Indiscernibles and Flatness in Monadically Stable and Monadically NIP Classes

Jan Dreier
TU Wien, Austria

Nikolas Mählmann
Universität Bremen, Germany

Sebastian Siebertz
Universität Bremen, Germany

Szymon Toruńczyk
University of Warsaw, Poland

Abstract

Monadically stable and monadically NIP classes of structures were initially studied in the context of model theory and defined in logical terms. They have recently attracted attention in the area of structural graph theory, as they generalize notions such as nowhere denseness, bounded cliquewidth, and bounded twinwidth.

Our main result is the – to the best of our knowledge first – purely combinatorial characterization of monadically stable classes of graphs, in terms of a property dubbed *flip-flatness*. A class $\mathscr{C}$ of graphs is flip-flat if for every fixed radius r, every sufficiently large set of vertices of a graph $G \in \mathscr{C}$ contains a large subset of vertices with mutual distance larger than r, where the distance is measured in some graph G' that can be obtained from G by performing a bounded number of flips that swap edges and non-edges within a subset of vertices. Flip-flatness generalizes the notion of uniform quasi-wideness, which characterizes nowhere dense classes and had a key impact on the combinatorial and algorithmic treatment of nowhere dense classes. To obtain this result, we develop tools that also apply to the more general monadically NIP classes, based on the notion of indiscernible sequences from model theory. We show that in monadically stable and monadically NIP classes indiscernible sequences impose a strong combinatorial structure on their definable neighborhoods. All our proofs are constructive and yield efficient algorithms.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Theory of computation → Finite Model Theory

Keywords and phrases stability, NIP, combinatorial characterization, first-order model checking

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.125

Category Track B: Automata, Logic, Semantics, and Theory of Programming

Related Version *Full Version*: `https://arxiv.org/abs/2206.13765` [15]

Funding *Nikolas Mählmann*: supported by the German Research Foundation (DFG) with grant greement No. 444419611.
Szymon Toruńczyk: supported by the project BOBR that is funded from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme with grant agreement No. 948057.

Acknowledgements We thank Édouard Bonnet, Jakub Gajarský, Stephan Kreutzer, Amer E. Mouawad and Alexandre Vigny for their valuable contributions to this paper. In particular, we thank Jakub Gajarský and Stephan Kreutzer for suggesting the notion of flip-flatness and providing a proof for the cases $r = 1, 2$.

1 Introduction

An important open problem in structural and algorithmic graph theory is to characterize those hereditary graph classes for which the model checking problem for first-order logic is tractable[1] [23, Section 8.2]. A result of Grohe, Kreutzer, and Siebertz [24] states that for monotone graph classes (that is, classes closed under removing vertices and edges), the limit of tractability is precisely captured by the notion of nowhere denseness, introduced by Nešetřil and Ossona de Mendez [29]. Examples of nowhere dense classes include the class of planar graphs, all classes that exclude a fixed minor, and classes with bounded expansion. Whereas these classes are sparse (for instance, they exclude some fixed biclique as a subgraph), the aforementioned problem seeks to generalize the result of Grohe, Kreutzer, and Siebertz to classes that are not necessarily sparse. Indeed, there are known hereditary graph classes that are not sparse, and for which the model checking problem is tractable, such as transductions of classes of bounded local cliquewidth [5], transductions of nowhere dense classes [16], or classes of ordered graphs (that is, graphs equipped with a total order) of bounded twinwidth [6].

So far, a complete picture is understood in two contexts: for monotone graph classes, where tractability coincides with nowhere denseness, and for hereditary classes of ordered graphs, where tractability coincides with bounded twinwidth. Despite the apparent dissimilarity of the combinatorial definitions of nowhere denseness and bounded twinwidth, those notions can be alternatively characterized in a uniform way in logical terms by the following notion, originating in model theory. Unless mentioned otherwise, all formulas are first-order formulas.

Say that a class $\mathscr{C}$ of graphs *transduces* a class $\mathscr{D}$ of graphs if for every $H \in \mathscr{D}$ there is some $G \in \mathscr{C}$ from which H can be obtained by performing the following steps: (1) coloring the vertices of G arbitrarily (2) interpreting a fixed formula $\varphi(x, y)$ (involving the edge relation and unary relations for the colors), thus yielding a new graph $\varphi(G)$ with the same vertices as G and edges uv such that $\varphi(u, v)$ holds, and finally (3) taking an induced subgraph of $\varphi(G)$. The transducability relation on graph classes is transitive, and classes that *do not* transduce the class of all graphs are called *monadically NIP*. For instance, the class of all bipartite graphs transduces the class of all graphs: to obtain an arbitrary graph G, consider its 1-subdivision, obtained by placing one vertex on each edge of G, thus yielding a bipartite graph H; then the formula $\varphi(x, y)$ expressing that x and y have a common neighbor defines a graph on $V(H)$ containing G as an induced subgraph. Hence, the class of bipartite graphs is not monadically NIP. On the other hand, all the graph classes mentioned earlier – nowhere dense classes and transductions thereof, classes of bounded twinwidth, or transductions of classes with bounded local cliquewidth – are monadically NIP. This suggests that monadic NIP might constitute the limit of tractability of the model checking problem. More precisely, the following has been conjectured[2].

▶ **Conjecture 1** ([1]). *Let $\mathscr{C}$ be a hereditary class of graphs. Then the model checking problem for first-order logic is fixed parameter tractable on $\mathscr{C}$ if and only if $\mathscr{C}$ is monadically NIP.*

Quite remarkably, among monotone graph classes, classes that are monadically NIP correspond precisely to nowhere dense classes [2], and among hereditary graph classes of ordered graphs, classes that are monadically NIP correspond precisely to classes of bounded

[1] more precisely, *fixed parameter-tractable*, that is, solvable in time $f(|\varphi|) \cdot |G|^c$, where φ is the input formula and G is the input graph, for some function $f : \mathbb{N} \to \mathbb{N}$ and constant c

[2] To the best of our knowledge the conjecture was first explicitly discussed during the open problem session of the Algorithms, Logic and Structure Workshop in Warwick, in 2016, see [1].

twinwidth [6]. One may tweak the definition of monadic NIP classes by considering other logics than first-order logic. For instance, for the counting extension CMSO_2 of monadic second-order logic, one recovers precisely the notion of classes of bounded cliquewidth [9], or classes of bounded treewidth if only monotone classes are considered. Thus, variations on the definition of monadic NIP recover important notions from graph theory: nowhere denseness, bounded twinwidth, bounded treewidth, and bounded cliquewidth.

Note that both implications in Conjecture 1 remain open. The conjecture is so far confirmed for monotone graph classes [24] (where monadically NIP classes are exactly the nowhere dense classes) and for hereditary classes of ordered graphs [6], tournaments [22], interval graphs and permutation graphs [7], (where monadically NIP classes are exactly the classes of bounded twinwidth). As a special important case, the conjecture predicts that all *monadically stable* graph classes are tractable. A class $\mathscr{C}$ is monadically stable if it does not transduce the class of all half-graphs, that is, graphs with vertices $a_1, b_1, \ldots, a_n, b_n$ such that a_i is adjacent to b_j if and only if $i \leqslant j$. Although much more restrictive than monadically NIP classes, monadically stable classes include all nowhere dense classes [2], and hence also all classes $\mathscr{D}$ that transduce in a nowhere dense class $\mathscr{C}$ (called structurally nowhere dense classes [19]). Those include dense graph classes, such as for instance squares of planar graphs. In fact, it is conjectured [30] that every monadically stable class of graphs is structurally nowhere dense.

These outlined connections between structural graph theory and model theory have recently triggered the interest to generalize combinatorial and algorithmic results from nowhere dense classes to structurally nowhere dense, monadically stable and monadically NIP classes, and ultimately to efficiently solve the model checking problem for first-order logic on these classes [5, 6, 19, 13, 21, 27, 30, 31, 32]. Logical results on monadically stable and monadically NIP classes in model theory include [3, 34, 8, 4].

Contribution

As discussed above, many central graph classes such as those with bounded cliquewidth, twinwidth or nowhere dense classes can be characterized both from a structural (i.e., graph theoretic) and a logical perspective. While monadically stable and monadically NIP graph classes are naturally defined via logic, structural characterizations have so far been elusive. In this work we take a step towards a structure theory for monadically stable and monadically NIP classes of graphs, which is the basis for their future algorithmic and combinatorial treatment, in particular, a tool for approaching Conjecture 1, as we discuss later.

Flatness. Our main result, Theorem 3, provides a purely combinatorial characterization of monadically stable graph classes in terms of *flip-flatness*. Flip-flatness generalizes *uniform quasi-wideness*[3], introduced by Dawar in [10] in his study of homomorphism preservation properties. Uniform quasi-wideness was proved by Nešetřil and Ossona de Mendez [29] to characterize nowhere dense graph classes and is a key tool in the combinatorial and algorithmic treatment of these classes. To foster the further discussion, let us formally define this notion.

[3] Another reasonable name for flip-flatness would be *flip-wideness*. We avoided this name to prevent confusion with the recently introduced graph parameter *flip-width* [37], which is studied in the same context.

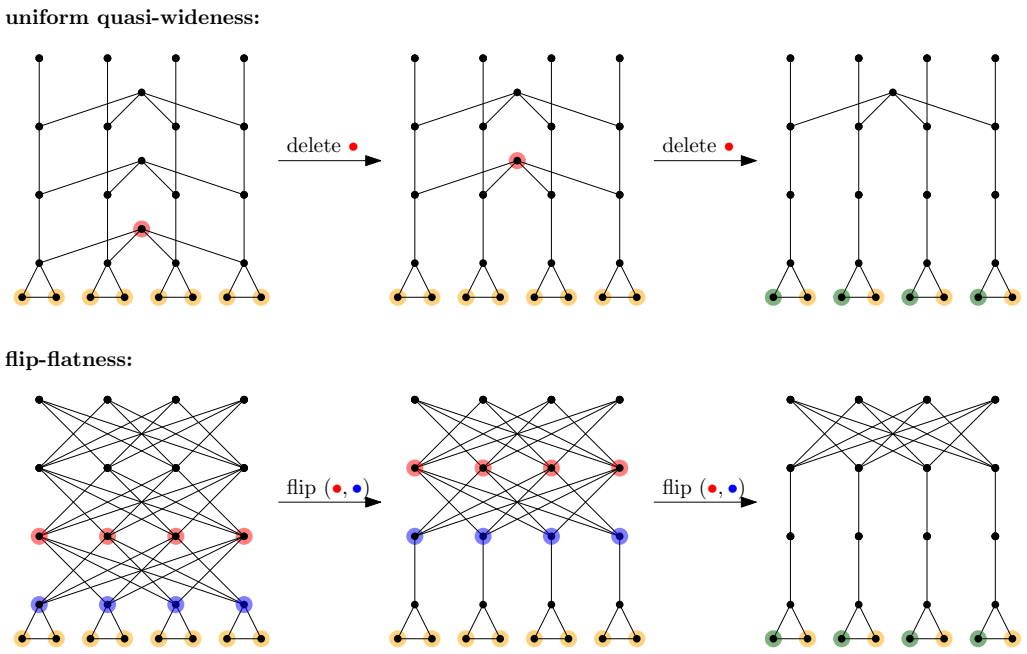

Figure 1 An example of uniform quasi-wideness (flip-flatness). Among the yellow vertices, we find a still large set of green vertices, that is distance-7 independent after deleting a bounded number of red vertices (performing a bounded number of flips between sets of red and blue vertices). Two key properties are illustrated: 1. the higher the desired independence distance, the more operations have to be performed; 2. we cannot hope to make all the yellow vertices distance-r independent with a bounded number of operations.

A set of vertices is *distance-r independent* in a graph if any two vertices in the set are at distance greater than r. A class of graphs $\mathscr{C}$ is *uniformly quasi-wide* if for every $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $s \in \mathbb{N}$ with the following property. For all $m \in \mathbb{N}$, $G \in \mathscr{C}$, and $A \subseteq V(G)$ with $|A| \geqslant N_r(m)$, there exists $S \subseteq V(G)$ with $|S| \leqslant s$ and $B \subseteq A \setminus S$ with $|B| \geqslant m$ such that B is distance-r independent in $G - S$ (see the top of Figure 1). Intuitively, for uniformly quasi-wide classes, in sufficiently large sets one can find large subsets that are distance-r independent after the deletion of a constant number of vertices. Uniform quasi-wideness is only suitable for the treatment of sparse graphs. Already very simple dense graph classes, such as the class of all cliques, are not uniformly quasi-wide.

Inspired by the notion of uniform quasi-wideness, Jakub Gajarský and Stephan Kreutzer proposed the new notion of *flip-flatness*. Roughly speaking, flip-flatness generalizes uniform quasi-wideness by replacing in its definition the deletion of a bounded size set of vertices by the inversion of the edge relation between a bounded number of (arbitrarily large) vertex sets. Let us make this definition more precise. A *flip* in a graph G is specified by a pair of sets $\mathsf{F} = (A, B)$ with $A, B \subseteq V(G)$. We write $G \oplus \mathsf{F}$ for the graph with the same vertices as G, and edges uv such that $uv \in E(G)$ xor $(u, v) \in (A \times B) \cup (B \times A)$. For a set $F = \{\mathsf{F}_1, \ldots, \mathsf{F}_n\}$ of flips, we write $G \oplus F$ for the graph $G \oplus \mathsf{F}_1 \oplus \cdots \oplus \mathsf{F}_n$.

▶ **Definition 2.** *A class of graphs $\mathscr{C}$ is* flip-flat *if for every $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $s_r \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathscr{C}$, and $A \subseteq V(G)$ with $|A| \geqslant N_r(m)$, there exists a set F of at most s_r flips and $B \subseteq A$ with $|B| \geqslant m$ such that B is distance-r independent in $G \oplus F$.*

Intuitively, for flip-flat classes in sufficiently large sets one can find large subsets that are distance-r independent after a constant number of flips (see bottom of Figure 1). For example the class of all cliques is flip-flat, requiring only a single flip to make the whole vertex set distance-∞ independent. Our main result is the following purely combinatorial characterization of monadic stability.

▶ **Theorem 3.** *A class of graphs is monadically stable if and only if it is flip-flat.*

Notably, our proof is algorithmic and yields polynomial bounds in the following sense.

▶ **Theorem 4.** *Every monadically stable class $\mathscr{C}$ of graphs is flip-flat, where for every r, the function N_r is polynomial. Moreover, given an n-vertex graph $G \in \mathscr{C}$, $A \subseteq V(G)$, and $r \in \mathbb{N}$, we can compute a subset $B \subseteq A$ and a set of flips F that makes B distance-r independent in $G \oplus F$ in time $\mathcal{O}(f_{\mathscr{C}}(r) \cdot n^3)$ for some function $f_{\mathscr{C}}$.*

Just as uniform quasi-wideness provided a key tool for the algorithmic treatment of nowhere dense graph classes, in particular for first-order model checking [24], we believe that the characterization by flip-flatness will provide an important step towards the algorithmic treatment of monadically stable classes. We leave as an open question whether a similar characterization of monadically NIP classes exists.

Indiscernible sequences. Our study is based on *indiscernible sequences*, which are a fundamental tool in model theory. An indiscernible sequence is a (finite or infinite) sequence $I = (\bar{a}_1, \bar{a}_2, \ldots)$ of tuples of equal length of elements of a fixed (finite or infinite) structure, such that any two finite subsequences of I that have equal length, satisfy the same formulas. More generally, for a set of formulas Δ, the sequence I is Δ-*indiscernible* if for each formula $\varphi(\bar{x}_1, \ldots, \bar{x}_k)$ from Δ either all subsequences of I of length k satisfy φ, or no subsequence of length k satisfies φ. For example, any enumeration of vertices forming a clique or an independent set in a graph is Δ-indiscernible for $\Delta = \{E(x_1, x_2)\}$ containing only the edge relation. Also, any increasing sequence of rationals in the structure $(\mathbb{Q}, <)$ is Δ-indiscernible for Δ being the set of all formulas. We use indiscernible sequences to obtain new insights about monadically stable and monadically NIP classes.

It was shown by Blumensath [4] that in monadically NIP classes, any fixed element interacts with the tuples of an indiscernible sequence in a very regular way. We give a new finitary proof of Blumensath's result. Building on this, we develop our main technical tool, Theorem 11, where we show that the regular properties of the indiscernible sequences extend to their disjoint definable neighborhoods (see Section 3.2 for details). A result similar to Theorem 11, also for disjoint definable neighborhoods in monadically NIP classes, plays a key role in [6, 35].

Apart from powering our algorithmic proof of flip-flatness, Theorem 11 has already found further applications in monadically stable classes of graphs: it was recently used to obtain improved bounds for Ramsey numbers [14] and to prove an algorithmic game-characterization of these classes, called *Flipper game* [20]. The paper [20] provides two proofs for this game characterization. One is constructive and algorithmic and builds on our work. The other builds on model theoretic tools; it is non-constructive but self-contained and highlights additional properties of monadically stable graph classes, including a second (though non-constructive) proof of Theorem 3. The algorithmic version of the Flipper game plays a crucial role for proving that the first-order model checking problem is fixed parameter tractable for structurally nowhere dense classes of graphs [16]. This suggests that our developed techniques may be an important tool towards resolving Conjecture 1.

2 Preliminaries

We use standard notation from graph theory and model theory and refer to [11] and [25] for extensive background. We write $[m]$ for the set of integers $\{1, \dots, m\}$.

Relational structures and graphs. A *(relational) signature* Σ is a set of relation symbols, each with an associated non-negative integer, called its *arity*. A Σ-*structure* $\mathbb{A}$ consists of a *universe*, which is a non-empty, possibly infinite set, and *interpretations* of the symbols from the signature: each relation symbol of arity k is interpreted as a k-ary relation over the universe. By a slight abuse of notation, we do not differentiate between structures and their universes and between relation symbols and their interpretations. By $\mathscr{C}$ we denote classes of structures over a fixed signature. Unless indicated otherwise, $\mathscr{C}$ may contain finite and infinite structures.

A *monadic extension* Σ^+ of a signature Σ is any extension of Σ by unary predicates. A unary predicate will also be called a *color*. A *monadic expansion* or *coloring* of a Σ-structure $\mathbb{A}$ is a Σ^+-structure $\mathbb{A}^+$, where Σ^+ is a monadic extension of Σ, such that $\mathbb{A}$ is the Σ-*reduct* of $\mathbb{A}^+$, that is, the Σ-structure obtained from $\mathbb{A}^+$ by removing all relations with symbols in $\Sigma^+ \setminus \Sigma$. When $\mathscr{C}$ is a class of Σ-structures and Σ^+ is a monadic extension of Σ, we write $\mathscr{C}[\Sigma^+]$ for the class of all possible Σ^+-expansions of structures from $\mathscr{C}$.

A *graph* is a finite structure over the signature consisting of a binary relation symbol E, interpreted as the symmetric and irreflexive edge relation.

First-order logic. We say that two tuples $\bar{a}, \bar{b}$ of elements are φ-*connected* in a structure $\mathbb{A}$ if $\mathbb{A} \models \varphi(\bar{a}, \bar{b})$. We call the set $N_\varphi^{\mathbb{A}}(\bar{a}) = \{\bar{b} \in \mathbb{A}^{|\bar{y}|} \ : \ \mathbb{A} \models \varphi(\bar{a}, \bar{b})\}$ the φ-*neighborhood* of $\bar{a}$. We simply write $N_\varphi(\bar{a})$ when $\mathbb{A}$ is clear from the context.

Let $\Phi(\bar{x})$ be a finite set of formulas with free variables $\bar{x}$. A Φ-*type* is a conjunction $\tau(\bar{x})$ of formulas in Φ or their negations, such that every formula in Φ occurs in τ either positively or negatively. More precisely, $\tau(\bar{x})$ is a formula of the following form, for some subset $A \subseteq \Phi$:

$$\bigwedge_{\varphi \in A} \varphi(\bar{x}) \wedge \bigwedge_{\varphi \in \Phi \setminus A} \neg\varphi(\bar{x}).$$

Note that for every $|\bar{x}|$-tuple $\bar{a}$ of elements of $\mathbb{A}$, there is exactly one Φ-type $\tau(\bar{x})$ such that $\mathbb{A} \models \tau(\bar{a})$.

Stability and NIP. While we already defined monadic stability and NIP in terms of transductions in the introduction, let us also give the equivalent original definition. A formula $\varphi(\bar{x}, \bar{y})$ has the k-*order property* on a class $\mathscr{C}$ of structures if there are $\mathbb{A} \in \mathscr{C}$ and two sequences $(\bar{a}_i)_{1 \leqslant i \leqslant k}$, $(\bar{b}_i)_{1 \leqslant i \leqslant k}$ of tuples of elements of $\mathbb{A}$, such that for all $i, j \in [k]$

$$\mathbb{A} \models \varphi(\bar{a}_i, \bar{b}_j) \quad \iff \quad i \leqslant j.$$

The formula φ is said to have the *order property* on $\mathscr{C}$ if it has the k-order property for all $k \in \mathbb{N}$. The class $\mathscr{C}$ is called *stable* if no formula has the order-property on $\mathscr{C}$. A class $\mathscr{C}$ of Σ-structures is *monadically stable* if for every monadic extension Σ^+ of Σ the class $\mathscr{C}[\Sigma^+]$ is stable.

Similarly, a formula $\varphi(\bar{x}, \bar{y})$ has the k-*independence property* on a class $\mathscr{C}$ if there are $\mathbb{A} \in \mathscr{C}$, a size k set $A \subseteq \mathbb{A}^{|\bar{x}|}$ and a sequence $(\bar{b}_J)_{J \subseteq A}$ of tuples of elements of $\mathbb{A}$ such that for all $J \subseteq A$ and for all $\bar{a} \in A$

$$\mathbb{A} \models \varphi(\bar{a}, \bar{b}_J) \quad \iff \quad \bar{a} \in J.$$

We then say that A is *shattered* by φ. We define the *independence property* (IP), classes with the *non-independence property* (NIP classes), and *monadically NIP* classes as expected. Note that every (monadically) stable class is (monadically) NIP. Baldwin and Shelah proved that in the definitions of monadic stability and monadic NIP, one can alternatively rely on formulas $\varphi(x,y)$ with just a pair of singleton variables, instead of a pair of tuples of variables [3, Lemma 8.1.3, Theorem 8.1.8].

We call a relation R *definable* on a structure $\mathbb{A}$ if $R = \{\bar{a} \in \mathbb{A}^{|\bar{x}|} : \mathbb{A} \models \varphi(\bar{a})\}$ for some formula $\varphi(\bar{x})$. The following is immediate from the previous definitions.

▶ **Lemma 5.** *Let $\mathscr{C}$ be a monadically stable (monadically NIP) class of Σ-structures, let Σ^+ be a monadic extension of Σ and let $\mathscr{D}$ be the hereditary closure of any expansion of $\mathscr{C}[\Sigma^+]$ by definable relations. Then also $\mathscr{D}$ is monadically stable (monadically NIP).*

A formula $\varphi(x,\bar{y})$ has *pairing index* k on a class $\mathscr{C}$ of structures if there is $\mathbb{A} \in \mathscr{C}$ and two sequences $(a_{ij})_{1\leqslant i<j\leqslant k}$ and $(\bar{b}_i)_{1\leqslant i\leqslant k}$ such that for all $1 \leqslant i < j \leqslant k$ and $\ell \in [k]$

$$\mathbb{A} \models \varphi(a_{ij}, \bar{b}_\ell) \quad \iff \quad \ell \in \{i,j\}.$$

We require that x is a single free variable, while $\bar{y}$ is allowed to be a tuple of variables. Intuitively, from a graph theoretic perspective, if a formula has unbounded pairing index on a class $\mathscr{C}$, then it can encode arbitrarily large 1-subdivided cliques in $\mathscr{C}$, where the principal vertices are represented by the tuples $\bar{b}_\ell$ and the subdivision vertices are represented by single elements a_{ij}. As discussed in the introduction, this is sufficient to encode arbitrary graphs in $\mathscr{C}$ by using an additional color predicate. In this case, $\mathscr{C}$ cannot be monadically NIP. The above reasoning is formalized in the following characterization.

▶ **Lemma 6.** *A class $\mathscr{C}$ of Σ-structures is monadically NIP if and only if for every monadic extension Σ^+ of Σ every Σ^+-formula $\varphi(x,\bar{y})$ has bounded pairing index on $\mathscr{C}[\Sigma^+]$.*

Indiscernible sequences. Let $\mathbb{A}$ be a Σ-structure and $\varphi(\bar{x}_1,\dots,\bar{x}_m)$ a formula. We say a sequence $(\bar{a}_i)_{1\leqslant i\leqslant n}$ of tuples from $\mathbb{A}$ (where all $\bar{x}_i$ and $\bar{a}_j$ have the same length) is a *φ-indiscernible sequence* of length n, if for every two sequences of indices $i_1 < \cdots < i_m$ and $j_1 < \cdots < j_m$ from $[n]$ we have

$$\mathbb{A} \models \varphi(\bar{a}_{i_1},\dots,\bar{a}_{i_m}) \quad \iff \quad \mathbb{A} \models \varphi(\bar{a}_{j_1},\dots,\bar{a}_{j_m}).$$

For a set of formulas Δ we call a sequence *Δ-indiscernible* if it is φ-indiscernible for all $\varphi \in \Delta$. For finite Δ, by (iteratively applying) Ramsey's theorem we can extract a Δ-indiscernible sequence from any sequence. In general structures however, in order to extract a large Δ-indiscernible sequence we must initially start with an enormously large sequence. To the best of our knowledge the following theorem goes back to Ehrenfeucht and Mostowski [17].

▶ **Theorem 7.** *Let Δ be a finite set of formulas. There exists a function f such that every sequence of elements of length at least $f(m)$ (in a structure with a signature matching Δ) contains a Δ-indiscernible subsequence of length m.*

In stable classes we can efficiently find polynomially large (that is, $f(m) = m^{O(1)}$) indiscernible sequences of elements (see also [28, Theorem 3.5]). In the full version of the paper we observe that the run time is essentially bounded by the time it takes to evaluate the formulas from Δ on a small polynomial part of the input structure.

3 Technical overview

Our work is organized as follows. Section 3.1 introduces useful combinatorial properties of indiscernibles in monadically NIP classes. We extend those properties in Section 3.2, where we prove our main tool, Theorem 11. We use this tool in Section 3.3 to prove flip-flatness. Due to space constraints, we mostly provide proof sketches, which should convey the key ideas. All missing proofs can be found in the full version of the paper [15].

3.1 Indiscernibles in monadically stable and monadically NIP classes

In classical model theory, instead of classes of finite structures, usually infinite structures are studied. For a single infinite structure $\mathbb{A}$, we say $\mathbb{A}$ is (monadically) stable/NIP, if the class $\{\mathbb{A}\}$ is (monadically) stable/NIP. In this context, an indiscernible sequence is usually assumed to be of infinite length and indiscernible over the set of all first-order formulas. Using Ω to denote the set of all first-order formulas, we will denote the latter property as Ω-*indiscernibility*.

It is well-known that indiscernible sequences can be used to characterize stability and NIP for infinite structures. To characterize NIP, define the *alternation rank* of a formula $\varphi(\bar{x}, \bar{y})$ over a sequence I in a structure $\mathbb{A}$ as the maximum k such that there exists a (possibly non-contiguous) subsequence $(\bar{a}_1, \ldots, \bar{a}_{k+1})$ of I and a tuple $\bar{b} \in \mathbb{A}^{|\bar{x}|}$ such that for all $i \in [k]$ we have $\mathbb{A} \models \varphi(\bar{b}, \bar{a}_i) \iff \mathbb{A} \models \neg\varphi(\bar{b}, \bar{a}_{i+1})$. A structure $\mathbb{A}$ is NIP if and only if every formula has finite alternation rank over every Ω-indiscernible sequence in $\mathbb{A}$ [33, Theorem 12.17]. Stable classes can be characterized in a similar way. We say the *exception rank* of a formula $\varphi(\bar{x}, \bar{y})$ over a sequence I in a structure $\mathbb{A}$ is the maximum k such that there exists a tuple $\bar{b} \in \mathbb{A}^{|\bar{x}|}$ such that for k distinct tuples $\bar{a}_i \in I$ we have $\mathbb{A} \models \varphi(\bar{b}, \bar{a}_i)$ and for k other distinct tuples $\bar{a}_i \in I$ we have $\mathbb{A} \models \neg\varphi(\bar{b}, \bar{a}_i)$. A structure $\mathbb{A}$ is stable if and only if every formula has finite exception rank over every Ω-indiscernible sequence in $\mathbb{A}$ [33, Corollary 12.24].

Hence, NIP and stablility can be characterized by the interaction of tuples of elements with indiscernible sequences. For *monadically* NIP structures, Blumensath [4] shows that the interaction of single elements with indiscernible sequences is even more restricted.

▶ **Theorem 8** ([4, Corollary 4.13]). *In every monadically NIP structure $\mathbb{A}$, for every formula $\varphi(x, \bar{y})$, where x is a single free variable, and Ω-indiscernible sequence $(\bar{a}_i)_{i\in\mathbb{N}}$ of $|\bar{y}|$-tuples the following holds: for every element $b \in \mathbb{A}$ there exists an exceptional index $\mathrm{ex}(b) \in \mathbb{N}$ and two truth values $t_<(b), t_>(b) \in \{0, 1\}$ such that*

$$\begin{aligned} &\text{for all } i < \mathrm{ex(b)} : \mathbb{A} \models \varphi(b, \bar{a}_i) \iff t_<(b) = 1, \text{ and} \\ &\text{for all } i > \mathrm{ex(b)} : \mathbb{A} \models \varphi(b, \bar{a}_i) \iff t_>(b) = 1. \end{aligned}$$

See Figure 2 for examples. In particular, the alternation rank of the formulas $\varphi(x, \bar{y})$ with a single free variable x over every Ω-indiscernible sequence in a monadically NIP structure is at most 2.

Theorem 8 will be a crucial tool for proving flip-flatness. However, as we strive for algorithmic applications, we have to develop an effective, computable version that is suitable for handling classes of finite structures. Instead of requiring Ω-indiscernibility, we will specifically consider Δ-indiscernible sequences with respect to special sets $\Delta = \Delta_k^\Phi$ that we define soon. These sets Δ_k^Φ will strike the right balance of being on the one hand sufficiently rich to allow us to derive structure from them, but on the other hand sufficiently simple such that we can efficiently evaluate formulas from Δ_k^Φ, making our flip-flatness result algorithmic.

Figure 2 Two monadically NIP structures. On the left: the infinite half-graph, where Theorem 8 applies for the edge relation with $t_<(b) = 0$ and $t_>(b) = 1$. On the right: the infinite matching, where we have $t_<(b) = t_>(b) = 0$ but a differing truth value at index $\mathrm{ex}(b) = 3$.

Fix a finite set $\Phi(x, \bar{y})$ of formulas of the form $\varphi(x, \bar{y})$ where x is a single variable. A Φ-*pattern* is a finite sequence $(\varphi_i)_{1 \leqslant i \leqslant k'}$, where each formula $\varphi_i(x, \bar{y})$ is a boolean combination of formulas $\varphi(x, \bar{y}) \in \Phi(x, \bar{y})$. Given a Φ-pattern $(\varphi_i)_{i \in [k']}$, the following formula expresses that, for a given sequence of k' tuples, each of length $|\bar{y}|$, there is some element that realizes that pattern (see Figure 3 for an example):

$$\gamma_{(\varphi_1,\dots,\varphi_{k'})}(\bar{y}_1, \dots, \bar{y}_{k'}) = \exists x. \bigwedge_{i \in [k']} \varphi_i(x, \bar{y}_i).$$

For a finite set of formulas $\Phi(x, \bar{y})$ and an integer k we define Δ_k^Φ to be the set of all formulas $\gamma_{(\varphi_1,\dots,\varphi_{k'})}$, where $(\varphi_1, \dots, \varphi_{k'})$ is a Φ-pattern of length $k' \leqslant k$. Note that the set Δ_k^Φ is finite, as (up to equivalence) there are only finitely many boolean combinations of formulas in Φ. We write Δ_k^φ for $\Delta_k^{\{\varphi\}}$.

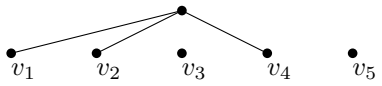

Figure 3 Example of an $\{E\}$-pattern: a graph satisfying $G \models \gamma_{(E,E,\neg E,E,\neg E)}(v_1, v_2, v_3, v_4, v_5)$.

We can now state a finitary version of Theorem 8 as follows. For the convenient use later on, we state the result not for single formulas φ but for Φ-types.

▶ **Theorem 9.** *For every monadically NIP class $\mathscr{C}$ of structures and finite set of formulas $\Phi(x, \bar{y})$ there exist integers n_0 and k, such that for every $\mathbb{A} \in \mathscr{C}$ the following holds. If $I = (\bar{a}_i)_{1 \leqslant i \leqslant n}$ is a Δ_k^Φ-indiscernible sequence of length $n \geqslant n_0$ in $\mathbb{A}$, and $b \in \mathbb{A}$, then there exists an exceptional index $\mathrm{ex}(b) \in [n]$ and Φ-types $\tau_-(x, \bar{y})$ and $\tau_+(x, \bar{y})$ such that*

$\mathbb{A} \models \tau_-(b, \bar{a}_i)$ *holds for all* $1 \leqslant i < \mathrm{ex}(b)$, *and*
$\mathbb{A} \models \tau_+(b, \bar{a}_i)$ *holds for all* $\mathrm{ex}(b) < i \leqslant n$.

Our proof uses different tools than [4]. It is combinatorial, fully constructive, and gives explicit bounds on n_0 and k as well as the required set of formulas Δ_k^Φ. One may alternatively finitize Theorem 8 via compactness, but then we do not obtain these crucial properties.

For the more restricted monadically stable classes we can give even stronger guarantees: for every element $b \in \mathbb{A}$, the types do not alternate and we have $\tau_- = \tau_+$.

▶ **Theorem 10.** *For every monadically stable class $\mathscr{C}$ of structures and finite set of formulas $\Phi(x, \bar{y})$ there exist integers n_0 and k, such that for every $\mathbb{A} \in \mathscr{C}$ the following holds. If $I = (\bar{a}_i)_{1 \leqslant i \leqslant n}$ is a Δ_k^Φ-indiscernible sequence of length $n \geqslant n_0$ in $\mathbb{A}$, and $b \in \mathbb{A}$, then there exists an exceptional index $\mathrm{ex}(b) \in [n]$ and a Φ-type τ, such that*

$\mathbb{A} \models \tau(b, \bar{a}_i)$ *for all* $i \in [n]$ *with* $i \neq \mathrm{ex}(b)$.

Building on Theorem 9, we give a full proof of Theorem 10. In order to give an intuition about the interaction of monadic stability and Δ_k^Φ-indiscernibility, we sketch a standalone proof here.

Proof sketch of Theorem 10. For simplicity, we will focus on the case were Φ is a singleton set $\{\varphi(x,\bar{y})\}$. Let k be the smallest number such that the bound for the pairing index of φ and $\neg\varphi$ on $\mathscr{C}$ is less than k. Such a bound exists since $\mathscr{C}$ is in particular monadically NIP. Assume towards a contradiction that there exists a sufficiently long Δ_k^φ-indiscernible sequence I in a structure $\mathbb{A} \in \mathscr{C}$ and an element $b \in \mathbb{A}$ that is φ-connected to at least two tuples in I and not φ-connected to at least two other tuples in I. By symmetry, we can assume that b is not φ-connected to the majority of I. We therefore find a length-k subsequence $\bar{a}_1, \dots, \bar{a}_k$ of I and two distinct indices $i_0, j_0 \in [k]$, such that b is φ-connected exactly to the i_0th and j_0th element of $\bar{a}_1, \dots, \bar{a}_k$. We say that the i_0th and j_0th element are *φ-paired* by b. Our goal is to derive a contradiction by finding also for every other pair i, j of indices an element that φ-pairs them, witnessing that φ has pairing index at least k in $\mathscr{C}$.

In stable classes, every sufficiently long Δ_k^φ-indiscernible sequence I is also *totally* Δ_k^φ-*indiscernible*: This means that every permutation I' of I is again Δ_k^φ-indiscernible, with the formulas from Δ_k^φ taking the same truth values on I' as on I. Intuitively speaking, if permuting two elements in an indiscernible sequence would change the truth value of a formula, then this formula orders these two elements and by indiscernibility also orders a large part of the sequence, contradicting stability. A proof of the statement for Ω-indiscernibles can be found in [36, Lemma 9.1.1].

The formula $\gamma(\bar{y}_1, \dots, \bar{y}_k) :=$ "$\bar{y}_{i_0}$ and $\bar{y}_{j_0}$ are φ-paired by some element" is contained in Δ_k^φ and holds on $\bar{a}_1, \dots, \bar{a}_k$, as witnessed by b. By total indiscernibility, we may permute $\bar{a}_1, \dots, \bar{a}_k$ and the formula still holds. By swapping a_{i_0} with a_i and a_{j_0} with a_j, we obtain for arbitrary i, j that "$\bar{a}_i$ and $\bar{a}_j$ are φ-paired by some element". This witnesses that φ has pairing index at least k, a contradiction. ◀

For the more general monadically NIP case, we cannot rely on total indiscernibility. Instead, we bound the alternation rank by pairing tuples located around alternation points, leading to a similar but more involved reasoning.

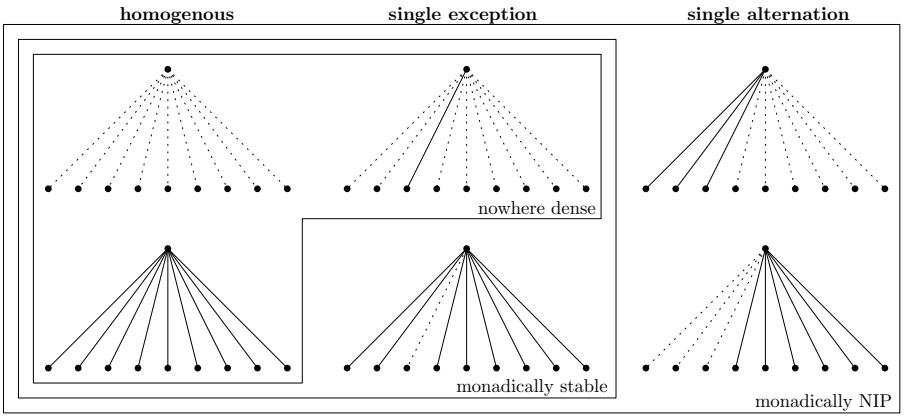

Figure 4 All types of neighborhoods a vertex (top) can have in indiscernible sequences (bottom) in different graph classes.

We conclude this section by mentioning that a behavior similar to Theorem 10 was already observed in [26] for the edge relation in nowhere dense classes. Considering only the edge relation, we depict the different possible behavior of vertices towards indiscernible sequences in different classes of graphs in Figure 4. In monadically NIP (or stable) classes, the same six (or four) patterns apply for connections with respect to any formula $\varphi(x, y)$.

3.2 Disjoint definable neighborhoods

In the previous section we have seen that in monadically stable and monadically NIP classes, every element is very homogeneously connected to all but at most one element of an indiscernible sequence. At a first glance, this exceptional behavior towards one element seems to be erratic and standing in the way of combinatorial and algorithmic applications. However, it turns out that it can be exploited to obtain additional structural properties.

The key observation is that elements that are "exceptionally connected" towards a single element of an indiscernible sequence, inherit some of the good properties of that sequence. We will give a simple example to demonstrate this idea. Let G be a graph containing certain red vertices R and blue vertices B, such that the edges between R and B describe a matching (see Figure 5, left).

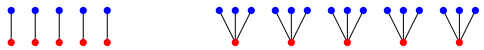

Figure 5 Examples of one-to-one and many-to-one connections in a graph. On the left: a matching. On the right: a star forest.

Given a formula $\varphi(x_1, \dots, x_q)$, define

$$\hat{\varphi}(x_1, \dots, x_q) := \exists b_1, \dots, b_q \in B. \varphi(b_1, \dots, b_q) \wedge \bigwedge_{i \in [q]} E(x_i, b_i).$$

Take a $\hat{\varphi}$-indiscernible sequence R' among R. Obviously, every vertex in the blue neighborhood B' of R' has one exceptional connection towards R', that is, towards its unique matching neighbor in R'. It is now easy to see that B' is a φ-indiscernible sequence in G: for every sequence of red vertices $a_1, \dots, a_q \in R'$ and their unique blue matching neighbors $b_1, \dots, b_q \in B'$ we have $G \models \hat{\varphi}(a_1, \dots, a_q)$ if and only if $G \models \varphi(b_1, \dots, b_q)$.

This example sketches how first-order definable one-to-one connections towards elements of an indiscernible sequence preserve indiscernibility. The more general case however is the many-to-one case, where we have a set of elements, each of which is exceptionally connected to a single element of an indiscernible sequence I, while allowing multiple elements to be exceptionally connected to the same element of I (think, for example, of R and B being the centers and leaves of a star-forest as depicted in Figure 5, right). Our notion of such many-to-one connections will be that of *disjoint α-neighborhoods*. Recall, that for a formula $\alpha(\bar{x}, y)$, the α-neighborhood $N_\alpha(\bar{a})$ of a tuple $\bar{a}$ is defined as the set of all b satisfying $\alpha(\bar{a}, b)$. We say a sequence J of $|\bar{x}|$-tuples has *disjoint α-neighborhoods* if $N_\alpha(\bar{a}_1) \cap N_\alpha(\bar{a}_2) = \varnothing$ for all distinct $\bar{a}_1, \bar{a}_2 \in J$.

As the main technical tool of this paper, we prove for monadically NIP classes that every sequence of disjoint α-neighborhoods contains a large subsequence that exerts strong control over its neighborhood. This lifts the strongly regular behaviour of idiscernible sequences to sequences of disjoint α-neighborhoods. The main result of this section, Theorem 11, is the backbone of our flip-flatness proof and states the following. In every large sequence of disjoint α-neighborhoods, we will find a still-large subsequence such that the φ-connections

of every element a towards all α-neighborhoods in the subsequence can be described by a bounded set of sample elements in the following sense. After possibly omitting one exceptional neighborhood at index ex(a) depending on a, the φ-connections are completely homogeneous (in the stable case) or alternate at most once (in the NIP case). A related (but orthogonal) result was also proved in [38, Lemma 64] using tools from model theory.

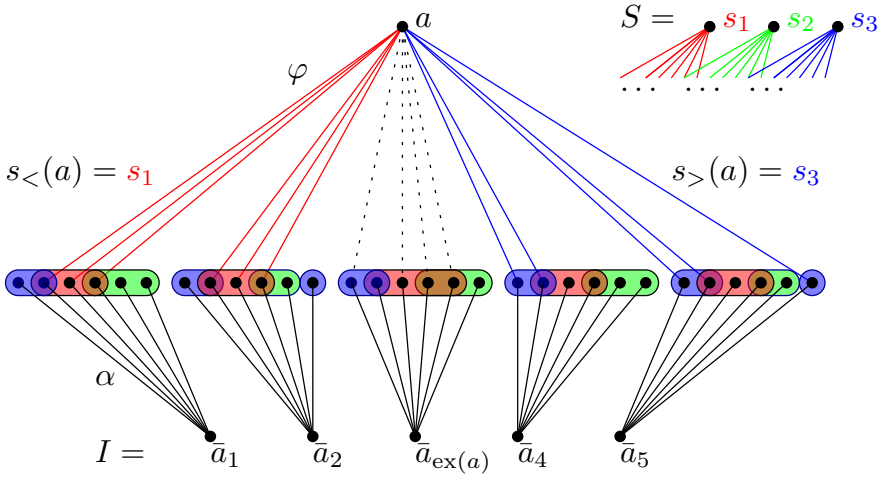

Figure 6 A visualization of Theorem 11. On the bottom: a sequence I with disjoint α-neighborhoods. On the top right: a small set of sample elements S. The φ-neighborhoods of the elements in S in the α-neighborhood of I are colored accordingly. Note that their φ-neighborhoods can overlap. The φ-neighborhood of a in the α-neighborhood of I is equal to the φ-neighborhood of $s_<(a) = s_1$ before the exceptional index $\mathrm{ex}(a)$. After it, it is equal to the φ-neighborhood of $s_>(a) = s_3$. For the α-neighborhood of $\bar{a}_{\mathrm{ex}(a)}$ we make no claims.

▶ **Theorem 11.** *For every monadically NIP class of structures $\mathscr{C}$, and formulas $\varphi(x, y)$ and $\alpha(\bar{x}, y)$, there exists a function $N : \mathbb{N} \to \mathbb{N}$ and an integer k such that for every $m \in \mathbb{N}$ every structure $\mathbb{A} \in \mathscr{C}$ and every finite sequence of tuples $J \subseteq \mathbb{A}^{|\bar{x}|}$ of length $N(m)$ whose α-neighborhoods are disjoint the following holds. There exists a subsequence $I = (\bar{a}_i)_{1 \leqslant i \leqslant m} \subseteq J$ of length m and a set $S \subseteq \mathbb{A}$ of at most k sample elements such that for every element $a \in \mathbb{A}$ there exists an exceptional index $\mathrm{ex}(a) \in [m]$ and a pair $s_<(a), s_>(a) \in S$ such that*

$$\text{for all } 1 \leqslant i < \mathrm{ex(a)} : N_\alpha(\bar{a}_i) \cap N_\varphi(a) = N_\alpha(\bar{a}_i) \cap N_\varphi(s_<(a)), \text{ and}$$
$$\text{for all } \mathrm{ex(a)} < i \leqslant m : N_\alpha(\bar{a}_i) \cap N_\varphi(a) = N_\alpha(\bar{a}_i) \cap N_\varphi(s_>(a)).$$

If $a \in N_\alpha(\bar{a}_i)$ for some $\bar{a}_i \in I$, then $i = \mathrm{ex}(a)$.

If $\mathscr{C}$ is monadically stable, then $s_<(a) = s_>(a)$ for every $a \in \mathbb{A}$.

A visualization of the NIP case of Theorem 11 is provided in Figure 6. An important ingredient of the proof is the following Ramsey-type result due to Ding et al. [12, Corollary 2.4]. We say a bipartite graph with sides $a_1, \ldots, a_\ell$ and $b_1, \ldots, b_\ell$ forms

- a *matching* of order ℓ if a_i and b_j are adjacent if and only if $i = j$ for all $i, j \in [\ell]$,
- a *co-matching* of order ℓ if a_i and b_j are adjacent if and only if $i \neq j$ for all $i, j \in [\ell]$,
- a *ladder* of order ℓ if a_i and b_j are adjacent if and only if $i \leqslant j$ for all $i, j \in [\ell]$.

We call two distinct vertices u and v *twins* in a graph G they have the same neighborhood with regard to the edge relation, i.e. $N_E^G(u) \setminus \{v\} = N_E^G(v) \setminus \{u\}$.

▶ **Theorem 12** ([12, Corollary 2.4]). *There exists a function $Q : \mathbb{N} \to \mathbb{N}$ such that for every $\ell \in \mathbb{N}$ and for every bipartite graph $G = (L, R, E)$ without twins, where L has size at least $Q(\ell)$, contains a matching, co-matching, or ladder of order ℓ as an induced subgraph.*

Equipped with Theorem 10 from the previous subsection and the above Theorem 12, we will now give a proof sketch of the monadically stable case of Theorem 11.

Proof sketch of Theorem 11 (monadically stable case). We will build the sequence I by inductively extracting indiscernible subsequences I_i of J and growing a set of sample elements $S_i = \{s_1, \dots, s_i\}$. During the induction, we maintain the following invariant: every two distinct sample elements from S_i have pairwise different φ-neighborhoods in the α-neighborhood of every tuple of I_i. We start with $I_0 = J$ and $S_0 = \varnothing$ where this trivially holds. Let us now describe the inductive construction of I_{i+1} and S_{i+1}. We first add a constant symbol for every element in S_i to the signature of $\mathbb{A}$. This extended signature will allow us to express for all $j \in [i]$ the formula $\psi_j(x, \bar{y})$ stating that x has the same φ-neighborhood as s_j in the α-neighborhood of $\bar{y}$, i.e., $N_\alpha(\bar{y}) \cap N_\varphi(x) = N_\alpha(\bar{y}) \cap N_\varphi(s_j)$. We collect all these formulas into the set $\Phi = \{\psi_j \ : \ j \in [i]\}$ and build I_{i+1} by extracting a Δ_k^Φ-indiscernible sequence from I_i, for the appropriate value of k given by Theorem 10. Now by Theorem 10, for every element $b \in \mathbb{A}$, there exists a Φ-type $\tau(x, \bar{y})$ such that b is τ-connected to all but at most one tuple from I_{i+1}. Since the Φ-types capture information about the connections relative to vertices in S_i, one of the following two cases applies for every $b \in \mathbb{A}$.

1. The φ-neighborhood of b is different to the φ-neighborhood of every element of S_i in the α-neighborhood around all but at most one tuple of I_{i+1}.
2. For some $j \in [i]$, the φ-neighborhood of b is equal to the φ neighborhood of s_j in the α-neighborhood around all but at most one tuple of I_{i+1} (whose index will $\mathrm{ex}(b)$).

If the first case applies for some $b \in \mathbb{A}$, then we can build S_{i+1} by setting $s_{i+1} := b$, which satisfies our invariant (after we possibly drop one more element from I_{i+1}). If the second case applies for every $b \in \mathbb{A}$, every element is represented by S_i, and the construction stops with $I := I_{i+1}$ and $S := S_i$.

It now remains to show that the construction stops after less than k steps for some k depending only on $\mathscr{C}$, φ, and α. To this end, we show that for every $\ell \in \mathbb{N}$, there exists $k \in \mathbb{N}$ with the following property. If there exists I_k and S_k satisfying the invariant of our construction, then we can derive a formula ψ from φ and α which has pairing index at least ℓ in an expansion of $\mathbb{A}$ with a constant number of colors (in the full proof we use 3 colors). As $\mathscr{C}$ is monadically NIP, the pairing index of ψ is bounded, which also yields a bound for k.

Assume we are given I_k and S_k, such that all the elements from S_k have different φ-neighborhoods, in the α-neighborhoods around the tuples of I_k. By choosing k large enough, repeated application of Theorem 12 and the pigeonhole principle, we find a subset $S' = \{s_1, \dots, s_{\ell^2}\}$ of S_k and a subsequence $I' = \{\bar{a}_1, \dots, \bar{a}_\ell\}$ of I_k with the following property. For each $j \in [\ell]$, there exists a subset N'_j of the α-neighborhood of $\bar{a}_j$, such that the φ-connections from S' to the N'_j all form a matching, all form a co-matching, or all form a ladder. Assume the φ-connections form a matching. Let us now show that the formula

$$\psi(x, \bar{y}) := \exists z \in N_\varphi(x) \cap N_\alpha(\bar{y}).\ R(z)$$

has pairing index at least ℓ in the class $\mathscr{C}$ extended with an additional unary predicate R. The situation is depicted in Figure 7. The formulas φ and α interpret a large 1-subdivided biclique in $\mathbb{A}$ as follows. The tuples from I' and the elements from S' form the principle vertices. The subdivision vertices are formed by $\bigcup_{i \in [\ell]} N'_i$. Due to the assumed matchings, each subdivision element has exactly one incoming φ-connection from S'. Due to the tuples

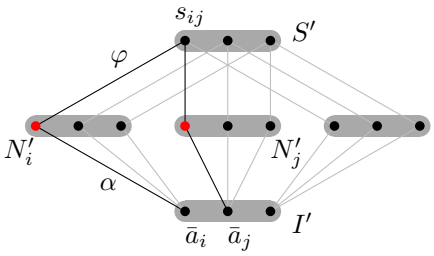

Figure 7 Pairing the elements of I'.

in I' having disjoint α-neighborhoods, each subdivision element has exactly one incoming α-connection from the elements of I'. Since I' has size ℓ and S' has size ℓ^2 we can assign to every pair $\bar{a}_i, \bar{a}_j \in I'$ a unique element $s_{ij} \in S'$. By marking with the predicate R the two unique subdivision elements from $N_\varphi(s_{ij}) \cap N'_i$ and $N_\varphi(s_{ij}) \cap N'_j$, we get that s_{ij} is ψ-connected only to $\bar{a}_i$ and $\bar{a}_j$ among I'. As desired, S' and I' witness that ψ has pairing index at least ℓ.

In case the φ-connections form a co-matching or a ladder, we can replace φ with a derived formula φ', such that the φ'-connections form a matching in a coloring of $\mathbb{A}$, which completes the proof sketch for the size bound of S.

Lastly, in order to ensure that $a \in N_\alpha(\bar{a}_i)$ implies $i = \mathrm{ex}(a)$, we can augment the formula $\psi_j(x, \bar{y})$ in the construction to return false whenever x is contained in $N_\alpha(\bar{y})$. This way, we enforce the single exception given by Theorem 10 to be on a_i. ◀

To prove the more general NIP case, we replace Theorem 10 with Theorem 9, which introduces the possibility for a vertex a to be sampled by two vertices $s_<(a)$ and $s_>(a)$, alternating around $\mathrm{ex}(a)$. We also have to slightly adapt the search for the next sample element s_j using an additional indiscernibility argument.

3.3 Flatness in monadically stable classes of graphs

In this section we use Theorem 11 to characterize monadically stable graph classes in terms of flip-flatness. We start with the forward direction, restated for convenience.

▶ **Theorem 4.** *Every monadically stable class $\mathscr{C}$ of graphs is flip-flat, where for every r, the function N_r is polynomial. Moreover, given an n-vertex graph $G \in \mathscr{C}$, $A \subseteq V(G)$, and $r \in \mathbb{N}$, we can compute a subset $B \subseteq A$ and a set of flips F that makes B distance-r independent in $G \oplus F$ in time $\mathcal{O}(f_{\mathscr{C}}(r) \cdot n^3)$ for some function $f_{\mathscr{C}}$.*

Partial proof. Let $\mathscr{C}$ be a monadically stable class of graphs and $r \in \mathbb{N}$. We want to show that in every graph $G \in \mathscr{C}$, in every set $A \subseteq V(G)$ of size $N_r(m)$ we find a subset $B \subseteq A$ of size at least m and a set F of at most s_r flips such that B is distance-r independent in $G \oplus F$. We will first inductively describe how to obtain the set of vertices B and the set of flips F and bound the runtime and values for N_r and s_r later. In the base case we have $r = 0$. We can pick $B := A$ and $F := \varnothing$ and there is nothing to show.

In the inductive case we assume the result is proved for r and extend it to $r + 1$. Let $r = 2i + p$ for some $i \in \mathbb{N}$ and parity $p \in \{0, 1\}$. Apply the induction hypothesis to obtain s_r flips F_r and a set A_r that is distance-r independent in $G_r := G \oplus F_r$. Note that for every fixed number t of flips, since flips are definable by coloring and a quantifier-free formula, the class $\mathscr{C}_t$ of all graphs obtainable from graphs of $\mathscr{C}$ by at most t flips is monadically stable by Lemma 5. Hence, G_r comes from the monadically stable class $\mathscr{C}_{s_r}$. Our goal is to

find a set of flips F'_{r+1} (of fixed finite size which will determine the number s_{r+1}) together with a set $A_{r+1} \subseteq A_r$, that is distance-$(r+1)$ independent in $G_{r+1} := G_r \oplus F'_{r+1}$. Then $G_r = G \oplus F_{r+1}$ with $F_{r+1} = F_r \cup F'_{r+1}$.

Since the elements in A_r have pairwise disjoint distance-i neighborhoods, we can apply Theorem 11 to A_r with $\varphi(x,y) = E(x,y)$ and $\alpha(x,y) = \mathrm{dist}_{\leqslant i}(x,y)$. Since we are in the monadically stable case, this yields a subset $A_{r+1} \subseteq A_r$ and a small set of sample vertices S such that for every vertex $a \in V(G)$ there exists $s(a) \in S$ and $\mathrm{ex}(a) \in A_{r+1}$ such that a has the same edge-neighborhood as $s(a)$ in the distance-i neighborhoods of all elements from $A_{r+1} \setminus \{\mathrm{ex}(a)\}$. We now do a case distinction depending on whether r is even or odd.

The odd case: $r = 2i+1$. For every $s \in S$ let C_s be the set containing every vertex a for which we have $s(a) = s$ and that is at distance at least $i+1$ from every vertex in A_{r+1}. Let D_s be the set containing every edge-neighbor of s that is at distance exactly i from one of the vertices in A_{r+1}. We now build the set of flips F'_{r+1} by adding for every sample vertex $s \in S$ the flip (C_s, D_s).

Let us now argue that A_{r+1} is distance-$(r+1)$ independent in $G_{r+1} := G_r$. We only flip edges in G_r between pairs of vertices a, b such that a has distance (in G_r) at least $i+1$ and b has distance exactly i to A_{r+1}. It follows that A_{r+1} remains distance-$(2i+1)$ independent in G_{r+1}. Assume towards a contradiction that there exists a path $P = (a, a_1, \ldots, a_i, u, b_i, \ldots, b_1, b)$ of length $r+1 = 2i+2$ between two vertices $a, b \in A_{r+1}$ in G_{r+1}. The distance between a and a_i (resp. b and b_i) is not affected by the flips. Only the connection between a_i (resp. b_i) and u can possibly be impacted. Additionally, note that $u \in C_{s(u)}$.

Since a and b are distinct we have that either $\mathrm{ex}(u) \neq a$ or $\mathrm{ex}(u) \neq b$. By symmetry we can assume the former case. As a_i is in the distance-i neighborhood of a, we have $G_r \models E(u, a_i) \iff G_r \models E(s(u), a_i)$. Observe that if $E(s(u), a_i)$ holds in G_r, then $a_i \in D_{s(u)}$ and therefore the edge (u, a_i) was removed by the flips and is not in G_{r+1}. Similarly, if $E(s(u), a_i)$ does not hold in G_r, then $a_i \notin D_{s(u)}$ and the edge (u, a_i) was not introduced by any flip. We can conclude that P is not a path in G_{r+1}, and that A_{r+1} is distance-$(r+1)$ independent in G_{r+1}.

The even case and runtime analysis. The even distance creates a symmetry that requires additional care to handle, but otherwise proceeds similarly to the odd case. The arguments concerning size bounds and runtimes can be found there as well. The runtime bound crucially uses the fact that we take Δ_k^Φ-indiscernible sequences only with respect to formulas Φ that can be evaluated in polynomial time. ◀

We have shown that for graph classes, monadic stability implies flip-flatness. We now show that the reverse holds as well. We will use the following statement, which is an immediate consequence of Gaifman's locality theorem [18]. For an introduction of the locality theorem see for example [23, Section 4.1].

▶ **Corollary 13** (of [18, Main Theorem]). *Let $\varphi(x,y)$ be a formula. Then there are numbers $r, t \in \mathbb{N}$, where r depends only on the quantifier-rank of φ and t depends only on the signature and quantifier-rank of φ, such that every (colored) graph G can be vertex-colored using t colors in such a way that for any two vertices $u, v \in V(G)$ with distance greater than r in G, $G \models \varphi(u,v)$ depends only on the colors of u and v. We call r the* Gaifman radius *of φ.*

▶ **Lemma 14.** *Every flip-flat class of graphs is monadically stable.*

Proof sketch of Lemma 14. Assume towards a contradiction that there exists a class $\mathscr{C}$ that is not monadically stable but flip-flat. Then there exists a formula $\sigma(x, y)$ and a graph $G \in \mathscr{C}$, such that σ defines an order on a large vertex set A in a coloring G^+. Let r be the Gaifman radius of σ. By flip-flatness there exists a (still large) subset B of A and a bounded size set of flips F, such that B is distance-r independent in $H := G^+ \oplus F$. Let H^+ be the graph where we have marked the flips F with colors in H. We can rewrite σ to a formula σ' of the same quantifier rank such that for all $u, v \in V(G)$ we have $H^+ \models \sigma'(u, v)$ if and only if $G^+ \models \sigma(u, v)$. In particular, σ' orders B in H^+. By Corollary 13 and the pidgeonhole principle, there must be two distinct vertices $u, v \in B$ such that $H^+ \models \sigma'(u, v)$ if and only if $H^+ \models \sigma'(v, u)$. However as σ' orders B, we also have $H^+ \models \sigma'(u, v)$ if and only if $H^+ \not\models \sigma'(v, u)$; a contradiction. ◀

From Theorem 4 and Lemma 14 we conclude the following.

▶ **Theorem 3.** *A class of graphs is monadically stable if and only if it is flip-flat.*

References

1 Algorithms, Logic and Structure Workshop in Warwick – Open Problem Session. URL: `https://warwick.ac.uk/fac/sci/maths/people/staff/daniel_kral/alglogstr/openproblems.pdf`, 2016. [Online; accessed 23-Jan-2023].

2 Hans Adler and Isolde Adler. Interpreting nowhere dense graph classes as a classical notion of model theory. *European Journal of Combinatorics*, 36:322–330, 2014.

3 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.

4 Achim Blumensath. Simple monadic theories and indiscernibles. *Mathematical Logic Quarterly*, 57(1):65–86, 2011. `doi:10.1002/malq.200910121`.

5 Édouard Bonnet, Jan Dreier, Jakub Gajarský, Stephan Kreutzer, Nikolas Mählmann, Pierre Simon, and Szymon Torunczyk. Model checking on interpretations of classes of bounded local cliquewidth. In Christel Baier and Dana Fisman, editors, *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2 - 5, 2022*, pages 54:1–54:13. ACM, 2022. `doi:10.1145/3531130.3533367`.

6 Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. Twin-width IV: Ordered graphs and matrices. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, pages 924–937, New York, NY, USA, 2022. Association for Computing Machinery. `doi:10.1145/3519935.3520037`.

7 Édouard Bonnet, Dibyayan Chakraborty, Eun Jung Kim, Noleen Köhler, Raul Lopes, and Stéphan Thomassé. Twin-width viii: delineation and win-wins, 2022. `doi:10.48550/arXiv.2204.00722`.

8 Samuel Braunfeld and Michael Laskowski. Characterizations of monadic NIP. *Transactions of the American Mathematical Society, Series B*, 8(30):948–970, 2021.

9 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. `doi:10.1016/j.jctb.2006.04.003`.

10 Anuj Dawar. Homomorphism preservation on quasi-wide classes. *Journal of Computer and System Sciences*, 76(5):324–332, 2010.

11 Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.

12 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`.

13 Jan Dreier, Jakub Gajarský, Sandra Kiefer, Michał Pilipczuk, and Szymon Toruńczyk. Tree-like decompositions for transductions of sparse graphs. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2022)*, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3531130.3533349.

14 Jan Dreier, Nikolas Mählmann, Amer E. Mouawad, Sebastian Siebertz, and Alexandre Vigny. Combinatorial and Algorithmic Aspects of Monadic Stability. In Sang Won Bae and Heejin Park, editors, *33rd International Symposium on Algorithms and Computation (ISAAC 2022)*, volume 248 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:17, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ISAAC.2022.11.

15 Jan Dreier, Nikolas Mählmann, Sebastian Siebertz, and Szymon Toruńczyk. Indiscernibles and flatness in monadically stable and monadically nip classes. *arXiv preprint*, 2022. arXiv:2206.13765.

16 Jan Dreier, Nikolas Mählmann, and Sebastian Siebertz. First-order model checking on structurally sparse graph classes, 2023. doi:10.48550/arXiv.2302.03527.

17 Andrzej Ehrenfeucht and Andrzej Mostowski. Models of axiomatic theories admitting automorphisms. *Fundamenta mathematicae*, 1(43):50–68, 1956.

18 Haim Gaifman. On local and non-local properties. In *Proceedings of the Herbrand Symposium*, volume 107 of *Stud. Logic Found. Math.*, pages 105–135. Elsevier, 1982.

19 Jakub Gajarskỳ, Stephan Kreutzer, Jaroslav Nešetřil, Patrice Ossona De Mendez, Michał Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. First-order interpretations of bounded expansion classes. *ACM Transactions on Computational Logic (TOCL)*, 21(4):1–41, 2020.

20 Jakub Gajarský, Nikolas Mählmann, Rose McCarty, Pierre Ohlmann, Michał Pilipczuk, Wojciech Przybyszewski, Sebastian Siebertz, Marek Sokołowski, and Szymon Toruńczyk. Flipper games for monadically stable graph classes. *arXiv preprint*, 2023. arXiv:2301.13735.

21 Jakub Gajarský, Michał Pilipczuk, and Szymon Toruńczyk. Stable graphs of bounded twin-width. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '22, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3531130.3533356.

22 Colin Geniet and Stéphan Thomassé. First order logic and twin-width in tournaments and dense oriented graphs, 2022. doi:10.48550/arXiv.2207.07683.

23 Martin Grohe. Logic, graphs, and algorithms. *Logic and Automata*, 2:357–422, 2008.

24 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *Journal of the ACM (JACM)*, 64(3):1–32, 2017.

25 Wilfrid Hodges. *A Shorter Model Theory*. Cambridge University Press, 1997.

26 Stephan Kreutzer, Roman Rabinovich, and Sebastian Siebertz. Polynomial kernels and wideness properties of nowhere dense graph classes. *ACM Transactions on Algorithms (TALG)*, 15(2):1–19, 2018.

27 O-joung Kwon, Michał Pilipczuk, and Sebastian Siebertz. On low rank-width colorings. *European Journal of Combinatorics*, 83:103002, 2020.

28 M. Malliaris and S. Shelah. Regularity lemmas for stable graphs. *Transactions of the American Mathematical Society*, 366(3):1551–1585, 2014. URL: http://www.jstor.org/stable/23813167.

29 Jaroslav Nešetřil and Patrice Ossona De Mendez. On nowhere dense graphs. *European Journal of Combinatorics*, 32(4):600–617, 2011.

30 Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. Rankwidth meets stability. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2014–2033. SIAM, 2021.

31 Jaroslav Nešetřil, Patrice Ossona de Mendez, and Sebastian Siebertz. Structural properties of the first-order transduction quasiorder. In *30th EACSL Annual Conference on Computer Science Logic (CSL 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.

32 Jaroslav Nešetřil, Roman Rabinovich, Patrice Ossona de Mendez, and Sebastian Siebertz. Linear rankwidth meets stability. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1180–1199. SIAM, 2020.

33 Bruno Poizat. *A Course in Model Theory: An Introduction to Contemporary Mathematical Logic*. Springer New York, 2000.

34 Saharon Shelah. Monadic logic: Hanf numbers. In *Around classification theory of models*, pages 203–223. Springer, 1986.

35 Pierre Simon and Szymon Toruńczyk. Ordered graphs of bounded twin-width, 2021. doi:10.48550/arXiv.2102.06881.

36 Katrin Tent and Martin Ziegler. *A course in model theory*, volume 40 of *Lecture Notes in Logic*. Cambridge University Press, 2012.

37 Symon Toruńczyk. Flip-width: Cops and robber on dense graphs, 2023. arXiv:2302.00352.

38 Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. Twin-width iv: ordered graphs and matrices, 2021. arXiv:2102.03117.