijia Chen and Jörg Flum. On the ordered conjecture. In *27th Annual IEEE Symposium on Logic in Computer Science*, pages 225–234. IEEE, 2012.

13 Josef Cibulka and Jan Kyncl. Füredi-Hajnal limits are typically subexponential. *CoRR*, abs/1607.07491, 2016. URL: `http://arxiv.org/abs/1607.07491`, `arXiv:1607.07491`.

14 Bruno Courcelle. Monadic second-order definable graph transductions: a survey. *Theoretical Computer Science*, 126(1):53–75, 1994. URL: `https://www.sciencedirect.com/science/article/pii/0304397594902682`, `doi:https://doi.org/10.1016/0304-3975(94)90268-2`.

15 Bruno Courcelle and Sang-il Oum. Vertex-minors, monadic second-order logic, and a conjecture by seese. *Journal of Combinatorial Theory, Series B*, 97(1):91–126, 2007. URL: `https://www.sciencedirect.com/science/article/pii/S0095895606000463`, `doi:https://doi.org/10.1016/j.jctb.2006.04.003`.

16 Guoli Ding, Bogdan Oporowski, James G. Oxley, and Dirk Vertigan. Unavoidable minors of large 3-connected binary matroids. *J. Comb. Theory, Ser. B*, 66(2):334–360, 1996. `doi:10.1006/jctb.1996.0026`.

17 Rodney G. Downey, Michael R. Fellows, and Udayan Taylor. The parameterized complexity of relational database queries and an improved characterization of W[1]. In Douglas S. Bridges, Cristian S. Calude, Jeremy Gibbons, Steve Reeves, and Ian H. Witten, editors, *First Conference of the Centre for Discrete Mathematics and Theoretical Computer Science, DMTCS 1996, Auckland, New Zealand, December, 9-13, 1996*, pages 194–213. Springer-Verlag, Singapore, 1996.

18 Kord Eickmeyer and Ken-ichi Kawarabayashi. FO model checking on map graphs. In *International Symposium on Fundamentals of Computation Theory*, pages 204–216. Springer, 2017.

19 J. Flum and M. Grohe. *Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series)*. Springer-Verlag, Berlin, Heidelberg, 2006.

20 Jakub Gajarský, Petr Hlinený, Daniel Lokshtanov, Jan Obdrzálek, and M. S. Ramanujan. A new perspective on FO model checking of dense graph classes. *CoRR*, abs/1805.01823, 2018. URL: `http://arxiv.org/abs/1805.01823`, `arXiv:1805.01823`.

21 Jakub Gajarský, Petr Hlinený, Jan Obdrzálek, Daniel Lokshtanov, and M. S. Ramanujan. A new perspective on FO model checking of dense graph classes. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 176–184, 2016. `doi:10.1145/2933575.2935314`.

22 Jakub Gajarský and Daniel Kráľ. Recovering sparse graphs. In *43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018.

23 Jakub Gajarský, Stephan Kreutzer, Jaroslav Nesetril, Patrice Ossona de Mendez, Michal Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. First-order interpretations of bounded expansion classes. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, pages 126:1–126:14, 2018. `doi:10.4230/LIPIcs.ICALP.2018.126`.

24 Sylvain Gravier, Frédéric Maffray, Jérôme Renault, and Nicolas Trotignon. Ramsey-type results on singletons, co-singletons and monotone sequences in large collections of sets. *Eur. J. Comb.*, 25(5):719–734, 2004. `doi:10.1016/j.ejc.2003.10.004`.

25 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017. `doi:10.1145/3051095`.

26 Benjamin Gunby and Dömötör Pálvölgyi. Asymptotics of pattern avoidance in the Klazar set partition and permutation-tuple settings. *Eur. J. Comb.*, 82, 2019. `doi:10.1016/j.ejc.2019.07.003`.

27 Petr Hliněný, Filip Pokrývka, and Bodhayan Roy. FO model checking on geometric graphs. *Computational Geometry*, 78:1–19, 2019.

28 Russell Impagliazzo and Ramamohan Paturi. On the Complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. `doi:10.1006/jcss.2000.1727`.

29 Martin Klazar. The Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture. In *Formal power series and algebraic combinatorics*, pages 250–255. Springer, 2000.

30 Martin Klazar. On growth rates of permutations, set partitions, ordered graphs and other objects. *Electron. J. Comb.*, 15(1), 2008. URL: `http://www.combinatorics.org/Volume_15/Abstracts/v15i1r75.html`.

31 Phokion G. Kolaitis and Moshe Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *Proceedings of the Seventh Annual IEEE Symposium on Logic in Computer Science*, pages 46–57. IEEE, 1992.

32 Stephan Kreutzer and Siamak Tazari. Lower bounds for the complexity of monadic second-order logic. In *Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11-14 July 2010, Edinburgh, United Kingdom*, pages 189–198. IEEE Computer Society, 2010. `doi:10.1109/LICS.2010.39`.

33 Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Comb. Theory, Ser. A*, 107(1):153–160, 2004. `doi:10.1016/j.jcta.2004.04.002`.

34 Jaroslav Nesetril and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012. doi:10.1007/978-3-642-27875-4.

35 Frank P. Ramsey. On a problem of formal logic. In *Proc. London Math. Soc. series 2*, volume 30 of *264-286*, 1930.

36 Neil Robertson and P.D Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986. URL: https://www.sciencedirect.com/science/article/pii/0095895686900304, doi:https://doi.org/10.1016/0095-8956(86)90030-4.

37 Detlef Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 6(6):505–526, 1996.

38 Saharon Shelah. *Monadic logic: Hanf Numbers. In: Around Classification Theory of Models*, volume vol 1182 of *Lecture Notes in Mathematics*. Springer, Berlin, Heidelberg, 1986.

39 Pierre Simon. A note on stability and NIP in one variable, 2021. arXiv:2103.15799.

40 Pierre Simon and Szymon Torunczyk. Ordered graphs of bounded twin-width. *CoRR*, abs/2102.06881, 2021. URL: https://arxiv.org/abs/2102.06881, arXiv:2102.06881.