Monadically Stable and Monadically Dependent Graph Classes

Characterizations and Algorithmic Meta-Theorems

Nikolas Mählmann

Dissertation zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften
(Dr. rer. nat.)

1. Gutachter: Prof. Dr. Sebastian Siebertz
2. Gutachter: Dr. habil. Patrice Ossona de Mendez
Datum der Verteidigung: 5. September 2024

Universität Bremen
Fachbereich 3 - Mathematik und Informatik

Abstract

A graph class is *monadically stable* if it does not encode the class of all linear orders using first-order logic and vertex colors. This includes many sparse classes like planar graphs, bounded degree, bounded tree-width, and nowhere dense classes, but also dense classes like map graphs. More generally, a class is *monadically dependent* (also known as *monadically NIP*) if it does not encode the class of all graphs. This includes the aforementioned monadically stable classes, and also classes of bounded clique- or twin-width. Originating in model theory, monadic stability and dependence have predominantly been studied on infinite structures. In this thesis we combine tools from combinatorics and logic, to develop a theory for monadically stable and monadically dependent classes of finite graphs that is well suited for their algorithmic treatment.

We obtain the following structure/non-structure dichotomy. On the structure side, we characterize monadic stability and monadic dependence by two Ramsey-theoretic properties called *flip-flatness* and *flip-breakability*. This gives rise to a larger framework: natural restrictions of flip-flatness and flip-breakability characterize nowhere denseness, bounded clique- and tree-width, and shrub- and tree-depth. On the non-structure side, we characterize monadic stability and monadic dependence by explicitly listing few families of forbidden induced subgraphs.

We show the algorithmic applicability of our characterizations by proving new tractability and hardness results for the *first-order model checking problem*. Given a graph G and a first-order formula φ, we want to check whether G satisfies φ. It is conjectured that a hereditary graph class admits fixed-parameter tractable model checking if and only if it is monadically dependent. Building on flip-flatness, we prove a game characterization of monadic stability called *Flipper game*. Using the game tree of the Flipper game as a decomposition of the input graph, we show that first-order model checking is fixed-parameter tractable on every monadically stable graph class. This confirms an important case of the tractability side of the model checking conjecture. Using the forbidden induced subgraph characterization for monadically dependent classes, we completely resolve the hardness side: we show that first-order model checking is AW[$*$]-hard on every hereditary graph class that is not monadically dependent.

Zusammenfassung

Eine Graphklasse ist *monadisch stabil*, wenn sie nicht die Klasse aller linearen Ordnungen mithilfe von Prädikatenlogik und Knotenfarben kodiert. Viele Graphklassen sowohl mit geringer als auch mit hoher Kantendichte haben diese Eigenschaft. Beispiele sind Klassen in denen der Maximalgrad oder die Baumweite beschränkt ist. Auch planare Graphen, Map Graphen, und nowhere dense Klassen sind monadisch stabil. Darüber hinaus ist eine Graphklasse *monadisch abhängig* (oder auch *monadisch NIP*), wenn sie nicht die Klasse aller Graphen kodiert. Alle monadisch stabilen Klassen und auch Klassen mit beschränkter Cliquen- oder Zwillingsweite sind monadisch abhängig. Monadische Stabilität und monadische Abhängigkeit haben ihren Ursprung in der Modelltheorie und wurden bisher hauptsächlich auf unendlichen Strukturen untersucht. In dieser Arbeit kombinieren wir Werkzeuge aus der Kombinatorik und der Logik, um eine algorithmische Theorie für monadisch stabile und monadisch abhängige Klassen endlicher Graphen zu entwickeln.

Wir beweisen die folgende strukturelle Dichotomie. Auf der einen Seite charakterisieren wir monadische Stabilität und monadische Abhängigkeit durch zwei Ramsey-theoretische Eigenschaften, die wir *Flip-Flatness* und *Flip-Breakability* nennen. Wir präsentieren diesen Zusammenhang als Teil eines größeren Frameworks: Natürliche Varianten von Flip-Flatness und Flip-Breakability charakterisieren Nowhere Denseness und beschränkte Cliquenweite, Baumweite, Strauchtiefe, und Baumtiefe. Auf der anderen Seite beschreiben wir minimale Familien von induzierten Subgraphen, die zu monadischer Instabilität und monadischer Unabhängigkeit führen.

Wir zeigen, dass unsere neuen Charakterisierungen sich gut für algorithmische Anwendungen eignen, indem wir neue obere und untere Komplexitätsschranken für das *Model Checking Problem der Prädikatenlogik* beweisen. In diesem Problem sollen wir für einen gegebenen Graphen G und Satz φ in Prädikatenlogik entscheiden, ob φ wahr auf G ist. Es wird vermutet, dass Model Checking fixed-parameter tractable auf einer hereditären Graphklasse ist, genau dann, wenn diese Klasse monadisch abhängig ist. Aufbauend auf *Flip-Flatness* charakterisieren wir monadische Stabilität durch das sogenannte *Flipper Spiel*. Indem wir den Spielbaum des Flipper Spiels als Zerlegung für den Eingabegraphen benutzen, können wir zeigen, dass Model Checking fixed-parameter tractable auf allen monadisch stabilen Graphklassen ist. Damit etablieren wir die obere Komplexitätsschranke für einen wichtigen Spezialfall der Model Checking Vermutung. Mithilfe unserer Charakterisierung durch induzierte Subgraphen können wir die untere Komplexitätsschranke sogar im allgemeinen Fall beweisen: Wir zeigen, dass Model Checking auf jeder hereditären, monadisch unabhänigen Graphklasse AW[$*$]-schwer ist.

Acknowledgements

First, I want to thank my supervisor Sebastian Siebertz for his support and guidance. You introduced me to this field of research and gave me all the freedom and opportunities needed to pursue it. These three years of being a PhD student in your group have been a fantastic journey!

Many thanks go to Jan Dreier and Szymon Toruńczyk who quickly became close collaborators, mentors, and friends. I always look forward to our meetings, where discovery is right around the corner. I also want to thank Michał Pilipczuk for hosting my stay at the University of Warsaw. This gratitude extends to the whole group in Warsaw for welcoming me so warmly. I had a great time and want to return soon.

I thank Patrice Ossona de Mendez for reviewing this thesis. I thank my co-authors Édouard Bonnet, Jan Dreier, Ioannis Eleftheriadis, Jakub Gajarský, Stephan Kreutzer, Rose McCarty, Amer Mouawad, Pierre Ohlmann, Michał Pilipczuk, Wojciech Przybyszewski, Sebastian Siebertz, Marek Sokołowski, Szymon Toruńczyk, and Alexandre Vigny. Working with all of you was a pleasure. Special thanks go to Édouard, Jan, Jakub, Stephan, and Szymon, for "adopting" me at the Dagstuhl sparsity workshop when I just started my PhD. Writing our LICS paper together, I learned a lot from you, and it led to my first conference travel.

Next I want to thank the present and past members of our group in Bremen. Special thanks go to Mario for being the best office mate; to Nicole, for her great pastries and our chaotic but fun journey to Montpellier; and to Alex, for his excellent teaching during my master's studies.

Last but not least, this thesis would not have been possible without the most important people in my life: I thank my friends for keeping my spirits up, my sisters Franziska and Emily for putting up with me, my parents Teresa and Dirk for their encouragement and support, and Lisa for her love and patience.

Contents

Part I

Prelude

Chapter 1

Motivation

Given a graph G and a sentence φ in first-order logic, the *(first-order) model checking problem* asks whether G satisfies φ (written as $G \models \varphi$). Here, the graph G is represented as a structure whose universe consists of the vertices of G and where the edges of G are modeled by a symmetric and irreflexive binary relation $E(x, y)$. Using this representation, many parameterized graph problems can be elegantly formulated as first-order sentences. For example, the sentences

$$\exists x_1 \dots \exists x_k \bigwedge_{i \neq j \in [k]} E(x_i, x_j) \quad \text{and} \quad \exists x_1 \dots \exists x_k \forall y \bigvee_{i \in [k]} E(x_i, y) \vee x_i = y$$

ask whether the graph G contains a clique or dominating set of size k, respectively. An algorithm for the first-order model checking problem can therefore solve k-Clique, k-Dominating-Set, and every other problem definable in first-order logic. Instead of solving a single algorithmic problem, model checking algorithms provide a uniform way to solve whole classes of problems. For this reason, they are often referred to as *algorithmic meta-theorems* [46, 47, 54].

This flexibility comes at a price. A naive branching algorithm solves the first-order model checking problem in time $n^{O(q)}$, where n is the number of vertices of G and q is the quantifier rank of φ. Assuming the exponential time hypothesis (ETH) [52, 53], this running time cannot be improved: the k-Clique problem is first-order definable with k quantifiers, but cannot be solved in time $f(k) \cdot n^{o(k)}$ for any function f [11]. Here we assume the input graph is an arbitrary graph for which no further structure is known. In other words, the hardness result only concerns the model checking problem on the *class of all graphs*[1]. However, graphs arising from real world applications are often highly structured [19]. It is therefore natural to consider the model checking problem where the domain is restricted to graph classes where better runtimes are possible. This raises the following question.

> Which graph classes admit *fixed-parameter tractable* model checking?

We say the model checking problem is *fixed-parameter tractable*[2] on a graph class $\mathcal{C}$ if it can be solved in time $f(|\varphi|) \cdot n^c$ for some function f and constant c, on every n-vertex graph from $\mathcal{C}$ and sentence φ. Hence, we demand the dependence of the running time on the size of the graph to be a fixed polynomial that is independent of the input sentence. Note that the naive $n^{O(q)}$ algorithm does not meet this criterion.

Tree- and Clique-Width. The prototypical examples of algorithmic meta-theorems are the results of Courcelle et al. stating that the model checking problem for the more expressive

[1] A *graph class* is a (usually infinite) set of graphs. We identify isomorphic graphs.
[2] We discuss further variants of fixed-parameter tractability in Section 12.4.

monadic second-order logic is fixed-parameter tractable on every class of *bounded tree-* or *clique-width* [13, 14]. The *tree-width* of a graph measures how similar it is to a tree: the lower the tree-width, the more tree-like the graph and the better its decomposability. A graph class $\mathcal{C}$ has *bounded* tree-width, if there is a bound k such that every graph in $\mathcal{C}$ has tree-width at most k. *Clique-width* generalizes the concept of tree-width, by introducing suitable decompositions for dense graphs. Tree- and clique-width are expected to characterize tractable monadic second-order model checking on *monotone*[3] and *hereditary*[4] graph classes [42, 56]. As we will see next, the less expressive first-order logic allows model checking on a much wider range of graph classes. A hierarchy of graph classes where first-order model checking is conjectured to be fixed-parameter tractable is depicted in Figure 1.1. We will discuss the various graph classes in the following.

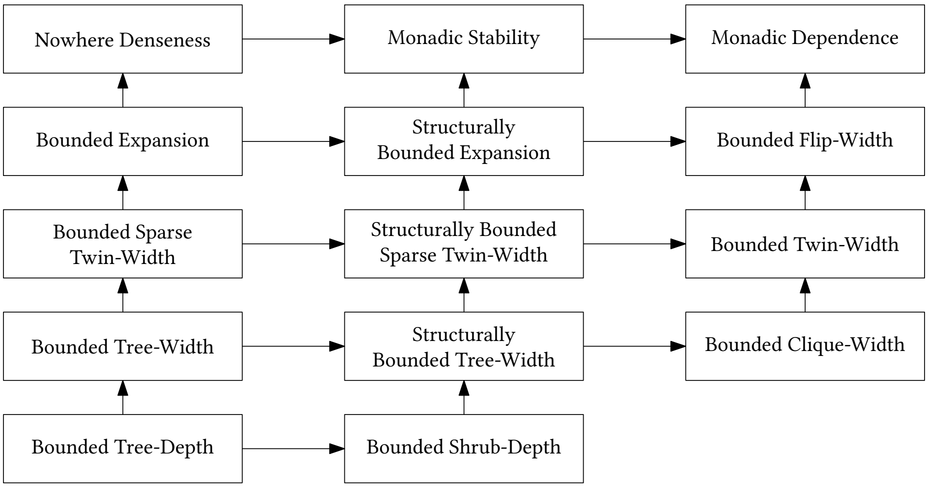

Figure 1.1: The hierarchy of class properties for which the first-order model checking problem is conjectured to be fixed-parameter tractable. An arrow $P_1 \to P_2$ between two properties means that every graph class that has property P_1 also has property P_2. In particular, *bounded tree-depth* is the most restricted class property in this hierarchy and *monadic dependence* is the most general one. *Sparse twin-width* [6, 41], *bounded expansion* [65], and *flip-width* [83] are not discussed in this thesis. We still include them in this hierarchy for completeness. We discuss the "hole" in the bottom right corner of the hierarchy in Chapter 17.

Nowhere Denseness. A long line of research [77, 33, 35, 17, 31] has culminated in the result of Grohe, Kreutzer, and Siebertz, who showed that first-order model checking is fixed-parameter tractable on every *nowhere dense* graph class [48]. Nowhere dense classes were introduced by Nešetřil and Ossona de Mendez [61] as a general framework for characterizing *sparsity* of graph classes. In particular, nowhere dense classes include all classes that have bounded degree, have bounded tree-width, are planar, exclude a minor, or have bounded expansion. The model checking result for nowhere dense classes marks the exact limit of tractability in monotone graph classes, as first-order model checking is AW[$*$]-hard on every monotone class that is not nowhere dense [31, 54].

> A **monotone** graph class admits fixed-parameter tractable model checking if and only if it is **nowhere dense** (assuming FPT $\neq$ AW[$*$]). [48, 31, 54]

Here, FPT $\neq$ AW[$*$] is a standard complexity assumption, equivalent to the assumption that the class of all graphs does not admit fixed-parameter tractable model checking.

[3] A graph class is *monotone* if it is closed under taking subgraphs, i.e., deleting vertices and edges.
[4] A graph class is *hereditary* if it is closed under taking induced subgraphs, i.e., deleting vertices.

While nowhere denseness is well suited to capture tractability on sparse classes, it fails to do so on dense classes: model checking is tractable on every class of bounded clique-width, but already the simple class of all cliques (which has clique-width 1) is not nowhere dense. Here we see that, in the statement above, the focus on monotone classes is very restrictive. The class of all cliques is not monotone and the only monotone class which contains all cliques is the class of all graphs. In order to establish tractability on more general graph classes, we must therefore relax the monotonicity restriction, which gives rise to the following question.

> Which **hereditary** graph classes admit fixed-parameter tractable model checking?

The restriction to hereditary classes is well suited for the study of dense graphs. The class of all cliques is hereditary, and taking the hereditary closure of a class does not increase its clique-width.

Twin-Width. For a long time, not many significant examples of dense classes with tractable model checking apart from classes of bounded clique-width were known. This changed with the introduction of the graph parameter *twin-width* [8]. Classes of bounded twin-width generalize classes of bounded clique-width and classes excluding a minor. There is a fixed-parameter tractable model checking algorithm for classes of bounded twin-width, that requires as additional input a *contraction sequence* witnessing that the input graph has low twin-width [8]. As of June 2024, it is not known how to efficiently compute or approximate suitable contraction sequences in general, but they can be efficiently computed in classes of *ordered graphs*[5]. There, twin-width forms the exact tractability limit for the model checking problem.

> A **hereditary** class of **ordered** graphs admits fixed-parameter tractable model checking if and only if it has **bounded twin-width** (assuming FPT $\neq$ AW$[*]$). [7]

For classes of unordered graphs, bounded twin-width is incomparable to nowhere denseness and fails to characterize tractability in hereditary classes of (unordered) graphs. This is witnessed by the class of graphs with maximum degree 3. This class is nowhere dense and therefore admits efficient model checking, but it has unbounded twin-width [6]. Towards characterizing tractable model checking on hereditary graph classes, we ask:

> Which notion generalizes both nowhere denseness and bounded twin-width?

Structural Nowhere Denseness. In order to generalize nowhere denseness we introduce *transductions*, a notion that originates from model theory. A transduction T_φ is an operation that is specified by a symmetric, irreflexive, binary first-order formula $\varphi(x, y)$ over the signature of colored graphs. It maps an input graph G to a set of output graphs $T_\varphi(G)$ obtained by

1. coloring G,
2. replacing the edge relation of G by the relation defined by φ in G,
3. taking an induced subgraph.

See Figure 1.2 for an example. A formal definition is given in the preliminaries (Chapter 4).

Transductions provide a flexible way to define graph transformations using logic. For example the formulas $\varphi_1(x, y) := \neg E(x, y)$ and $\varphi_2(x, y) := \text{dist}(x, y) \leqslant 2$ specify the transductions that produce the complement and the square of a graph, respectively. Let $\mathcal{C}, \mathcal{D}, \mathcal{E}$ be graph classes. We say $\mathcal{C}$ *transduces* $\mathcal{D}$ (or $\mathcal{D}$ is *transducible* in $\mathcal{C}$) if there exists a transduction T_φ such that $\mathcal{D} \subseteq T_\varphi(\mathcal{C})$, where $T_\varphi(\mathcal{C}) := \bigcup_{G \in \mathcal{C}} T_\varphi(G)$. This relation is transitive: if $\mathcal{C}$ transduces $\mathcal{D}$ and $\mathcal{D}$ transduces $\mathcal{E}$, then also $\mathcal{C}$ transduces $\mathcal{E}$.

[5] *Ordered graphs* are graphs equipped with an additional relation that expresses a linear order on the vertex set and is accessible in first-order logic. Adding an order to a graph can increase both the twin-width of the graph and the expressiveness of first-order logic on this graph.

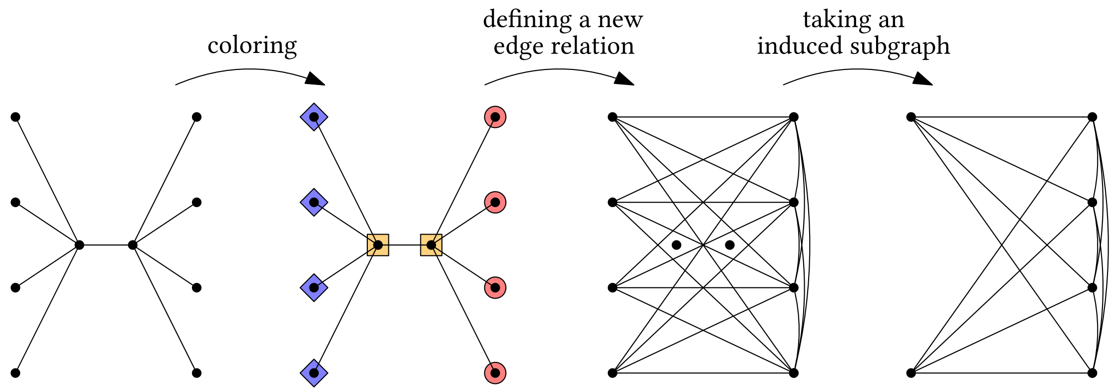

Figure 1.2: Depiction of a transduction. On the very left: the graph G. On the very right: A graph contained in $T_\varphi(G)$ for the formula $\varphi(x, y) = (\text{dist}(x, y) = 3) \vee (\text{Red}(x) \wedge \text{Red}(y))$.

The *class property*[6] of having bounded twin-width is *transduction-closed*: Every class transducible in a class of bounded twin-width again has bounded twin-width [8]. Transduction-closed class properties are very robust as they are preserved under all operations expressible by transductions. In particular, as each graph class transduces its hereditary closure (i.e. the closure under taking induced subgraphs), transduction-closed properties are well suited for studying hereditary graph classes. Nowhere denseness is not a transduction-closed class property. For example the nowhere dense class of edgeless graphs transduces the class of all cliques that is not nowhere dense. To generalize nowhere dense classes, we define *structurally nowhere dense* classes as graph classes transducible from nowhere dense classes [38, 63]. Structural nowhere denseness is a transduction-closed class property due to the transitivity of the transduction relation. It vastly generalizes nowhere denseness: While nowhere dense classes are necessarily sparse, structurally nowhere dense classes can also be dense. For example the class of map graphs is transducible from the class of planar graphs [12]. It has been conjectured that the fixed-parameter tractable model checking can be lifted from nowhere dense classes to structurally nowhere dense classes [37]. However, until now, not even special cases like k-Independent-Set could be solved.

Monadic Stability. Is structural nowhere denseness the sought after generalization of nowhere denseness and bounded twin-width? The answer to this question is "no". To obtain this answer we have to find a class $\mathcal{H}$ satisfying the following two conditions.

$\mathcal{H}$ is not transducible from a nowhere dense class and $\mathcal{H}$ has bounded twin-width.

It is already non-obvious that there exists a class that is not transducible from any nowhere dense class. Note that if no such class would exist, then every class would be structurally nowhere dense and this class property would be trivial. A suitable class $\mathcal{H}$ is provided by a result of Podweski and Ziegler [72] and Adler and Adler [1], who built a bridge between structural graph theory and model theory by showing that nowhere dense graph classes are *monadically stable*. A class is monadically stable if it does not transduce the class of all *half-graphs* (see Figure 1.3).

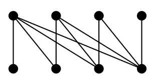

Figure 1.3: The half-graph of order n consists of vertices $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$ such that for all $i, j \in [n]$, a_i is adjacent to b_j if and only if $i \leqslant j$. Depicted: the half-graph of order 4.

[6] Formally, a *class property* is a (usually infinite) set of graph classes. Examples discussed so far include 1. nowhere denseness: the set of all nowhere dense classes, 2. bounded twin-width: the set of all classes with bounded twin-width.

Monadic stability was first studied by Baldwin and Shelah [4] in the context of stability theory, a branch of model theory pioneered by Shelah [79]. One can think of monadically stable classes as "orderless": It is not hard to show from the definition of a half-graph, that we can equivalently characterize monadically stable classes as those that do not transduce the class of all linear orders. The definition of monadic stability exemplifies that transductions can not only be used in *constructive* way, but also in a *restrictive* way:

- Constructive: a class is simple (e.g., structurally nowhere dense) if it is transducible from something simple (e.g., nowhere dense classes).
- Restrictive: a class is simple (e.g., monadically stable) if it does not transduce something complicated (e.g., the class of all half-graphs, the class of all linear orders).

By transitivity of the transduction relation and [72, 1], monadic stability is a transduction-closed non-trivial class property that includes all structurally nowhere dense classes. It is an open question whether there exist classes that are monadically stable but not structurally nowhere dense, and it is conjectured, that both notions coincide [68, 41]. It is easy to prove that the (by definition not monadically stable) class of all half-graphs has bounded twin-width (and even bounded clique-width). This shows that monadically stable and bounded twin-width classes are also incomparable.

Monadic Dependence. We have seen that monadic stability is a transduction-closed class property that generalizes (structural) nowhere denseness but is incomparable to bounded twin-width. It is now time to introduce *monadic dependence*[7] as a common generalization of both monadic stability and bounded twin-width. A graph class is monadically dependent if it does not transduce the class of all graphs. By definition, monadic dependence is the most general transduction-closed class property that is not trivial. It generalizes both monadic stability and bounded twin-width:

- Monadically stable classes are monadically dependent, because they do not transduce the class of all half-graphs (by definition).
- Bounded twin-width classes are monadically dependent, because they do not transduce the class of graphs with maximum degree 3 (by transduction-closure).

As previously discussed, monadic stability can be seen as the "orderless" fragment of monadic dependence. A graph class has bounded twin-width if and only if its graphs can be ordered such that the resulting class of ordered graphs is still monadically dependent [7]. Hence, bounded twin-width can be seen as the "orderable" fragment of monadic dependence. Note that the two notions are not mutually exclusive: for example classes of bounded tree-width or the class of planar graphs are both monadically stable and also have bounded twin-width.

Like monadic stability, monadic dependence originates in stability theory [4]. It was recently shown by Braunfeld and Laskowski that on hereditary classes of structures, monadic stability and monadic dependence coincide with respectively *stability* and *dependence*, two important dividing lines in stability theory used to separate tame from wild structures [9]. Remarkably, in both cases where the tractability limit of the first-order model checking problem has been determined, it is precisely captured by monadic dependence. Monadic dependence …

- …is equivalent to nowhere denseness on monotone classes [1].
- …is equivalent to bounded twin-width on classes of ordered graphs [7].

[7] Monadic dependence is also known as *monadic NIP* where NIP stands for "not the independence property".

Based on these connections between monadic dependence and structural graph theory, the following conjecture has been posed (see for example [2, 37, 7, 5]) and is now the central open question in the area.

Conjecture 1.1. *A hereditary graph class admits fixed-parameter tractable model checking if and only if it is monadically dependent (assuming* $\mathrm{FPT} \neq \mathrm{AW}[*]$*).*

Both directions of this conjecture have been open, with the tractability side being unsolved even for structurally nowhere dense and monadically stable classes. Originating in model theory, both monadic stability and monadic dependence are defined in terms of logic and have mostly been studied on infinite structures. The biggest issue blocking the progress of Conjecture 1.1 is a missing combinatorial understanding of monadic stability and monadic dependence in classes of finite graphs.

Chapter 2

Contribution

In this thesis we build a bridge between model theory and structural graph theory, by providing several combinatorial characterizations of monadic stability and monadic dependence. Our characterizations are well suited for the algorithmic treatment of classes of finite graphs. We utilize them to make significant progress on both the tractability and hardness side of Conjecture 1.1. On the tractability side, we greatly extend the state of the art by showing that all monadically stable classes admit fixed-parameter tractable first-order model checking. We completely resolve the hardness side by showing that first-order model checking is AW$[*]$-hard on every hereditary class that is *monadically independent* (i.e., not monadically dependent). Additionally, we reveal striking conceptual similarities between monadic dependence and many other well known class properties from structural graph theory. We systematically study these similarities under the name "breakability framework". We first state our main theorems and then explain the statements in detail in the remainder of the chapter.

Combinatorial Characterizations of Monadic Stability

Theorem 2.1. *For every graph class $\mathcal{C}$, the following are equivalent.*

(1) $\mathcal{C}$ is monadically stable.

(2) $\mathcal{C}$ is flip-flat.

(3) For every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs

- *all flipped star r-crossings of order k, and*
- *all flipped clique r-crossings of order k, and*
- *all flipped half-graphs of order k.*

(4) For every $r \in \mathbb{N}$ there exists $\ell \in \mathbb{N}$ such that Flipper wins the radius-r budget-2 Flipper game in at most ℓ rounds on every graph from $\mathcal{C}$.

Model Checking on Monadically Stable Classes

Theorem 2.2. *There is an algorithm that, given a graph G and a first-order sentence φ, decides whether $G \models \varphi$, and has the following property.*

For every monadically stable class $\mathcal{C}$, there exists a function $f : \mathbb{N} \times \mathbb{R} \to \mathbb{N}$ such that on any n-vertex graph $G \in \mathcal{C}$ and sentence φ the algorithm runs in time $f(|\varphi|, \varepsilon) \cdot n^{6+\varepsilon}$ for every $\varepsilon > 0$.

Combinatorial Characterizations of Monadic Dependence

Theorem 2.3. *Let $\mathcal{C}$ be a graph class. Then the following are equivalent.*

(1) $\mathcal{C}$ is monadically dependent.

(2) $\mathcal{C}$ is flip-breakable.

(3) For every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs
   - *all flipped star r-crossings of order k, and*
   - *all flipped clique r-crossings of order k, and*
   - *all flipped half-graph r-crossings of order k, and*
   - *the comparability grid of order k.*

(4) The hereditary closure of $\mathcal{C}$ does not efficiently interpret the class of all graphs.

Model Checking Hardness on Monadically Independent Classes

Theorem 2.4. *The first-order model checking problem is* $\mathrm{AW}[*]$*-hard on every hereditary, monadically independent graph class.*

The Breakability Framework

Theorem 2.5. *For every graph class $\mathcal{C}$, the following holds.*

(1) $\mathcal{C}$ is flip-breakable	*if and only if it is*	*monadically dependent.*
(2) $\mathcal{C}$ is flip-flat	*if and only if it is*	*monadically stable.*
(3) $\mathcal{C}$ is deletion-breakable	*if and only if it is*	*nowhere dense.*
(4) $\mathcal{C}$ is deletion-flat	*if and only if it is*	*nowhere dense.*
(5) $\mathcal{C}$ is dist.-∞ flip-breakable	*if and only if it has*	*bounded clique-width.*
(6) $\mathcal{C}$ is dist.-∞ flip-flat	*if and only if it has*	*bounded shrub-depth.*
(7) $\mathcal{C}$ is dist.-∞ deletion-breakable	*if and only if it has*	*bounded tree-width.*
(8) $\mathcal{C}$ is dist.-∞ deletion-flat	*if and only if it has*	*bounded tree-depth.*

It is correct that both deletion-breakability and -flatness correspond to nowhere denseness.

In rest of the chapter, we explain the above statements in detail.

2.1 Flip-Flatness and Flip-Breakability

Our first characterization of monadic stability and monadic dependence is by two Ramsey-theoretic properties called *flip-flatness* and *flip-breakability*. These two properties are natural generalizations of *uniform quasi-wideness*, a property that characterizes nowhere denseness.

Definition 2.6 (Uniform Quasi-Wideness)**.** A graph class $\mathcal{C}$ is *uniformly quasi-wide* if for every radius $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $k_r \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N_r(m)$ there exist sets $S \subseteq V(G)$ with $|S| \leqslant k_r$ and $A \subseteq W \setminus S$ with $|A| \geqslant m$ such that for every two distinct vertices $u, v \in A$:

$$\text{dist}_{G-S}(u, v) > r.$$

Here $G - S$ denotes the induced subgraph of G, where the vertices S have been removed. Intuitively, uniform quasi wideness states that in every huge set W, after the removal of a small set of vertices S, we find a still large subset A of W whose vertices have pairwise distance greater than r. For growing values of r, S may grow and A may shrink. An example is depicted in Figure 2.3 on page 13. Uniform quasi-wideness is a variant of the *quasi-wideness* property introduced by Dawar in the context of homomorphism preservation theorems [16]. Nešetřil and Ossona de Mendez showed that uniform quasi-wideness characterizes nowhere denseness [61, 60].

Fact 2.7 ([61, 60])**.** *A graph class is nowhere dense if and only if it is uniformly quasi-wide.*

Since then, uniform quasi-wideness has become one of the main tools in algorithm design for nowhere dense classes [18, 48, 55].

Flip-Flatness

In order to capture also dense classes, we need a more powerful operation than vertex deletion. For two graphs G and H on the same vertex set, we say H is a *k-flip* of G if it can be obtained by partitioning the vertices of G into at most k parts and complementing the edge relation between some pairs of parts. See Figure 2.1 for an example and Chapter 4 for a formal definition.

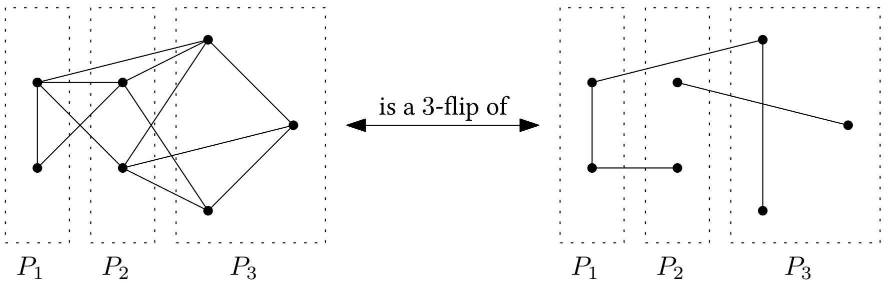

Figure 2.1: Two graphs that are 3-flips of each other. This is witnessed by the *flip partition* $\{P_1, P_2, P_3\}$ where we flip the pairs (P_1, P_2), (P_2, P_3), and (P_3, P_3). In particular, we allow a part to be flipped with itself.

The flip operation is powerful enough to simplify many dense graphs. For example a biclique is a 2-flip of an edgeless graph. However, it is still a very tame operation. It is reversible in the sense that if H is a k-flip of G, then also G is a k-flip of H. Moreover, the edge relation of H is first-order definable in a k-coloring of G, where the colors mark the parts of the flip partition. In particular every graph class transduces the class of all its k-flips.

The following definition was suggested by Jakub Gajarský and Stephan Kreutzer (private communication). It generalizes uniform quasi-wideness, by replacing vertex deletions by flips.

Definition 2.8 (Flip-Flatness)**.** A graph class $\mathcal{C}$ is *flip-flat* if for every radius $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $k_r \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N_r(m)$ there exist a subset $A \subseteq W$ with $|A| \geqslant m$ and a k_r-flip H of G such that for every two distinct vertices $u, v \in A$:

$$\mathrm{dist}_H(u, v) > r.$$

An example is depicted in Figure 2.3. Every uniformly quasi-wide graph class is also flip-flat[1]: instead of deleting a vertex set S in a graph G, we can isolate it in the $2^{|S|}$-flip of G that flips each vertex of S with its neighborhood. As one of the main results of this thesis we show that flip-flatness characterizes monadic stability.

Theorem 2.9. *A graph class is monadically stable if and only if it is flip-flat.*

This characterization is the (to the best of our knowledge) first purely combinatorial characterization of monadic stability. Similar to how uniform quasi-wideness plays a key role in algorithm design for nowhere dense classes, flip-flatness and the techniques developed to prove it are key ingredients of our model checking algorithm for monadically stable graph classes.

Flip-Breakability

Since flip-flatness characterizes monadic stability, the class of all half-graphs cannot be flip-flat. In particular if H is a k-flip of a half-graph and $A \subseteq V(H)$ a size $8k$ set, then at least two vertices of A have distance at most 2 to each other. However, for every half-graph there is a 3-flip that breaks it into two large connected components, as shown in Figure 2.2.

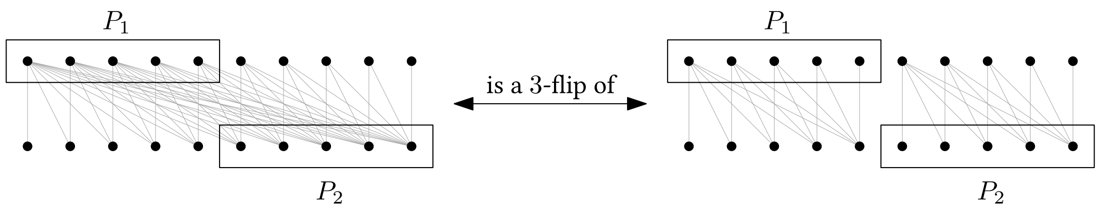

Figure 2.2: A half-graph and one of its 3-flips where we have flipped between the parts P_1 and P_2. The remaining vertices belong to the part P_3 that is not drawn.

This shows that while we cannot produce a large set of vertices that have pairwise high distance in half-graphs, we can instead produce two large sets that are at high distance from each other. Based on this observation, we introduce the following generalization of flip-flatness that we call *flip-breakability*.

Definition 2.10 (Flip-Breakability)**.** A graph class $\mathcal{C}$ is *flip-breakable* if for every radius $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $k_r \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N_r(m)$ there exist subsets $A, B \subseteq W$ with $|A|, |B| \geqslant m$ and a k_r-flip H of G such that:

$$\mathrm{dist}_H(A, B) > r.$$

Here $\mathrm{dist}_H(A, B) > r$ means that $\mathrm{dist}_H(a, b) > r$ for all $a \in A$ and $b \in B$. As one of the main result of this thesis, we show that this notion precisely captures monadic dependence.

Theorem 2.11. *A graph class is monadically dependent if and only if it is flip-breakable.*

[1]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 [83], which is studied in the same context.

This characterization is the (to the best of our knowledge) first purely combinatorial characterization of monadic dependence. Due to the similarity to uniform quasi-wideness and flip-flatness, we believe that flip-breakability is an important step towards establishing fixed-parameter-tractable model checking on monadically dependent classes.

The Breakability Framework

The definitions of uniform quasi-wideness, flip-flatness, and flip-breakability all follow a similar pattern. We modify the graph using either *(1) flips or vertex deletions* and demand that the resulting subset is either *(2) flat or broken*, that is, either pairwise separated or separated into two large sets. We now introduce *the type of separation* as an additional parameterization for this pattern. The type of separation can be either *(3) distance-r or distance-∞*. While *distance-r* separation demands that the vertices or sets have distance at least r from each other (as seen in Definitions 2.6, 2.8 and 2.10), *distance-∞* separation demands the vertices or sets to be in different connected components of the graph. This is formalized by the following definition.

Definition 2.12. A graph class $\mathcal{C}$ is *distance-∞ flip-breakable*, if there exists a function $N : \mathbb{N} \to \mathbb{N}$ and a constant $k \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N(m)$ there exist subsets $A, B \subseteq W$ with $|A|, |B| \geqslant m$ and a k-flip H of G such that in H, no two vertices $a \in A$ and $b \in B$ are in the same connected component.

We show that each of the eight possible combinations of (1), (2) and (3) characterizes a well-studied class property. The results are summarized in the following Table 2.1. We refer to Part IV for formal definitions.

		flatness		breakability	
dist-r	flip-	monadic stability	(Thm. 2.1)	monadic dependence	(Thm. 2.3)
	deletion-	nowhere denseness	[16, 61]	nowhere denseness	(Thm. 13.2)
dist-∞	flip-	bd. shrub-depth	(Thm. 16.2)	bd. clique-width	(Thm. 14.2)
	deletion-	bd. tree-depth	(Thm. 17.2)	bd. tree-width	(Thm. 15.2)

Table 2.1: Variants of flip-breakability.

Notably, the right column of Table 2.1 consists of the conjectured tractability limits of the model checking problem, where the distance-r and distance-∞ variants correspond to first-order and monadic second-order logic, and the flip and deletion variants correspond to hereditary and monotone classes. This further indicates the importance of flip-breakability in the context of model checking. It is interesting to see that the seemingly more general deletion-breakability collapses to deletion-flatness (i.e., uniform quasi-wideness), as both properties characterize nowhere denseness.

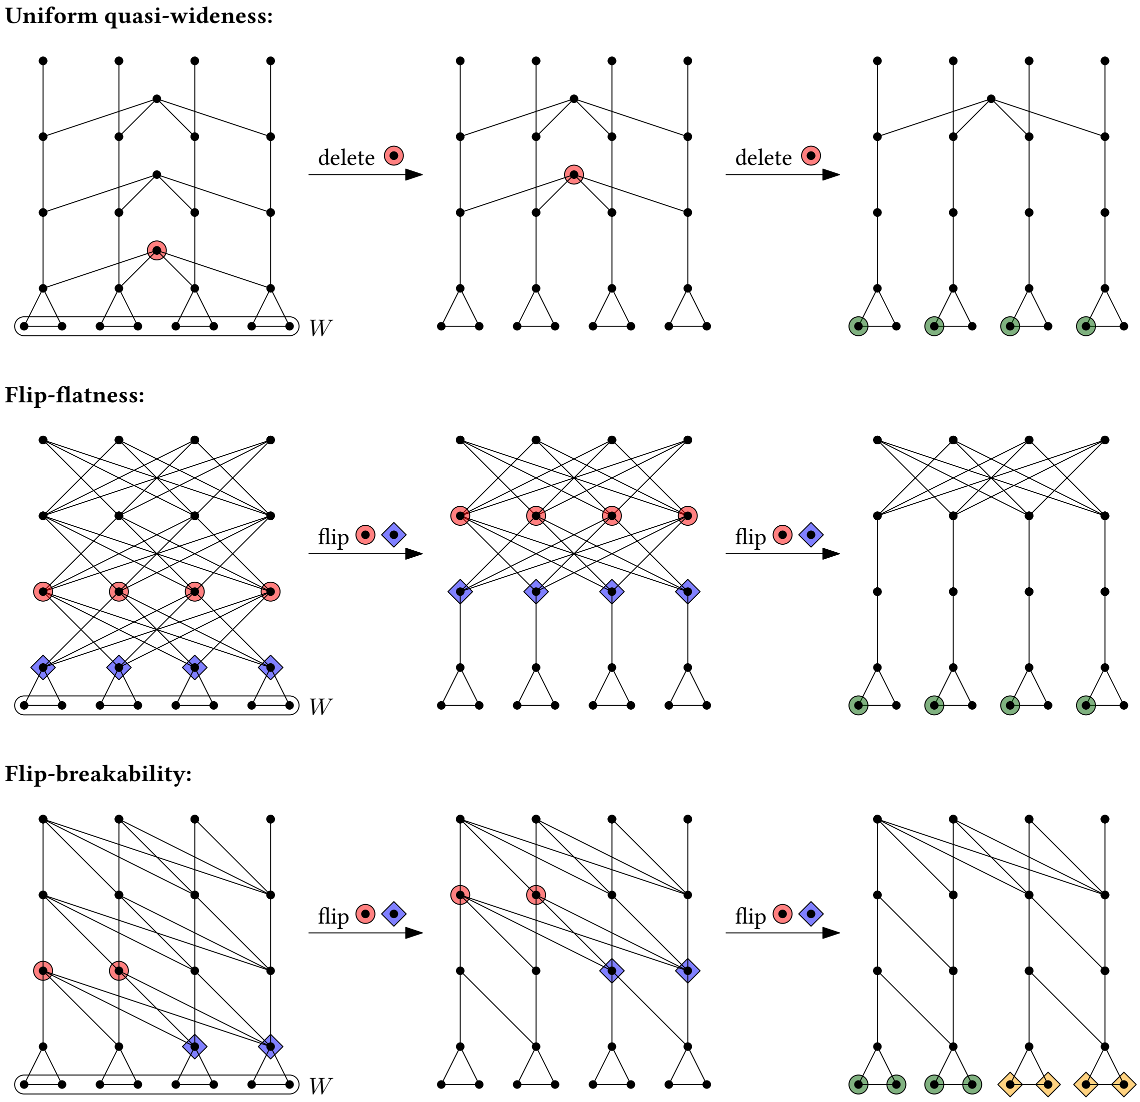

Figure 2.3: Examples for uniform quasi-wideness, flip-flatness, and flip-breakability for radius $r = 7$.

Uniform quasi-wideness: After deleting a few vertices (marked with red circles), a large subset of W (marked with green circles) has pairwise distance greater than 7.

Flip-flatness: After flipping between few sets of vertices (marked with red circles and blue diamonds), a large subset of W (marked with green circles) has pairwise distance greater than 7.

Flip-breakability: After flipping between few sets of vertices (marked with red circles and blue diamonds), two large subsets of W (marked with green circles and yellow diamonds) have distance greater than 7 from each other.

2.2 Forbidden Induced Subgraphs and Hardness

Our second characterization of monadic stability and monadic dependence is by explicitly listing few families for forbidden induced subgraphs.

Forbidden Patterns

Before stating our characterizations, we first introduce the necessary definitions.

Flipped Crossings. For $r \geqslant 1$, the *star r-crossing* of order n is the r-subdivision of $K_{n,n}$ (the biclique of order n). More precisely, it consists of *roots* $a_1, \dots, a_n$ and $b_1, \dots, b_n$ together with r-vertex paths $\{\pi_{i,j} : i, j \in [n]\}$ that are pairwise vertex-disjoint (see Figure 2.4). We denote the two endpoints of a path $\pi_{i,j}$ by start$(\pi_{i,j})$ and end$(\pi_{i,j})$. We require that roots appear on no path, that each root a_i is adjacent to $\{\text{start}(\pi_{i,j}) : j \in [n]\}$, and that each root b_j is adjacent to $\{\text{end}(\pi_{i,j}) : i \in [n]\}$. The *clique r-crossing* of order n is the graph obtained from the star r-crossing of order n by turning the neighborhood of each root into a clique. Moreover, we define the *half-graph r-crossing* of order n similarly to the star r-crossing of order n, where each root a_i is instead adjacent to $\{\text{start}(\pi_{i',j}) : i', j \in [n], i \leqslant i'\}$, and each root b_j is instead adjacent to $\{\text{end}(\pi_{i,j'}) : i, j' \in [n], j \leqslant j'\}$. Each of the three r-crossings contains no edges other than the ones described.

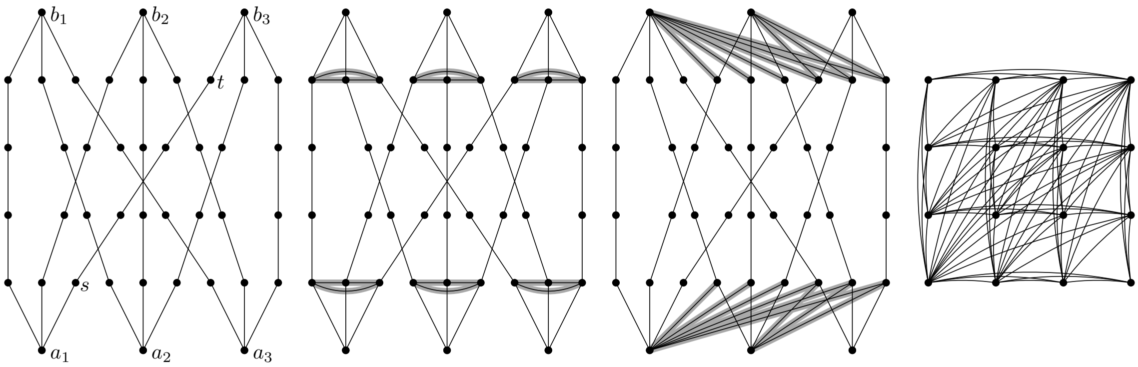

Figure 2.4: From left to right: the star/clique/half-graph 4-crossing of order 3 and the comparability grid of order 4. In the star crossing, the vertices start$(\pi_{1,3})$ and end$(\pi_{1,3})$ are marked s and t. In the clique and half-graph crossing, the edges differing from the star crossing are highlighted.

We also need to consider *flipped* versions of the above patterns. To this end, we partition the vertices of star, clique, and half-graph r-crossings into *layers*: The 0th layer consists of the vertices $\{a_1, \dots, a_n\}$. The lth layer, for $l \in [r]$, consists of the lth vertices of the paths $\{\pi_{i,j} : i, j \in [n]\}$ (that is, the 1st and rth layer, respectively, are $\{\text{start}(\pi_{i,j}) : i, j \in [n]\}$ and $\{\text{end}(\pi_{i,j}) : i, j \in [n]\}$). Finally, the $(r+1)$th layer consists of the vertices $\{b_1, \dots, b_n\}$. A *flipped* star/clique/half-graph r-crossing is a graph obtained from a star/clique/half-graph r-crossing by performing a flip where the parts of the flip partition are the layers of the r-crossing. Note that while there is only one star/clique/half-graph r-crossing of order n, there are multiple flipped star/clique/half-graph r-crossings of order n. Their number is however bounded by $2^{(r+2)^2}$: an upper bound for the number of possible flips for a fixed flip partition of size $(r+2)$.

Flipped Half-Graphs. Recall that the half-graph of order k is the graph on vertices $a_1, \dots, a_k$ and $b_1, \dots, b_k$ where a_i and b_j are adjacent if and only if $i \leqslant j$. A *flipped* half-graph of order k is a 2-flip of a half-graph of order k, where the flip partition has parts $P_1 = \{a_1, \dots, a_k\}$ and $P_2 = \{b_1, \dots, b_k\}$. See Figure 2.5 for an example.

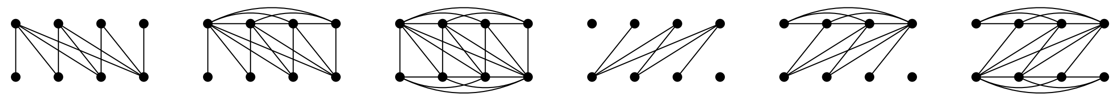

Figure 2.5: The flipped half-graphs of order 4.

Comparability Grids. At last, the *comparability grid* of order n consists of vertices $\{a_{i,j} : i, j \in [n]\}$ and edges between vertices $a_{i,j}$ and $a_{i',j'}$ if and only if either $i = i'$, or $j = j'$, or $i < i' \Leftrightarrow j < j'$. See Figure 2.4 for an example. The reader might wonder why we do not define a *flipped* variant of the comparability grid. This is due to the fact that any huge flipped comparability grid contains a still large non-flipped comparability grid as an induced subgraph.

Characterizations

We are now ready to present our characterizations by forbidden induced subgraphs. For comparison, we fist state a similar characterization of nowhere denseness by forbidden (not necessarily induced) subgraphs.

Fact 2.13 ([61]). *Let $\mathcal{C}$ be a graph class. Then $\mathcal{C}$ is nowhere dense if and only if for every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as a subgraph*

- *the star r-crossing of order k.*

For monadically stable classes, we obtain the following characterization.

Theorem 2.14. *Let $\mathcal{C}$ be a graph class. Then $\mathcal{C}$ is monadically stable if and only if for every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs*

- *all flipped star r-crossings of order k, and*
- *all flipped clique r-crossings of order k, and*
- *all flipped half-graphs of order k.*

In monadically dependent classes, we obtain the following characterization.

Theorem 2.15. *Let $\mathcal{C}$ be a graph class. Then $\mathcal{C}$ is monadically dependent if and only if for every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs*

- *all flipped star r-crossings of order k, and*
- *all flipped clique r-crossings of order k, and*
- *all flipped half-graph r-crossings of order k, and*
- *the comparability grid of order k.*

The flip-flatness and flip-breakability characterizations tell us which structure can be found in monadically stable and dependent classes. They are therefore useful to obtain tractability results. The forbidden induced subgraph characterizations complement this result. They show us which induced subgraphs are contained in classes that are *not* monadically stable or dependent. They are therefore useful to obtain hardness results.

Hardness

We have seen that classes containing arbitrarily large flipped crossings or comparability grids are not monadically dependent. It is instructive to show that for every fixed $r > 1$, the class of all star r-crossings transduces the class of all bipartite graphs.

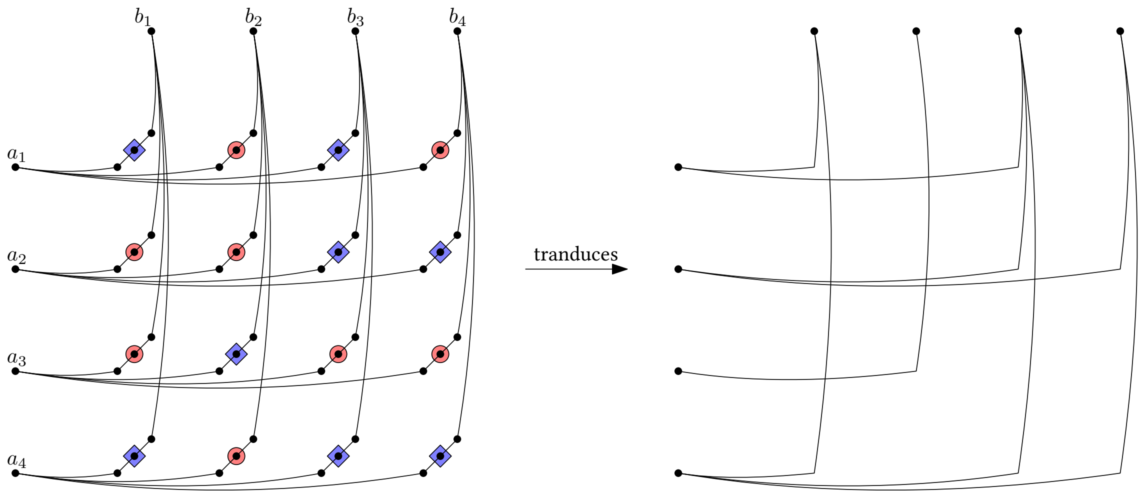

Figure 2.6: Transducing an arbitrary bipartite graph of order 4 from the star 3-crossing of order 4. A suitable transduction would be $\varphi(x, y) =$ "x and y are connected by a path of length 4 that contains a blue vertex".

In Figure 2.6 a star-crossing is presented in a grid-like layout. This layout is the origin of the name "crossing". Using this layout, it is easy to see that the star r-crossing of order n transduces every bipartite graph of order n: we mark the paths we want to keep or remove by colors. A suitable transduction would be

$$\varphi(x, y) := \text{"}x \text{ and } y \text{ are connected by a path of length } r + 1 \text{ containing a blue vertex".}$$

As this formula only depends on r (and not on n), the class of all star-r crossings transduces the class of all bipartite graphs. As the class of all bipartite graphs transduces the class of all graphs and by transitivity, we have shown that the class of all star-r crossings is not monadically dependent. Similarly, also *flipped* star/clique/half-graph crossings as well as comparability grids admit grid-like encodings of bipartite graphs using first-order logic. In particular, a transduction can reverse the flips.

In Figure 2.6, instead of using colors, one can encode bipartite graphs by removing vertices from the source graph: we delete paths between vertices that should not be connected in the target graph. Using our forbidden induced subgraphs characterization, we generalize this observation to all monadically independent classes as follows.

Theorem 2.16. *A hereditary graph class is monadically independent if and only if it efficiently interprets the class of all graphs.*

Here an *interpretation* is a more restrictive version of a transduction. In particular, interpretations do not include a coloring step. Moreover, a graph class $\mathcal{C}$ *efficiently* interprets a graph class $\mathcal{D}$, if there is a polynomial time algorithm that calculates for any graph $G \in \mathcal{D}$ a suitable preimage $H \in \mathcal{C}$. Intuitively the result shows that in hereditary graph classes, where one can take induced subgraphs, the power of vertex coloring is not needed to encode arbitrary graphs in first-order logic. As a corollary, we obtain the hardness side of Conjecture 1.1.

Theorem 2.4. *The first-order model checking problem is* AW[$*$]*-hard on every hereditary, monadically independent graph class.*

Our results also reprove the following result by Braunfeld and Laskowski.

Fact 2.17 ([9]). *Let $\mathcal{C}$ be a hereditary graph class.*

- *$\mathcal{C}$ is monadically stable if and only if it is stable.*
- *$\mathcal{C}$ is monadically dependent if and only if it is dependent.*

Here, *stability* and *dependence* are notions from model theory that generalize monadic stability and monadic dependence. The proof by Braunfeld and Laskowski is much more general. They show a collapse of the monadic and non-monadic variants of stability and dependence on all hereditary classes of relational structures (and not just graphs). Similarly to our induced subgraphs, the proof of Braunfeld and Laskowski also exhibits grid-like configurations (called *pre-coding configurations*) in classes that are monadically independent. As pre-coding configurations are defined in terms of formulas with tuples of free variables, they have a rather logical than combinatorial flair. In particular, it is not clear how to obtain algorithmic hardness results like Theorem 2.4 from them.

2.3 Model Checking and the Flipper Game

As a main contribution of this thesis, we show that first-order model checking is fixed-parameter tractable on every monadically stable graph class.

Theorem 2.2. *There is an algorithm that, given a graph G and a first-order sentence φ, decides whether $G \models \varphi$, and has the following property.*
For every monadically stable class $\mathcal{C}$, there exists a function $f : \mathbb{N} \times \mathbb{R} \to \mathbb{N}$ such that on any n-vertex graph $G \in \mathcal{C}$ and sentence φ the algorithm runs in time $f(|\varphi|, \varepsilon) \cdot n^{6+\varepsilon}$ for every $\varepsilon > 0$.

We point out that the algorithm above is a single algorithm that works for every graph class, but it is only guaranteed to be efficient on monadically stable classes. It was shown in [64], and also follows from our characterizations, that monadic stability and monadic dependence are equivalent on *edge-stable*[2] classes. Hence, our tractability and hardness result for monadically stable and monadically independent classes can be combined to yield another fragment of hereditary classes where tractability precisely coincides with monadic dependence.

Theorem 2.18. *A hereditary and edge-stable graph class admits fixed-parameter tractable*[3] *model checking if and only if it is monadically stable (assuming* FPT $\neq$ AW$[*]$*).*

Note that in this result, the edge-stable fragment of hereditary graph classes strictly generalizes the monotone fragment, for which nowhere denseness was shown to be the tractability limit. This is due to the easy fact that any monotone, non-edge-stable graph class contains all bipartite graphs, and therefore does not admit tractable model checking.

A main ingredient for the proof of Theorem 2.2, which is of independent interest, is the *Flipper game*. The Flipper game is a game characterization of monadic stability. As with our previous characterizations, the Flipper game bears strong resemblance to a characterization of nowhere denseness by a game called the *Splitter game*. The radius-r budget-k Splitter game is played by two players *Splitter* and *Localizer* on a graph G called the *arena*. Intuitively, Splitter wants to show that G is simple, by decomposing it through vertex deletions. In each round of the game, Splitter deletes at most k vertices from the arena and then Localizer restricts the arena to an r-neighborhood: Localizer chooses a center vertex v and all vertices at distance greater than r from v are removed from the arena. Splitter wins the game once the arena contains at most one vertex. With every vertex deletion, the distances in the arena may grow. Therefore, deleting few vertices in Splitters move may result in the deletion of many vertices in Localizers move. See Figure 2.7 for an example play.

[2]A graph class $\mathcal{C}$ is *edge-stable* if and only if there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ contains no flipped half-graph of order k as an induced subgraph.

[3]We discuss different variants of fixed-parameter tractability in Section 12.4.

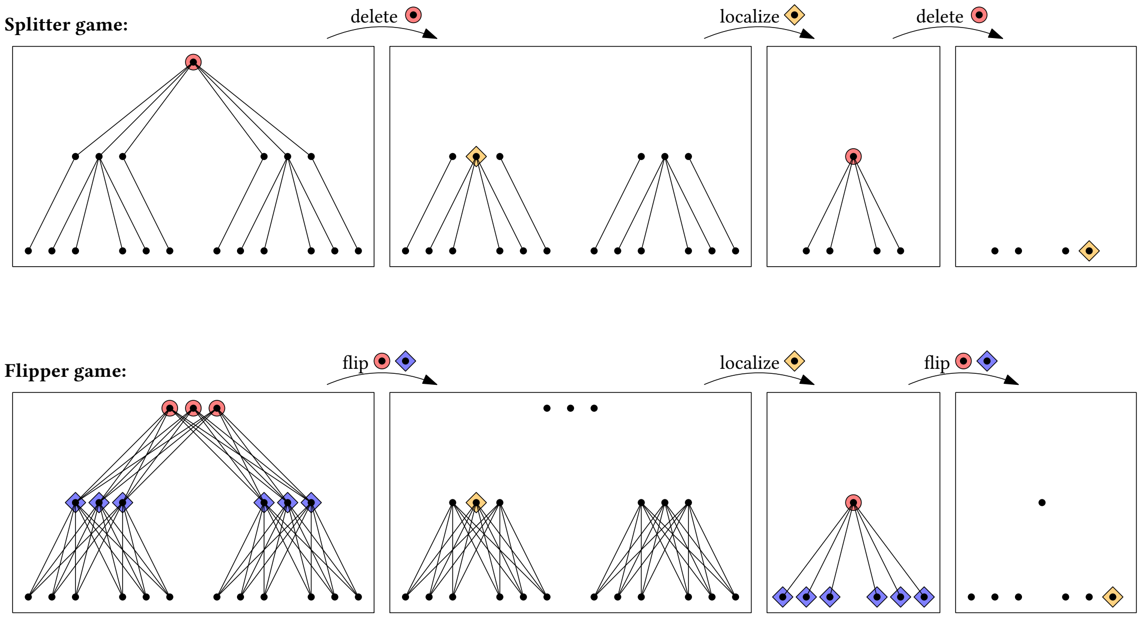

Figure 2.7: On the top: an example play of the radius-1 budget-1 Splitter game. On the bottom: an example play of the radius-1 budget-3 Flipper game. For Flippers moves, two parts of the size 3 flip partition are marked by red circles and blue diamonds. The third part consists of the remaining unmarked vertices.

Splitter wins on any finite graph, given the game is played for sufficiently many rounds. Grohe, Kreutzer, and Siebertz showed that in nowhere dense classes Splitter can always win in a constant number of rounds, independent of the size of the graph.

Fact 2.19 ([48]). *A graph class $\mathcal{C}$ is nowhere dense if and only if for every $r \in \mathbb{N}$ there is $\ell \in \mathbb{N}$ such that Splitter wins the radius-r budget-1 Splitter game in at most ℓ rounds on every graph from $\mathcal{C}$.*

Just as flip-flatness was obtained from uniform quasi-wideness by replacing vertex deletions with flips, we obtain the Flipper game from the Splitter game, by replacing Splitter with the player *Flipper* who, instead of deleting k vertices, performs a k-flip on the arena in every round. See Figure 2.7 for an example play. (The rules of the game are formally defined in Chapter 10.) We show that the Flipper game characterizes monadic stability.

Theorem 2.20. *A graph class $\mathcal{C}$ is monadically stable if and only if for every $r \in \mathbb{N}$ there is $\ell \in \mathbb{N}$ such that Flipper wins the radius-r budget-2 Flipper game in at most ℓ rounds on every graph from $\mathcal{C}$.*

In the model checking algorithm, we use the game-tree of the Flipper game as a bounded depth decomposition of the input graph into r-neighborhoods. We use locality of first-order logic to show that the r-neighborhoods preserve sufficient information to evaluate first-order sentences. As another important ingredient of the algorithm, we prove that monadically stable classes admit *sparse neighborhood covers* that can be used to cluster neighborhoods and thereby keep the size of the game-tree small.

Chapter 3

Bibliographic Remark

The content of this thesis is based on the following publications. All results should be considered equal contributions between the authors.

(P1) *Indiscernibles and Flatness in Monadically Stable and Monadically NIP Classes* [27]
joint work with Jan Dreier, Sebastian Siebertz, Szymon Toruńczyk
presented at ICALP 2023

(P2) *Flipper Games for Monadically Stable Graph Classes* [39]
joint work with Jakub Gajarský, Rose McCarty, Pierre Ohlmann, Michał Pilipczuk, Wojciech Przybyszewski, Sebastian Siebertz, Marek Sokołowski, Szymon Toruńczyk
presented at ICALP 2023

(P3) *First-Order Model Checking on Structurally Sparse Graph Classes* [26]
joint work with Jan Dreier, Sebastian Siebertz
presented at STOC 2023

(P4) *First-Order Model Checking on Monadically Stable Graph Classes* [24]
joint work with Jan Dreier, Ioannis Eleftheriadis, Rose McCarty, Michał Pilipczuk, Szymon Toruńczyk
accepted at FOCS 2024

(P5) *Flip-Breakability: A Combinatorial Dichotomy for Monadically Dependent Graph Classes* [28]
joint work with Jan Dreier, Szymon Toruńczyk
presented at STOC 2024

Parts of this thesis therefore correspond to or are identical to parts of the above publications:

- Part I: *Prelude* is based on (P1) - (P5).
- Part II: *Monadic Dependence* is based on (P5).
- Part III: *Monadic Stability* is based on (P1) - (P5).
  - Chapter 8: *Flip-Flatness* is based on (P1) and (P5).
  - Chapter 9: *Forbidden Induced Subgraphs* is based on (P4) and (P5).
  - Chapter 10: *Flipper Game* is based on (P2).
  - Chapter 11: *Neighborhood Covers* is based on (P4).
  - Chapter 12: *Model Checking* is based on (P3).
- Part IV: *The Breakability Framework* is based on (P5).

Results of publication (P4) may also appear in the future PhD thesis of Ioannis Eleftheriadis.

Note that the publications are listed in order of their first preprint publication date, which roughly corresponds to the order in which the results where discovered. The presentation in this thesis follows a different order. The more general monadically dependent classes are explored first. Some results for the more restricted monadically stable classes can then be obtained as corollaries.

Chapter 4

Preliminaries

We write $\mathbb{N}$ for the set of natural numbers $\{0, 1, 2, \ldots\}$. For $m \in \mathbb{N}$ we let $[m] = \{1, \ldots, m\}$. We write $\bar{x}, \bar{y}, \ldots$ for tuples of variables and $\bar{a}, \bar{b}, \bar{v}, \bar{w}, \ldots$ for tuples of elements and usually leave it to the context to determine the length of a tuple. We access the elements of a tuple using subscripts, that is, $\bar{x} = x_1 x_2 \ldots x_{|\bar{x}|}$.

Sequences

To address and order the combinatorial objects of this thesis, we use *indexing sequences*. These are sequences (usually denoted by I, J) of elements without duplicates. We denote the sequence $(1, \ldots, n)$ also sometimes by $[n]$. We write $I \subseteq J$ if I is a subsequence of J. We use the usual comparison operators $<, >$ to indicate the order of elements within a sequence. Given a sequence I, and an element $i \in I$, we denote by $\text{pred}_I(i)$ and $\text{succ}_I(i)$ the predecessor and successor of i in I. Moreover, if $I = (a_1, \ldots, a_n)$, we define $\text{tail}(I) := (a_2, \ldots, a_n)$.

Graphs

All graphs in this thesis are simple and undirected. Unless a graph G is considered to be an input to an algorithm, we do not need to assume that G is a finite graph. The *length* of a path equals its number of edges. The *distance* between two vertex vertices a and b in a graph G, denoted by $\text{dist}_G(a, b)$, is the length of a shortest path with endpoints a and b and ∞ if no such path exists. The *distance* between two vertex sets A and B in a graph G is defined as

$$\text{dist}_G(A, B) := \min_{a \in A, b \in B} \text{dist}_G(a, b).$$

We drop the subscripts if the graph G is clear from the context. Two vertex sets A and B are *non-adjacent* if $\text{dist}(A, B) > 1$. A set or tuple is *distance-r independent* if all its vertices have pairwise distance greater than r. The *(closed) r-neighborhood* and *open r-neighborhood* of a vertex v are denoted by

$$N_r^G[u] := \{v \in V(G) : \text{dist}_G(\{u\}, \{v\}) \leqslant r\}, \quad \text{and} \quad N_r^G(u) := N_r^G[u] \setminus \{u\}.$$

More generally, for a tuple (or set) $\bar{a}$ we let $N_r^G[\bar{a}] = \bigcup_{a \in \bar{a}} N_r^G[a]$. Again we drop the graph from the notation if it is clear from the context, and we drop the subscript if $r = 1$. We also call $N_r[v]$ the *r-ball* around v. We call two distinct vertices u and v *twins* in a graph G if $N^G(u) \setminus \{v\} = N^G(v) \setminus \{u\}$. The *complement graph* of a graph G is denoted by $\overline{G}$. By default, graphs have no colors, but we speak of *colored graphs* when we allow vertex-colors. In this case we treat the colors as a fixed set of unary predicates which partition the vertex set, that is, each vertex has exactly one color. We call G^+ an *s-coloring of G*, if G^+ is obtained by coloring G with

s many colors. A *graph class* is a (usually infinite) set of graphs. For a set of vertices $X \subseteq V(G)$, we write $G[X]$ for the subgraph of G induced by X, and $G - X$ for the subgraph of G induced by $V(G) - X$. A graph class is *monotone* if it is closed under taking subgraphs. It is *hereditary* if it is closed under taking induced subgraphs. For a graph G and two disjoint subsets U and V of its vertices, we say the bipartite graph with vertices U and V and edges $E(G) \cap \big((U \times V) \cup (V \times U)\big)$ is a *semi-induced* subgraph of G.

We call the bipartite graph with sides $a_1, \ldots, a_\ell$ and $b_1, \ldots, b_\ell$

- 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 *half-graph* of order ℓ if a_i and b_j are adjacent if and only if $i \leqslant j$ for all $i, j \in [\ell]$.

The *powerset graph* of order ℓ is the bipartite graph with sides $\{a_S : S \subseteq [\ell]\}$ and $b_1, \ldots, b_\ell$ such that a_S and b_i are adjacent if and only if $i \in S$ for all $i \in [\ell]$ and $S \subseteq [\ell]$. See Figure 4.1 for examples.

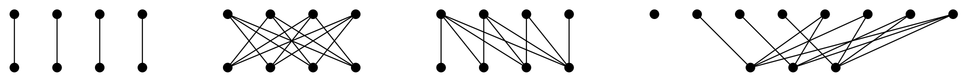

Figure 4.1: From left to right: the matching of order 4, the co-matching of order 4, the half-graph of order 4, and the powerset graph of order 3.

Flips

Fix a graph G and a partition $\mathcal{K}$ of its vertices. We will think of $\mathcal{K}$ as a coloring of the vertices of G. For every vertex $v \in V(G)$ we denote by $\mathcal{K}(v)$ the unique color $X \in \mathcal{K}$ satisfying $v \in X$. Let $F \subseteq \mathcal{K}^2$ be a symmetric relation. The *flip* $G \oplus F$ of G is defined as the (undirected) graph with vertex set $V(G)$, and edges defined by the following condition, for distinct $u, v \in V(G)$:

$$uv \in E(G \oplus F) \Leftrightarrow \begin{cases} uv \notin E(G) & \text{if } (\mathcal{K}(u), \mathcal{K}(v)) \in F, \\ uv \in E(G) & \text{otherwise.} \end{cases}$$

We call $G \oplus F$ a $\mathcal{K}$*-flip* of G. If $\mathcal{K}$ has at most k parts, we say that $G \oplus F$ a k*-flip* of G. A crucial property of flips is that they are reversible using first order-logic. We can recover the edges of the original graph in a coloring of its flip as follows. Let $H := G \oplus F$ and H^+ be the coloring of H where each part of $\mathcal{K}$ is assigned its own color. Define the symmetric binary formula

$$\varphi_{\mathcal{K},F}(x, y) := x \neq y \wedge \bigvee_{X,Y \in \mathcal{K}} x \in X \wedge y \in Y \wedge \big(E(x, y) \ \text{XOR} \ (X, Y) \in F\big).$$

We now have $G \models E(u, v) \Leftrightarrow H^+ \models \varphi_{\mathcal{K},F}(u, v)$.

Lemma 4.1. *Given a graph G and a k-flip H of G, we can compute in time $O(k \cdot |V(G)|^2)$ a partition $\mathcal{P}$ of $V(G)$ of size at most k and a symmetric relation $F \subseteq \mathcal{P}^2$ such that $H = G \oplus F$.*

Proof. Construct the graph $G_\oplus$ with vertices $V(G)$ and edges

$$uv \in E(G_\oplus) \Leftrightarrow \big(uv \in E(G) \ \text{XOR} \ uv \in E(H)\big).$$

This means two vertices are adjacent in $G_\oplus$ if and only if the adjacency between them was flipped in H. Let $\mathcal{P}$ be the partition of $V(G)$ where two vertices u and v are in the same part of $\mathcal{P}$ if and only if they are twins in $G_\oplus$ (i.e., if $N^{G_\oplus}(u) \setminus \{v\} = N^{G_\oplus}(v) \setminus \{u\}$).

Claim 4.2. $|\mathcal{P}| \leqslant k$.

Proof. Assume towards a contradiction that $|\mathcal{P}| > k$. Let $\mathcal{Q}$ be the partition of size at most k witnessing that H is a k-flip of G. By the pigeonhole principle, there are two distinct vertices u and v that are in distinct parts of $\mathcal{P}$ but in the same part of $\mathcal{Q}$. Being in different parts of $\mathcal{P}$, u and v are not twins in $G_\oplus$. By symmetry, we can therefore assume the existence of a third vertex w whose adjacency with u was flipped in H but whose adjacency with v was not flipped in H. A contradiction to u and v being in the same part of $\mathcal{Q}$. ■

By the above claim, $\mathcal{P}$ can be constructed in time $O(k \cdot |V(G)|^2)$: we iteratively classify the vertices into parts by comparing each vertex with one representative of each of the at most k parts created so far.

We observe that every two distinct parts $P, Q \in \mathcal{P}$ are either non-adjacent or fully adjacent in $G_\oplus$. Moreover, every single part $P \in \mathcal{P}$ either forms a clique or an independent set in $G_\oplus$. We define $F \subseteq \mathcal{P}^2$ as follows. For distinct parts $P, Q \in \mathcal{P}$ we set $(P, Q) \in F$ if and only if P and Q are fully adjacent in $G_\oplus$ and $(P, P) \in F$ if and only if P is a clique in $G_\oplus$. It is now easy to verify that $H = G \oplus F$. □

Lemma 4.3 (Transitivity of Flips). *Let G be a graph, H_1 be a k_1-flip of G and H_2 be a k_2-flip of H_1. Then H_2 is a $(k_1 \cdot k_2)$-flip of G.*

Proof. Let $\mathcal{P}_1, |\mathcal{P}_1| \leqslant k_1, F_1 \subseteq \mathcal{P}_1^2$ and $\mathcal{P}_2, |\mathcal{P}_2| \leqslant k_2, F_2 \subseteq \mathcal{P}_2^2$ be witnesses such that $H_1 = G \oplus F_1$ and $H_2 = H_1 \oplus F_2$. Let $\mathcal{P}$ be the common refinement of $\mathcal{P}_1$ and $\mathcal{P}_2$. Define the relation $F \subseteq \mathcal{P}^2$ as follows. For parts $P, Q \in \mathcal{P}$ let $P_1, Q_1 \in \mathcal{P}_1$ and $P_2, Q_2 \in \mathcal{P}_2$ be the unique parts containing P and Q. We set

$$(P, Q) \in F \Leftrightarrow \big((P_1, Q_1) \in F_1 \;\text{ XOR }\; (P_2, Q_2) \in F_2\big).$$

It is easy to verify that $H_2 = G \oplus F$. □

Logic

We use standard terminology from model theory and refer to [50] for extensive background. Every formula in this thesis will be a first-order formula over the signature of (possibly colored) graphs. We will often not explicitly write down the formulas if the properties they express are obviously expressible. For example, $x \in N_r[\bar{y}]$ stands for the first-order formula expressing that x is contained in the r-neighborhood of $\bar{y}$. Also, for a color predicate P, we often write $\exists x \in P\ \varphi$ as a shorthand for $\exists x\ P(x) \wedge \varphi$ and $\forall x \in P\ \varphi$ as a shorthand for $\forall x\ P(x) \to \varphi$. For a formula φ, we write $\text{free}(\varphi)$ for the set of free variables appearing in φ, and we write $\varphi(\bar{x})$ to indicate that the free variables of φ are in $\bar{x}$.

Every formula $\eta(x, y)$ on a graph G defines the relation $\eta(G) := \{(u, v) \in V(G)^2 : G \models \eta(u, v)\}$. Similarly, a formula $\nu(x)$ defines the set $\{v \in V(G) : G \models \varphi(v)\}$. We call a formula $\eta(x, y)$ symmetric and irreflexive if on all graphs the relation it defines is symmetric and irreflexive. Let $\alpha(x; y_1, \ldots, y_k)$ be a formula, with free variables partitioned into x and $y_1, \ldots, y_k$, as indicated by the semicolon. Given a graph G, vertices $v_1, \ldots, v_k$, and a set $U \subseteq V(G)$, we denote

$$\alpha(U; v_1, \ldots, v_k) := \{u \in U : G \models \alpha(u; v_1, \ldots, v_k)\}.$$

Let G^+ be a graph with colors $U_1, \ldots, U_l$. The *atomic type* of a tuple $\bar{v} = (v_1, \ldots, v_k)$ of vertices in G^+ is the quantifier-free formula $\alpha(x_1, \ldots, x_k)$ defined as the conjunction of all literals $\beta(x_1, \ldots, x_k)$ (that is, formulas $x_i = x_j$, $E(x_i, x_j)$, $U_1(x_i), \ldots, U_l(x_i)$, or their negations) such that $G^+ \models \beta(\bar{v})$. We write $\text{atp}_{G^+}(v_1, \ldots, v_k)$ to denote the atomic type of $\bar{v}$ in G^+.

Normalization

For every finite signature Σ, quantifier rank q, and tuple of free variables $\bar{x}$, up to equivalence there only exist a finite number of distinct formulas $\varphi(\bar{x})$ over Σ with quantifier rank at most q. Testing equivalence of first-order formulas is undecidable. However, given a formula we can compute an equivalent *normalized* formula of the same quantifier rank, such that again there only exist a finite number of distinct normalized formulas $\varphi(\bar{x})$ over Σ with quantifier rank at most q. In particular, the length of a normalized formula $\varphi(\bar{x})$ with quantifier rank q over Σ only depends on $|\bar{x}|$, q, and Σ. The normalization process works by renaming quantified variables, reordering boolean combinations into conjunctive normal form, and deleting duplicates from conjunctions and disjunctions. We will assume throughout this thesis that all appearing formulas are normalized. This also includes formulas which we construct ourselves: normalization is always performed implicitly as the last step of a construction.

Interpretations

The *interpretation* I is specified by two first-order formulas $\delta(x)$ and $\varphi(x, y)$, where φ is symmetric and irreflexive, and defines an operation that maps an input graph G to the output graph $I(G) := H$ such that:

$$V(H) := \{v \in V(G) \colon G \models \delta(v)\} \quad \text{and} \quad E(H) := \left\{uv \in V(H)^2 : G \models \varphi(u, v)\right\}.$$

We write $I_{\delta,\varphi}$ for the interpretation specified by $\delta(x)$ and $\varphi(x, y)$ and I_φ if $\delta(x) =$ "true" leaves the domain unchanged. We say a graph class $\mathcal{C}$ *interprets* a graph class $\mathcal{D}$, if there is an interpretation I such that for every graph $G \in \mathcal{D}$ there is a *preimage* $H \in \mathcal{C}$ with $I(H) = G$. Additionally, we say $\mathcal{C}$ *efficiently* interprets $\mathcal{D}$, if there is a polynomial time algorithm that given a graph $G \in \mathcal{C}$ computes a suitable preimage $H \in \mathcal{C}$. In particular, we require the size of H to be polynomial in the size of G.

First-order formulas can be naturally combined with interpretations. More precisely, given an interpretation $I := I_{\delta,\varphi}$ and a formula $\psi(\bar{x})$, we define $I(\psi)(\bar{x})$ to be the formula obtained by recursively rewriting ψ where we replace each

- atomic subformula $E(x, y)$ with $\varphi(x, y)$,
- existential quantification $\exists z : \alpha(\bar{x}, z)$ with $\exists z : \delta(z) \wedge \alpha(\bar{x}, z)$, and
- universal quantification $\forall z : \alpha(\bar{x}, z)$ with $\forall z : \delta(z) \rightarrow \alpha(\bar{x}, z)$.

We have the following standard fact.

Fact 4.4 (see, e.g., [50, Theorem 4.3.1]). *For every interpretation I, formula $\psi(\bar{x})$, graphs G and H satisfying $G = I(H)$, and tuple $\bar{a} \in V(G)^{|\bar{x}|}$,*

$$G \models \psi(\bar{a}) \qquad \textit{if and only if} \qquad H \models I(\psi)(\bar{a}).$$

We deduce two corollaries.

Corollary 4.5 (Transitivity). *Let $\mathcal{C}, \mathcal{D}, \mathcal{E}$ be classes of graphs such that $\mathcal{C}$ (efficiently) interprets $\mathcal{D}$ and $\mathcal{D}$ (efficiently) interprets $\mathcal{E}$. Then also $\mathcal{C}$ (efficiently) interprets $\mathcal{E}$.*

Corollary 4.6 (Reduction). *Let $\mathcal{C}$ be a graph class that efficiently interprets the class of all graphs. Then the first-order model checking problem is* AW$[*]$*-hard on $\mathcal{C}$.*

The second corollary additionally uses the fact that the first-order model checking problem is AW$[*]$-hard on the class of all graphs [22].

Transductions

A *transduction* T is specified by a first-order formula $\varphi(x, y)$ over the signature of k-colored graphs for some $k \in \mathbb{N}$ and defines an operation that maps an input graph G to the set of output graphs $T(G)$ that is the hereditary closure of the set $\bigcup\{I_\varphi(G^+) : G^+ \text{ is a } k\text{-coloring of } G\}$. This means each graph in $T(G)$ is obtained by

1. Coloring G. 2. Applying the interpretation I_φ. 3. Taking an induced subgraph.

Again we write T_φ for the transduction specified by φ. We say a class $\mathcal{C}$ *transduces* a class $\mathcal{D}$ if there is a transduction T such that $\mathcal{D} \subseteq \bigcup_{G \in \mathcal{C}} T(G)$. By refining colors and a small modification of Fact 4.4, this notion is transitive.

Stability and Dependence

We refer to the textbooks [3, 71, 79, 81, 80] for extensive background on classical stability theory. While we already defined monadic stability and dependence in terms of transductions in the introduction, let us also give the equivalent original definitions (restricted to graph classes). A formula $\varphi(\bar{x}, \bar{y})$ over the signature of (possibly colored) graphs has the *k-order property* on a class of (possibly colored) graphs $\mathcal{C}$ if there are $G \in \mathcal{C}$ and two sequences $(\bar{a}_i)_{i \in [k]}$, $(\bar{b}_j)_{j \in [k]}$ of tuples of vertices of G, such that for all $i, j \in [k]$

$$G \models \varphi(\bar{a}_i, \bar{b}_j) \quad \Leftrightarrow \quad i \leqslant j.$$

The formula φ has the *order property* on $\mathcal{C}$ if it has the k-order property on $\mathcal{C}$ for all $k \in \mathbb{N}$. The class $\mathcal{C}$ is *stable* if no formula has the order-property on $\mathcal{C}$. A graph class $\mathcal{C}$ is *monadically stable* if for every $\ell \in \mathbb{N}$ the class of all ℓ-colorings of graphs from $\mathcal{C}$ is stable.

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

$$G \models \varphi(\bar{a}, \bar{b}_J) \quad \Leftrightarrow \quad \bar{a} \in J.$$

We define the *independence property*, *dependent* classes, and *monadically dependent* classes as expected. Note that every (monadically) stable class is (monadically) dependent. Baldwin and Shelah proved that in the definitions of monadic stability and monadic dependence, one can alternatively rely on formulas $\varphi(x, y)$ with just a pair of singleton variables, instead of a pair of tuples of variables.

Fact 4.7 ([4, Lemma 8.1.3, Theorem 8.1.8])**.** *A graph class $\mathcal{C}$ is monadically stable (monadically dependent) if and only if for every $k \in \mathbb{N}$ every binary formula $\varphi(x, y)$ over the signature of k-colored graphs is stable (dependent) on the class of all k-colorings of graphs from $\mathcal{C}$.*

As a corollary we obtain the definition of monadic stability and dependence by transductions that was used in the introduction.

Corollary 4.8. *A graph class is monadically stable (monadically dependent) if and only if it does not transduce the class of all half-graphs (the class of all graphs).*

Locality of First-Order Logic

Two formulas $\varphi(\bar{x})$ and $\psi(\bar{x})$ are *equivalent* if for every graph G and every tuple $\bar{a} \in V(G)^{|\bar{x}|}$

$$G \models \varphi(\bar{a}) \Leftrightarrow G \models \psi(\bar{a}).$$

A first-order formula $\varphi(\bar{x})$ is *r-local* if every quantifier appearing in $\varphi(\bar{x})$ is of the form $\forall y \in N_r[\bar{x}]$ or $\exists y \in N_r[\bar{x}]$. Intuitively: whether the formula holds on a tuple, depends only on the r-neighborhood of the tuple. In particular, for every graph G and for all $\bar{a} \in V(G)^{|\bar{x}|}$

$$G \models \varphi(\bar{a}) \Leftrightarrow G\big[N_r[\bar{a}]\big] \models \varphi(\bar{a}).$$

A sentence ψ is *basic r-local* if there is $r' \leqslant r$ such that ψ is of the form

$$\exists x_1, \ldots, x_k \bigwedge_{1 \leqslant i < j \leqslant k} \mathrm{dist}(x_i, x_j) > 2r' \wedge \bigwedge_{1 \leqslant i \leqslant k} \varphi(x_i)$$

for an r-local formula φ.

Fact 4.9 (Gaifman's Locality Theorem [36]). *Every first-order formula $\varphi(\bar{x})$ of quantifier rank q is equivalent to a computable formula $\varphi^{\mathrm{loc}}(\bar{x})$ that is boolean combination of basic 7^q-local sentences and 7^q-local formulas.*

For an introduction to Gaifman's locality theorem see for example [46, Sec. 4.1]. We will need the following folklore consequence of Gaifman's locality theorem.

Lemma 4.10. *There exists a computable function $p : \mathbb{N}^2 \to \mathbb{N}$, such that for every quantifier rank q formula $\varphi(x, y)$ over the signature of k-colored graphs and every k-colored graph G, there exists a $p(q, k)$-coloring of G that for any two vertices $u, v \in V(G)$ with $\mathrm{dist}_G(u, v) > 2 \cdot 7^q + 1$, whether $G \models \varphi(u, v)$ depends only on the colors of u and v.*

We provide a proof for completeness. We first prove the following auxiliary statement.

Lemma 4.11. *For every r-local formula $\varphi(\bar{x})$ we can compute a formula $\varphi'(\bar{x})$ of the same quantifier rank that is boolean combination of r-local formulas with single free variables from $\bar{x}$, such that for every graph G and tuple $\bar{a} \in V(G)^{|x|}$ that is distance-$(2r+1)$ independent in G, we have*

$$G \models \varphi(\bar{a}) \Leftrightarrow G \models \varphi'(\bar{a}).$$

Proof. Every quantifier appearing in $\varphi(\bar{x})$ is of the form $\forall y \in N_r[\bar{x}]$ or $\exists y \in N_r[\bar{x}]$. Using the two equivalences

$$\forall y \in N_r[\bar{x}]\ \psi(\bar{x}, y) \equiv \bigwedge_{x \in \bar{x}} \forall y \in N_r[x]\ \psi(\bar{x}, y)$$

$$\exists y \in N_r[\bar{x}]\ \psi(\bar{x}, y) \equiv \bigvee_{x \in \bar{x}} \exists y \in N_r[x]\ \psi(\bar{x}, y)$$

we can assume that each quantifier is of the form $\forall y \in N_r[x]$ or $\exists y \in N_r[x]$ for some $x \in \bar{x}$. Up to renaming, we can assume each variable is quantified only once. Then, there is a function f that maps each variable y appearing in φ to a variable $f(y) \in \bar{x}$ such that $y \in N_r[f(y)]$ (we have $f(x) = x$ for each $x \in \bar{x}$). We inductively rewrite φ into a boolean combination φ' of formulas ψ such that for all $y_1, y_2 \in \mathrm{var}(\psi)$ we have $f(y_1) = f(y_2)$ where $\mathrm{var}(\psi)$ is the set of variables appearing in ψ. The rewriting replaces atoms $(y_1 = y_2)$ and $E(y_1, y_2)$ by "false" if $f(y_1) \neq f(y_2)$ and leaves them unmodified otherwise. This preserves the truth value of the atoms for any distance-$(2r+1)$ independent valuation $\bar{a}$ of $\bar{x}$: If $\bar{a}$ is distance-$(2r+1)$ independent then y_1 and y_2 are non-equal and non-adjacent in G, because

1. $f(y_1)$ and $f(y_2)$ have distance greater than $2r+1$ in G, and
2. $y_1 \in N_r[f(y_1)]$ and $y_2 \in N_r[f(y_2)]$.

The inductive step is trivial and it is easy to check that φ' has the desired properties. ☐

Proof of 4.10. Let $\Phi^t_{q,k}$ be the set of normalized formulas over the signature of k-colored graphs, with arity t and quantifier rank at most q. The size of $\Phi^t_{q,k}$ is bounded and computable. Therefore, there exists a computable function $f : \mathbb{N}^2 \to \mathbb{N}$ such that for every normalized formula $\varphi(x, y) \in \Phi^2_{q,k}$ the corresponding formula $\varphi^{\text{loc}}(x, y)$ obtained from Gaifman's locality theorem (Fact 4.9) has quantifier rank at most $f(q, k)$. We argue that we can set $p(q, k) := 2^m$ where $m := |\Phi^1_{f(q,k),k}|$.

Let $\varphi(x, y)$ and G be a formula and graph as in the statement of the lemma. We can assume φ is normalized. By Fact 4.9, $\varphi^{\text{loc}}(x, y)$ is a boolean combination of basic 7^q-local sentences and 7^q-local formulas, such that $\varphi^{\text{loc}}(x, y)$ is equivalent to $\varphi(x, y)$ and has quantifier rank at most $f(q, k)$. We obtain $\varphi'(x, y)$ from $\varphi^{\text{loc}}(x, y)$ by replacing in this boolean combination

- each basic local sentence ψ with the "true" or "false" atom, depending on whether $G \models \psi$,
- each local formula $\psi(x, y)$ with the boolean combination $\psi'(x, y)$ obtained by applying Lemma 4.11 to ψ.

Let $\Psi := \Phi^1_{f(q,k),k}$. Now $\varphi'(x, y)$ is a boolean combination of formulas from Ψ with single free variables. For all $u, v \in V(G)$ with $\text{dist}_G(u, v) > 2 \cdot 7^q + 1$, we have

$$G \models \varphi(u, v) \Leftrightarrow G \models \varphi'(u, v).$$

Let G^+ be the $p(q, k)$-coloring of G, where every vertex v is colored by the set $\text{col}(v) \subseteq \Psi$ that contains exactly those formulas $\alpha(x) \in \Psi$ with $G \models \alpha(v)$. It is now easy to see that this coloring has the desired properties: There is a relation $R \subseteq 2^\Psi \times 2^\Psi$ such that or all $u, v \in V(G)$ with $\text{dist}_G(u, v) > 2 \cdot 7^q + 1$, we have

$$G \models \varphi(u, v) \Leftrightarrow G \models \varphi'(u, v) \Leftrightarrow \big(\text{col}(u), \text{col}(v)\big) \in R. \qquad \square$$

Asymptotic Notation

As an analogue to the O-notation, we introduce two new notations that simplify our statements and proofs:

- $\text{const}(p_1, \dots, p_k)$ denotes a natural number, only depending on the parameters $p_1, \dots, p_k$.
- $U_{p_1,\dots,p_k}(n)$ denotes an anonymous function that is monotone, non-negative, and unbounded in n, and only depends on the parameters $p_1, \dots, p_k$.

More precisely, the ith occurrence of the notation $U_{p_1,\dots,p_k}(n)$ in the text should be interpreted as $f^i_{p_1,\dots,pk}(n)$, for some fixed unbounded function $f^i_{p_1,\dots,p_k} : \mathbb{N} \to \mathbb{N}$ that depends only on the parameters $p_1, \dots, p_k$ and i. We allow parameters of any kind, in particular, graph classes. We write $O_{p_1,\dots,p_k}(n)$ as a shorthand for $\text{const}(p_1, \dots, p_k) \cdot n$. To illustrate the use, we state Ramsey's theorem with this notation.

Example 4.12. For every k-coloring of the edges of the complete graph on vertex set $[n]$, there exists a set $X \subseteq [n]$ of size $U_k(n)$ which induces a monochromatic clique.

Ramsey Theory

For $\ell \in \mathbb{N}$ and a set I, let $I^{(\ell)}$ denote the set of subsets $J \subseteq I$ of size ℓ.

Fact 4.13 (Ramsey's Theorem, [74])**.** *For all $k, \ell, n \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for every*

$$c\colon [N]^{(\ell)} \to [k]$$

there is some $I \in [N]^{(n)}$ such that $c(J) = c(J')$ for all $J, J' \in I^{(\ell)}$.
Moreover, there is a function $f\colon \mathbb{N}^3 \to \mathbb{N}$ such that for every $k, \ell \in \mathbb{N}$ the function $m \mapsto f(k, \ell, m)$ is computable, monotone, and unbounded, and there is an algorithm that, given numbers $k, \ell, m \in \mathbb{N}$ and a function $c\colon [m]^{(\ell)} \to [k]$, computes in time $O_{k,\ell}(m^\ell)$ a set $I \subseteq [m]$ of size $f(k, \ell, m)$ such that $c(J) = c(J')$ for all $J, J' \in I^{(\ell)}$.

See, e.g., [74, Thm. C] or [44, Thm. 5.4] for a (standard) proof of Fact 4.13. The "moreover" part of the statement above follows by tracing the construction, which proceeds by induction on ℓ, where in each stage of the construction we iterate over some subset of $[m]$.

For a pair (a, b) of elements of a linearly ordered set $(A, \leqslant)$, let $\mathrm{otp}(a, b) \in \{<, =, >\}$ indicate whether $a < b$, $a = b$, or $a > b$ holds. For $\ell \geqslant 1$ and an ℓ-tuple of elements $a_1, \ldots, a_\ell$ of a linearly ordered set $(A, \leqslant)$, define the *order type* of $(a_1, \ldots, a_\ell)$, denoted $\mathrm{otp}(a_1, \ldots, a_\ell)$ as the tuple $(\mathrm{otp}(a_i, a_j))_{1 \leqslant i < j \leqslant \ell}$.

We now reformulate Ramsey's theorem using the U-notation, and also so that it talks about ℓ-tuples, rather than ℓ-element subsets.

Fact 4.14 (Reformulation of Ramsey's Theorem)**.** *For every k, ℓ, n and coloring*

$$c\colon [n]^\ell \to [k]$$

there is a subset $I \subseteq [n]$ of size $U_{k,\ell}(n)$ such that $c(a_1, \ldots, a_\ell)$ depends only on $\mathrm{otp}(a_1, \ldots, a_\ell)$, for all $(a_1, \ldots, a_\ell) \in I^\ell$.
Moreover, there is an algorithm that, given c, computes I in time $O_{k,\ell}(n^\ell)$.

The conclusion of the lemma means that there is a function f such that

$$c(a_1, \ldots, a_\ell) = f(\mathrm{otp}(a_1, \ldots, a_\ell)),$$

for all $(a_1, \ldots, a_\ell) \in I^\ell$.

Lemma 4.15 (Bipartite Ramsey Theorem)**.** *For every $k, \ell_1, \ell_2, n \in \mathbb{N}$ and coloring*

$$c\colon [n]^{\ell_1} \times [n]^{\ell_2} \to [k]$$

there are subsets $I_1, I_2 \subseteq [n]$ of size $U_{k,\ell_1,\ell_2}(n)$ such that $c(\bar{a}, \bar{b})$ depends only on $\mathrm{otp}(\bar{a})$ and $\mathrm{otp}(\bar{b})$, for all $\bar{a} \in I_1^{\ell_1}$ and $\bar{b} \in I_2^{\ell_2}$.

The conclusion of the lemma means that there is a function f such that

$$c(\bar{a}, \bar{b}) = f(\mathrm{otp}(\bar{a}), \mathrm{otp}(\bar{b})),$$

for all $\bar{a} \in I_1^{\ell_1}$ and $\bar{b} \in I_2^{\ell_2}$.

Proof. Set $\ell = \ell_1 + \ell_2$. The coloring c can be viewed as a coloring $c\colon [n]^\ell \to [k]$. By Lemma 4.14, there is a subset $I \subseteq [n]$ of size $U_{k,\ell}(n)$ such that the restriction of c to I is homogeneous, that is, $c(\bar{a})$ depends only on $\mathrm{otp}(\bar{a})$ for $\bar{a} \in I^\ell$. We can assume that $|I|$ is even. Let I_1 be the first $|I|/2$ elements of $I \subseteq [n]$, and I_2 be the remaining $|I|/2$ elements of I. Then $|I_1| = |I_2| \geqslant U_{k,\ell}(n)$. Let $(\bar{a}, \bar{b}), (\bar{a}', \bar{b}') \in I_1^{\ell_1} \times I_2^{\ell_2}$ be two pairs such that $\mathrm{otp}(\bar{a}) = \mathrm{otp}(\bar{a}')$ and $\mathrm{otp}(\bar{b}) = \mathrm{otp}(\bar{b}')$. Then $\mathrm{otp}(\bar{a}, \bar{b}) = \mathrm{otp}(\bar{a}', \bar{b}')$, since $\mathrm{otp}(a, b) = <$ for all $a \in I_1$ and $b \in I_2$. Therefore, $c(\bar{a}, \bar{b}) = c(\bar{a}', \bar{b}')$, by homogeneity of c restricted to I. □

Part II

Monadic Dependence

Outline Part II

In this part we prove various characterizations of monadically dependent graph classes as summarized in Theorem 2.3, which we restate here for convenience. The definitions of crossings and comparability grids can be found in Section 2.2.

Theorem 2.3. *Let $\mathcal{C}$ be a graph class. Then the following are equivalent.*

(1) $\mathcal{C}$ is monadically dependent.

(2) $\mathcal{C}$ is flip-breakable.

(3) For every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs

- *all flipped star r-crossings of order k, and*
- *all flipped clique r-crossings of order k, and*
- *all flipped half-graph r-crossings of order k, and*
- *the comparability grid of order k.*

(4) The hereditary closure of $\mathcal{C}$ does not efficiently interpret the class of all graphs.

We first prove the equivalence (1) ⇔ (2) ⇔ (3) in Chapters 5 and 6. The proofs are effective, giving us the following algorithmic version of flip-breakability.

Theorem 4.16. *For every monadically dependent graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$, there exists an unbounded function $f_r : \mathbb{N} \to \mathbb{N}$, a constant $k_r \in \mathbb{N}$, and an algorithm that, given a graph $G \in \mathcal{C}$ and $W \subseteq V(G)$, computes in time $O_{\mathcal{C},r}(|V(G)|^2)$ two subsets $A, B \subseteq W$ with $|A|, |B| \geqslant f_r(|W|)$ and a k_r-flip H of G such that:*

$$\mathrm{dist}_H(A, B) > r.$$

The equivalence with (4) is proven in Chapter 7. There we also obtain our hardness result for the first-order model checking problem.

Theorem 2.4. *The first-order model checking problem is* AW$[*]$*-hard on every hereditary, monadically independent graph class.*

Chapter 5

Flip-Breakability

In this chapter we define three tameness conditions for graph classes:

- *Prepattern-freeness*: the absence of certain combinatorial obstructions called *prepatterns*.
- *Insulation property*: the ability to guard any vertex set using so-called *insulators*.
- *Flip-breakability*: the ability to break any vertex set into two distant parts using few flips.

We show that for any graph class the following implications hold:

prepattern-free $\Rightarrow$ insulation-property $\Rightarrow$ flip-breakable $\Rightarrow$ mon. dependent

We later close the circle of implications in Chapter 6, where we show:

$\neg$prepattern-free $\Rightarrow$ large flipped crossings/comparability grids $\Rightarrow$ mon. independent

5.1 Grids and Insulators

Definition 5.1 (Grids)**.** Fix a non-empty sequence I and an integer $h \geqslant 1$. A *grid A indexed by I* and of *height h* in a graph G is a collection of pairwise disjoint sets $A[i, r] \subseteq V(G)$, for $i \in I$ and $r \in [h]$, called *cells*. Each grid is either tagged as *orderless* or *ordered*.

To facilitate notation, we often assume, up to renaming, the indexing sequence to be $I = (1, \dots, n)$. We do so in the following. For subsets $J \subseteq I$ and $R \subseteq [h]$, we write $A[J, R] = \bigcup_{i \in J, r \in R} A[i, r]$. We often use implicitly defined sets via wildcards and comparisons. For example $A[\leqslant i, *]$ stands for $A\big[\{1, \dots, i\}, [h]\big]$. In particular, we use $A[i, *] := \bigcup_{r \in [h]} A[i, r]$ and $A[*, r] := \bigcup_{i \in I} A[i, r]$, and refer to those sets as to *columns* and *rows* of A, respectively. In slight abuse of notation, we often write A instead of $A[*, *]$ to denote the set of all elements inside the grid. We define the *interior* of A as $\mathrm{int}(A) := A \setminus (A[1, *] \cup A[n, *] \cup A[*, h])$. It is intentional and will later be important that the cells $A[2, 1], \dots, A[n-1, 1]$ of the bottom row are part of the interior. Moreover, we say two columns $A[i, *]$ and $A[j, *]$ are *close*, if $|i - j| \leqslant 1$ and two rows $A[*, r]$ and $A[*, t]$ are *close*, if $|r - t| \leqslant 1$. Two cells are *close* if their respective columns and rows are close. Note that unlike standard matrix notation, to highlight the hierarchical relationship between columns and rows, our notation $A[i, r]$ first mentions the column index $i \in I$ and then the row index $r \in [h]$.

Insulators. The following notion of *insulators* serves a twofold purpose. On the one hand, insulators enforce the necessary structure to obtain *flip-breakability* (see Section 5.4). On the other hand, in Chapter 6, we use insulators to build the patterns presented in Chapter 2.

Definition 5.2 (Insulators)**.** An *insulator* $\mathcal{A} = (A, \mathcal{K}, F, R)$ *indexed by a sequence* I of *height* h and *cost* k in a graph G consists of

- a grid A indexed by I and of height h,
- a partition $\mathcal{K}$ of $V(G)$ into at most k color classes,
- a symmetric relation $F \subseteq \mathcal{K}^2$ specifying a flip $G' := G \oplus F$,
- a relation $R \subseteq \mathcal{K}^2$.

We furthermore say that $\mathcal{A}$ is *orderless* (*ordered*), if A is *orderless* (*ordered*).

If $\mathcal{A}$ is *orderless*, we demand:

(U.1) For all $i \in I$ there exists $a_i \in V(G)$ such that

$$A[i, 1] = \{a_i\} \quad \text{and} \quad A[i, \leqslant r] = N_{r-1}^{G'}[a_i] \text{ for all } r \in [h].$$

▶ In particular, a column of an orderless insulator is just the $(h-1)$-ball in the k-flip G' around a single vertex sitting in the bottom cell of that column.

If $\mathcal{A}$ is *ordered*, we demand:

(O.1) If two vertices are in different rows of A, then they have different colors in $\mathcal{K}$.

▶ This technical property ensures that the rows of the insulator are sufficiently distinguishable. We will use this to argue that certain modifications on the insulator can be performed without increasing its cost.

(O.2) Every vertex $v \in A[i, r]$ with $r > 1$, $i \in I$ has a neighbor in the cell $A[i, r-1]$ in G'.

▶ The mandatory downward edge, together with (O.4.2) and (O.5), ensures that each column is cohesive: we will later observe that the columns of the insulator are first-order definable. This property is crucial to obtain hardness results.

(O.3) For every $v \notin A[*, *]$ and $X \in \mathcal{K}$ we require that v is *homogeneous* to $X \cap \text{int}(A)$ in G (that is, either all or no vertices in $X \cap \text{int}(A)$ are adjacent to v).

▶ The inside of the insulator is "insulated" from its outside: the adjacencies between the two are described using only colors.

(O.4) For every $u \in A[i, r]$ with $r < h$, $i \in I$ and $v \in A$ we have the following:
(Up to renaming, we assume $I = (1, \ldots, n)$.)

(O.4.1) If u and v are in rows that are not close and $u \in \text{int}(A)$,
then they are non-adjacent in G'.

(O.4.2) If $v \in A[<i, r-1] \cup A[>i, r+1]$, then u and v are non-adjacent in G'.

(O.4.3) If $v \in A[>i+1, \{r, r-1\}]$, then $G \models E(u, v) \Leftrightarrow (\mathcal{K}(u), \mathcal{K}(v)) \in R$.

(O.4.4) If $v \in A[<i-1, \{r, r+1\}]$, then $G \models E(u, v) \Leftrightarrow (\mathcal{K}(v), \mathcal{K}(u)) \in R$.

Otherwise, we make no claims regarding the adjacency of u and v.

▶ Properties (O.4.1), (O.4.3), and (O.4.4) provide vertical and horizontal insulation inside the insulator. Property (O.4.2) helps to keep each column cohesive. See Figure 5.1 for an illustration.

(O.5) There exists a bound $r < h$ and a k-flip H of G such that for every two distinct vertices $u, v \in A[*, 1]$ we have $N_r^H[u] \cap N_r^H[v] = \emptyset$. Moreover, there are vertices $\{b(v) \in N_r^H[v] : v \in A[*, 1]\}$ and $\{c_i \in V(G) : i \in I\}$ and a symbol $\sim \in \{\leqslant, \geqslant\}$ such that for every $i, j \in I$, $v \in A[j, 1]$ and $G \models E(c_i, b(v))$ if and only if $i \sim j$.

▶ This property orders the columns. It will later be used to first-order define the columns as intervals in the order.

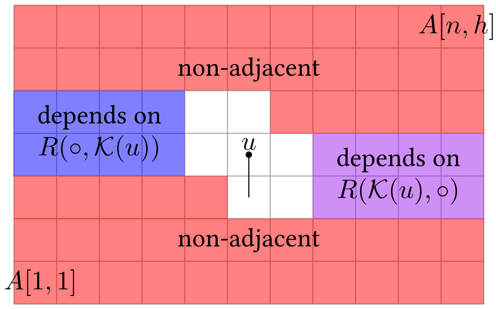

Figure 5.1: Illustration of how (O.4) controls the adjacency of a vertex u within the insulator.

Observation 5.3. Every orderless insulator $\mathcal{A} = (A, \mathcal{K}, F, R)$ where $F = R$ also satisfies the properties (O.2), (O.3), and (O.4) of an ordered insulator. Since the requirements of an orderless insulator put no restrictions on R, in the orderless case we can always assume $F = R$.

Two example insulators are depicted in Figure 5.2.

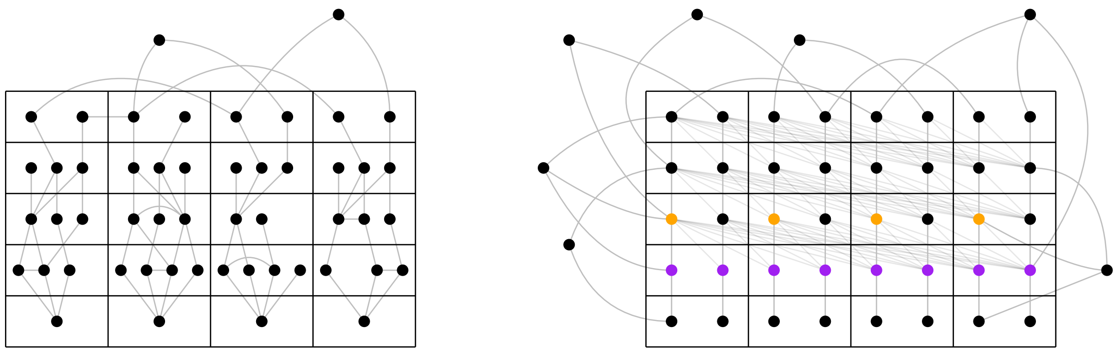

Figure 5.2: On the left: an orderless insulator. Each column is just the ball (in a flip) around the single vertex of its bottom cell. Apart from the boundary, there are no connections in between columns. On the right: an ordered insulator. It contains connections between columns, but they are well controlled by property (O.4). Property (O.5) is witnessed by the highlighted vertices. The purple $b(\cdot)$ vertices are contained in disjoint 1-balls around the vertices of the bottom cell. They are preordered by the yellow $c(\cdot)$ vertices.

We will often identify an insulator $\mathcal{A}$ and its underlying grid A, and write, for example $v \in \mathcal{A}$ to indicate that $v \in A[*,*]$. We start by observing basic properties of insulators. The following property will be crucial to obtain the model checking hardness result.

Lemma 5.4. *The columns of an insulator are definable in first-order logic in a coloring of G.*

More precisely: Let $\mathcal{A}$ be an insulator with grid A of cost k and height h indexed by I in a graph G. There exists a formula $\alpha(x, y)$ (depending only on k and h), a $\mathrm{const}(k, h)$*-coloring G^+ of G, and vertices $\{a_i : i \in I\}$ such that for each $i \in I$*

$$A[i,*] = \{v \in V(G) : G^+ \models \alpha(v, a_i)\}.$$

We do not rely directly on this lemma, as we will later further process insulators before proving hardness. However, we sketch a proof for instructive purposes.

Proof sketch for Lemma 5.4. All the formulas we write will work in a $\mathrm{const}(k, h)$-coloring G^+ of G, but we omit the details of the coloring to streamline the presentation. If $\mathcal{A}$ is orderless we use property (U.1). The a_i vertices are the singleton vertices in the bottom row of $\mathcal{A}$ and the formula $\alpha(x, y)$ is defined as

"x and y are at distance at most $h-1$ in the flip G' of G".

If $\mathcal{A}$ is ordered, we use (O.5). There is a formula $\beta(x,y)$ that defines $b(v)$ for each $v \in A[i,*]$:

"x is a b-vertex and contained in the r-ball around y in the flip H of G".

Building on β, there is a formula $\gamma(x,y)$ that, given $v \in A[j,*]$, defines $\{c_i : R(i,j), i \in I\}$:

"x is a c-vertex and adjacent to $b(y)$ in G".

As $R \in \{\leqslant, \geqslant\}$, we can now define a preorder $\prec$ on the vertices of the bottom row $A[*,1]$ which respects the column order by comparing their γ-neighborhoods:

$$x \prec y := \exists z : \gamma(z,x) \wedge \neg\gamma(z,y) \qquad \text{or} \qquad x \prec y := \exists z : \neg\gamma(z,x) \wedge \gamma(z,y)$$

depending on the choice of R. We choose an arbitrary vertex a_i of each cell $A[i,1]$ of the bottom row. Using the preorder, we can write a formula $\alpha_1(x,y)$ that defines $A[i,1]$ from a_i as its equivalence class in the preorder. For $1 < r \leqslant h$, we can now inductively write a formula $\alpha_r(x,y)$ that defines $A[i,\leqslant r]$ from a_i as all the vertices that are already in $A[i,<r]$, or adjacent to $A[i,<r]$ but not to $A[i-1,<r]$ in G'. Here we again use the preorder to define $A[i-1,<r]$ from a_i. The correctness for vertices inside $\mathcal{A}$ follows from the properties (O.2) and (O.4.2). Vertices from outside $\mathcal{A}$ can be marked in the coloring and ignored. Setting $\alpha := \alpha_h$ finishes the proof sketch. ☐

Definition 5.5 (Subgrids and subinsulators)**.** Let A be a grid indexed by a sequence J and of height h in a graph G. For every subsequence $I \subseteq J$ of length at least two we define the *subgrid* $A|_I$ as the grid indexed by tail(I) and of height h, containing the following cells. For all $i \in \text{tail}(I)$ and $r \in [h]$, depending on whether A is *orderless* or *ordered*, we respectively set

$$A|_I[i,r] := A[i,r] \quad \text{or} \quad A|_I[i,r] := \bigcup\{A[m,r] : m \in I \text{ and } \text{pred}_I(i) < m \leqslant i\}.$$

$A|_I$ is ordered (orderless) if and only if A is ordered (orderless). See Figure 5.3 for a depiction. For every insulator $(A|_I, \mathcal{K}, F, R)$ indexed by J, we moreover define the *subinsulator* $\mathcal{A}|_I := (A|_I, \mathcal{K}, F, R)$. The upcoming Lemma 5.9 will prove the validity of this definition.

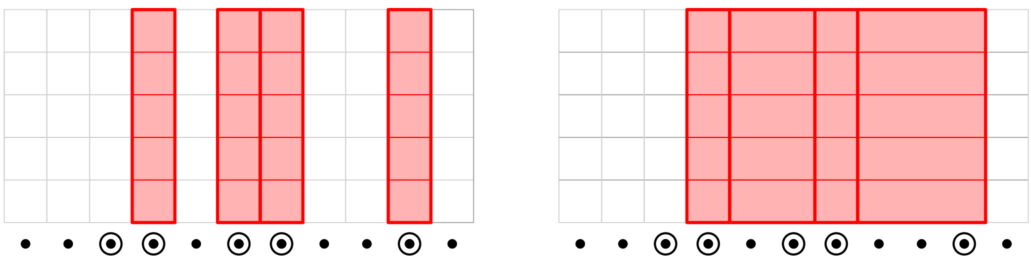

Figure 5.3: On the left/right: a subgrid of an orderless/ordered grid. The original grid A is depicted in gray. The dots at the bottom represent the sequence J indexing A. The subsequence $I \subseteq J$ is marked with circles. The subgrid $A|_I$ is overlaid in red. It is indexed by tail(I).

For an ordered insulator $\mathcal{A}$, in the definition of a subinsulator $\mathcal{A}|_I$, it is necessary that the indexing sequence tail(I) of $\mathcal{A}|_I$ excludes the first element of I: each $i \in I$ represents the interval $(\text{pred}_I(i), i]$ that is undefined for the first element of I. In orderless insulators this problem does not arise, but we choose to also exclude the first element of I to allow for uniform proofs which do not distinguish between the two. The following observation about subgrids is crucial to their definition and will later be used to build well-behaved subgrids using Ramsey-arguments.

Lemma 5.6. *Let A be a grid indexed by J and of height h in a graph G. For every subsequence $I \subseteq J$ of length at least two, in the subgrid $A|_I$, the content of the column $A|_I[i,*]$ depends only on A, i and $\text{pred}_I(i)$, instead of the whole sequence I. More precisely, there exists a function $M_A : J \times J \to 2^{V(G)}$ such that for every $I \subseteq J$ and $i \in \text{tail}(I)$ we have*

$$A|_I[i,*] = M_A(\text{pred}_I(i), i).$$

Proof. We define the function $\mu_A : J \times J \times [h] \to 2^{V(G)}$ as follows. For all $i < j \in J$ and every $r \in [h]$, depending on whether A is *orderless* or *ordered*, we respectively set

$$\mu_A(i,j,r) := A[j,r] \quad \text{or} \quad \mu_A(i,j,r) := \bigcup\{A[m,r] : m \in I \text{ and } i < m \leqslant j\}.$$

We can now define for all $i < j \in J$

$$M_A(i,j) := \bigcup_{r\in[h]} \mu_A(i,j,r).$$

□

Observation 5.7 (Transitivity). Let A be a grid indexed by I_0, I_1 be a subsequence of I_0, and I_2 be a subsequence of $\text{tail}(I_1)$. If $B = A|_{I_1}$ and $C = B|_{I_2}$ then $C = A|_{I_2}$.

Observation 5.8 (Monotonicity and Coverability). Let A be a grid indexed by J of height h and I be a subsequence of J. For all $i \in \text{tail}(I)$ and $r \in [h]$ we have

$$A[i,r] \ \subseteq\ A|_I[i,r] \ \subseteq\ \bigcup\{A[m,r] : m \in I \text{ and } \text{pred}_I(i) < m \leqslant i\}.$$

We finally show that taking a subinsulator preserves its good properties without increasing its cost.

Lemma 5.9. *Let $\mathcal{A} = (A, \mathcal{K}, F, R)$ be an insulator indexed by J on a graph G. For every subsequence $I \subseteq J$ also $\mathcal{A}|_I := (A|_I, \mathcal{K}, F, R)$ is an insulator on G.*

Proof. Up to renaming, we assume $J = (1,\ldots,n)$. Let $B := A|_I$ and $G' := G \oplus F$. In the orderless case, as the graph G' remains the same and B is obtained from A by just dropping columns, it is easy to see that (U.1) still holds. It remains to check the ordered case.

- To prove (O.1) and (O.4), we observe that for every $i \in I$ and $r \in [h]$

  $$B[*,r] \subseteq A[*,r] \quad \text{and} \quad B[<i,r] \subseteq A[<i,r] \quad \text{and} \quad B[>i,r] \subseteq A[>i,r].$$

  Property (O.1) follows directly. To prove, for example, (O.4.4), assume $u \in B[i,r]$ and $v \in B[<i-1,\{r,r+1\}]$. Then also $u \in B[> \text{pred}_I(i), r]$ and $v \in B[< \text{pred}_I(i), \{r,r+1\}]$. As argued above, $u \in A[> \text{pred}_I(i), r]$ and $v \in A[< \text{pred}_I(i), \{r,r+1\}]$. By property (O.4.4) of $\mathcal{A}$, we have $G \models E(u,v) \Leftrightarrow (\mathcal{K}(v), \mathcal{K}(u)) \in R$, as desired. The remaining statements of (O.4) follow similarly.
- To prove (O.2), let $u \in B[i,r]$ for $r > 1$ and let us show that u has a neighbor in $B[i,r-1]$ in G. By construction, we have $u \in A[i',r]$ for some $\text{pred}_I(i) < i' \leqslant i$. By (O.2) of $\mathcal{A}$, u has a neighbor v in $A[i',r-1]$ in G'. Again by construction, $A[i',r-1] \subseteq B[i,r-1]$, so also $v \in B[i,r-1]$.
- For (O.3) to hold we must check, for every $u \notin B$ and $X \in \mathcal{K}$, that u is homogeneous to $X_B := X \cap \text{int}(B)$ in G. As G' is a $\mathcal{K}$-flip of G, we can check the property in G' instead. If $u \notin A$ this holds as $X_B \subseteq \text{int}(A)$ and (O.3) was already true in $\mathcal{A}$.
  Assume now $u \in A[i,r] \setminus B$ for some $i \in J$ and $r \in [h]$. As we already established (O.1), we know that all vertices from X_B are in the same row $B[*,r'] \subseteq A[*,r']$ for some $r' \in [h]$.

If $|r - r'| > 1$ then u and X_B are non-adjacent in G' as $\mathcal{A}$ satisfies (O.4). Now assume $|r - r'| \leqslant 1$. Let $i_0, i_1, \ldots, i_n$ be the continuous subsequence of J where i_0 and i_n are the first and last elements of I. By construction, we have

$$B[*, *] = A[\{i_1, \ldots, i_n\}, *] \quad \text{and} \quad \text{int}(B) = A[\{i_2, \ldots, i_{n-1}\}, {<}h].$$

Since $u \notin B$ we have $i < i_1$ or $i_n < i$. Assume $i < i_1$. Then we have $X_B \subseteq A[{\leqslant}i_1, r'] \subseteq A[{>}i, r']$, and we can again use (O.4) of $\mathcal{A}$. Now if $r' = r + 1$ then u is non-adjacent to X_B in G'. If $r' \in \{r, r-1\}$ then u is adjacent to all of X_B if $(\mathcal{K}(u), X) \in R$ and non-adjacent to all of X_B otherwise. In each case u is homogeneous to X_B. The case where $i_n < j$ follows by a symmetric argument.

- The property (O.5) of $\mathcal{A}$ is witnessed by
  - a symbol $\sim \in \{\leqslant, \geqslant\}$,
  - a k-flip H, and
  - vertices $\{b(v) \in V(G) : v \in A[*, 1]\}$ and $\{c_i \in V(G) : i \in J\}$.

  To witness (O.5) of $\mathcal{B}$, we use $\sim$, H, $\{b(v) \in V(G) : v \in B[*, 1]\}$, and
  - $\{c_{i'} : i' = \text{succ}_J(\text{pred}_I(i)), i \in \text{tail}(I)\}$ if $(\sim) = (\leqslant)$, or
  - $\{c_i : i \in \text{tail}(I)\}$ if $(\sim) = (\geqslant)$. □

5.2 Prepatterns

Throughout the following sections, we will either make progress constructing large insulators, or will obtain the following kind of preliminary patterns, which are then processed further in Chapter 6.

Definition 5.10. Let $\mathcal{A}$ be an insulator indexed by a sequence K with grid A in a graph G. Say that G contains a *bi-prepattern* of order t on $\mathcal{A}$ if there exist

- index sequences $I, J \subseteq K$ of size t and with $|I| = |J| = t$,
- vertices $c_{i,j} \in V(G)$ for all $i \in I, j \in J$,
- quantifier-free formulas $\alpha_1(x; y, s_1)$ and $\alpha_2(x; y, s_2)$ with parameters $s_1, s_2 \in V(G)$,
- symbols $\sim_1, \sim_2 \in \{=, \neq\}$,

such that for all $i \in I, j \in J$,

$$i = \min\big\{i' \in I : \alpha_1(A[i', *]; c_{i,j}, s_1) \sim_1 \emptyset\big\},$$
$$j = \min\big\{j' \in J : \alpha_2(A[j', *]; c_{i,j}, s_2) \sim_2 \emptyset\big\}.$$

Let us give some intuition for the above definition. The bi-prepattern consists of two sequences of columns indexed by I and J. For each pair of columns $(i, j) \in I \times J$, there exists a private vertex $c_{i,j}$ pairing them up in the following sense. Column i is the first column in I, in which

- $c_{i,j}$ has no α_1-neighbor (if $\sim_1$ is $=$), or
- $c_{i,j}$ has an α_1-neighbor (if $\sim_1$ is $\neq$).

Similarly, j is the first column in J in which has an α_2-neighbor (resp. no α_2-neighbor). Figure 5.4 (*left*) illustrates this column pairing.

This pairing aspect will be used later to argue that large bi-prepatterns witness monadic independence (or equivalently: they are obstructions for monadic dependence). By Lemma 5.4, each column of the insulator can be first-order defined from a single representative vertex. Bi-prepatterns therefore logically resemble 1-subdivided bicliques, where the column representatives

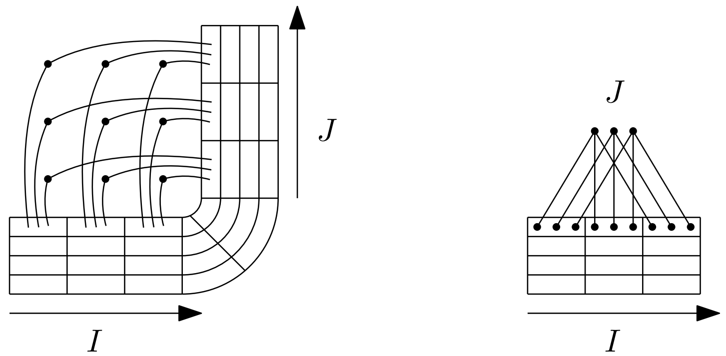

Figure 5.4: Schematic depiction of a bi-prepattern (left) and mono-prepattern (right).

are the principal vertices and the vertices $c_{i,j}$ form the subdivision vertices. Subdivided bicliques of unbounded size are witnesses for monadic independence: they encode arbitrary bipartite graphs by coloring the subdivision vertices. Therefore, bi-prepatterns may be thought of as obstructions to monadic dependence. Moreover, by the strong structure properties of insulators and since α_1 and α_2 are quantifier-free, we can use bi-prepatterns to later extract from them our concrete forbidden induced subgraph characterization.

In addition to the bi-prepatterns, our analysis produces a second kind of obstruction called a *mono-prepattern.*

Definition 5.11. Let $\mathcal{A}$ be an insulator indexed by a sequence K with grid A in a graph G. Say that G contains a *mono-prepattern* of order t on $\mathcal{A}$ if there exist

- index sequences I, J with $I \subseteq K$ and $|I| = |J| = t$,
- vertices $c_j \in V(G)$ for all $j \in J$,
- vertices $b_{i,j} \in A[i,*]$ for all $i \in I, j \in J$,
- a symbol $\sim \in \{=, \neq, \leqslant, <, \geqslant, >\}$,

such that for all $i \in I$ and $j, j' \in J$,

$$G \models E(c_j, b_{i,j'}) \Leftrightarrow j \sim j'.$$

As in a bi-prepattern, in a mono-prepattern the c_j vertices can be used to pair up columns $(i, j) \in I \times I$. While the bi-prepattern logically resembles a subdivided biclique (we pair elements from two sequences I and J), the mono-prepattern corresponds to a subdivided clique (we pair elements from the same sequence I). Figure 5.4 (*right*) illustrates mono-prepatterns.

Definition 5.12. G contains a *prepattern* of order t on an insulator $\mathcal{A}$ if it either contains a bi-prepattern of order t on $\mathcal{A}$ or a mono-prepattern of order t on $\mathcal{A}$.

Definition 5.13. A graph class $\mathcal{C}$ is *prepattern-free*, if for every $k, r \in \mathbb{N}$, there exists $t \in \mathbb{N}$ such that every graph $G \in \mathcal{C}$ does not contain prepatterns of order t on insulators of cost at most k and height at most r in G.

We later show in Proposition 6.47 that these obstructions are exhaustive: prepattern-freeness is equivalent to monadic dependence. In the following sections, we start by deriving structural properties of prepattern-free classes.

5.3 The Insulation Property

Towards proving that prepattern-free classes are flip-breakable, we first show that they have a more fine-grained structure property that we call the *insulation property.*

Definition 5.14. Let $\mathcal{A}$ be an insulator with grid A indexed by a sequence I in a graph G and let $W \subseteq V(G)$. We say that $\mathcal{A}$ *insulates* W if there is a bijection $f : W \to I$, such that for all $v \in W$

$$v \in A[f(v), 1].$$

A set W is (r, k)*-insulated* in G if there is an insulator $\mathcal{A}$ of height r and cost k that insulates W.

Definition 5.15. A graph class $\mathcal{C}$ has the *insulation property*, if for every radius $r \in \mathbb{N}$ there exist a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $k_r \in \mathbb{N}$ such that for every $m \in \mathbb{N}$, $G \in \mathcal{C}$, $W \subseteq V(G)$ with $|W| \geqslant N_r(m)$, there is a subset $W_\star \subseteq W$ of size at least m that is (r, k_r)-insulated in G.

More briefly: $\mathcal{C}$ has the insulation property if for every $r \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ there is a subset $W_\star \subseteq W$ of size $U_{\mathcal{C},r}(|W|)$ that is $(r, \mathrm{const}(\mathcal{C}, r))$-insulated in G.

The goal of this section is to prove the following.

Proposition 5.16. *Every prepattern-free graph class has the insulation property.*

We first notice that insulation for radius $r = 1$ is trivial.

Lemma 5.17. *Fix a graph G. Every set $W \subseteq V(G)$ is $(1, 1)$-insulated in G.*

Proof. Fix any enumeration $(a_1, \ldots, a_n)$ of W. Let A be the orderless grid indexed by the sequence $I := (1, \ldots, n)$ and of height 1 such that $A[i, 1] := \{a_i\}$ for all $i \in I$. Now $\mathcal{A} := (A, \{V(G)\}, \emptyset, \emptyset)$ is an orderless insulator with cost and height 1 that insulates W. □

In order to insulate sets with higher radii, we will grow insulators in height.

Definition 5.18 (Row-Extensions)**.** Let A be a grid indexed by I and of height h. We say a grid B is a *row-extension* of A, if it satisfies the following properties.

- B is indexed by I and has height $h + 1$.
- For all $i \in I$ and $j \in [h]$ we have $B[i, j] = A[i, j]$.
- Either A and B are both *orderless* or both *ordered*.

Similarly, we say an insulator $\mathcal{B}$ is a *row-extension* of an insulator $\mathcal{A}$, if the grid of $\mathcal{B}$ is a row-extension of the grid of $\mathcal{A}$.

Depending on whether the insulator at hand is orderless or ordered, we will create row-extensions using one of the following two insulator growing lemmas. To keep the presentation streamlined, the (quite technical) proofs of the two lemmas are deferred to Section 5.7.

Lemma 5.19 (Orderless Insulator Growing)**.** *Fix $k, t \in \mathbb{N}$. For every graph G and orderless insulator $\mathcal{A}$ indexed by J of cost k in G, we can compute a subsequence $I \subseteq J$ of length $U_t(|J|)$ such that either*

- *G contains a prepattern of order t on $\mathcal{A}|_I$,*
- *$\mathcal{A}|_I$ is orderable, or*
- *there exists a row-extension of $\mathcal{A}|_I$ of cost* $\mathrm{const}(k, t)$ *in G.*

Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes the sequence I and one of the three outcomes (a prepattern, witnesses for $\mathcal{A}|_I$ being orderable, or a row-extension) in time $O_{k,t}(|V(G)|^2)$.

The definition of an *orderable* insulator will be given shortly after the following lemma.

Lemma 5.20 (Ordered Insulator Growing). *Fix $k, t \in \mathbb{N}$. For every graph G and ordered insulator $\mathcal{A}$ with cost k, indexed by J in G, we can compute a subsequence $I \subseteq J$ of length $U_t(|J|)$ such that either*

- *G contains a prepattern of order t on $\mathcal{A}|_I$, or*
- *G contains a row-extension of $\mathcal{A}|_I$ with cost $\mathrm{const}(k,t)$.*

Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes the sequence I and one of the two outcomes (a prepattern or a row-extension) in time $O_{k,t}(|V(G)|^2)$.

Both insulator growing lemmas follow the same scheme. Given an insulator we find a large subinsulator that either witnesses a lot of non-structure (in form of a prepattern), or improves the structural guarantees of the original insulator (in form of a row-extension). In the orderless case a third outcome may appear: the subinsulator may be *orderable*. Orderable insulators are orderless insulators which can be converted into ordered ones, as made precise by the following definition and lemma.

Definition 5.21 (Orderable Insulators). Let $\mathcal{A}$ be an orderless insulator with grid A indexed by I in a graph G. We say that $\mathcal{A}$ is *orderable* if there exist vertices $\{b_i \in A[i,*] : i \in I\}$ and $\{c_i \in V(G) : i \in I\}$ and a symbol $\sim\, \in \{\leqslant, \geqslant\}$ such that for all $i, j \in I$

$$G \models E(c_i, b_j) \quad \text{if and only if} \quad i \sim j.$$

Lemma 5.22. *Let $\mathcal{A}$ be an orderless insulator of cost k and height h in a graph G that insulates a set $W \subseteq V(G)$. If $\mathcal{A}$ is orderable, then there exists an ordered insulator $\mathcal{B}$ of cost $k \cdot h$ and height h that also insulates W.*

Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes $\mathcal{B}$ in time $O_{k,h}(|V(G)|)$.

Proof. Let $\mathcal{A} = (A, \mathcal{K}, F, F)$ be as in the statement. Let B be the grid obtained by changing the tag of the grid A from *orderless* to *ordered*. Towards ensuring (O.1), let $\mathcal{K}_\star$ be the size $k \cdot h$ coloring obtained by taking the common refinement of the k-coloring $\mathcal{K}$ and an h-coloring that assigns vertices from different rows in A different colors. Let $F_\star \subseteq \mathcal{K}_\star^2$ be the corresponding refinement of F such that $G \oplus F = G \oplus F_\star$. We check that $\mathcal{B} := (B, \mathcal{K}_\star, F_\star, F_\star)$ is the desired ordered insulator. The insulator property (O.1) holds by construction. By Observation 5.3, (O.2), (O.3), and (O.4) carry over from $\mathcal{A}$. Finally, (U.1) ensures for the k-flip $H := G \oplus F$ and $v \in A[i,1]$ that $N^H_{h-1}[v] = A[i,*]$. As $\mathcal{A}$ moreover is orderable (cf. Definition 5.21), property (O.5) follows. The bound on the running time is trivial. □

As we know how to grow orderless insulators, turn orderless, orderable insulators into ordered ones, and grow ordered insulators, we can now grow arbitrarily high insulators. This yields a proof of Proposition 5.16, which we restate below with an added algorithmic conclusion.

Proposition 5.23. *Every prepattern-free graph class $\mathcal{C}$ has the insulation property.*

Moreover, there is an algorithm that, given a radius r, a graph $G \in \mathcal{C}$ and a set W, computes the subset $W_\star \subseteq W$ and a witnessing insulator in time $O_{\mathcal{C},r}(|V(G)|^2)$.

Proof. Fix a prepattern-free class $\mathcal{C}$ and $r \geqslant 1$. We prove that for every $G \in \mathcal{C}$ and $W \subseteq V(G)$ there is a subset $W_r \subseteq W$ of size $U_{\mathcal{C},r}(|W|)$ that is (r, k_r)-insulated in G, for some $k_r \leqslant \mathrm{const}(\mathcal{C}, r)$. This statement is proved by induction on $r \geqslant 1$. The base case of $r = 1$ follows immediately from Lemma 5.17.

In the inductive step, assume the statement holds for some $r \geqslant 1$; we prove it for $r+1$. Let $k_r \leqslant \mathrm{const}(\mathcal{C}, r)$ be the value obtained by inductive assumption. As $\mathcal{C}$ is prepattern-free, there is some number $t = \text{const}(\mathcal{C}, r)$ such that no graph $G \in \mathcal{C}$ contains a pattern of order t on insulators of cost at most $k_r \cdot r$ and height at most r in G.

Let $G \in \mathcal{C}$ and $W \subseteq V(G)$. By the inductive assumption, there is a subset $W_r \subseteq W$ of size $U_{\mathcal{C},r}(|W|)$ and an insulator $\mathcal{A}_r$ of height r and cost $\text{const}(\mathcal{C}, r)$ which insulates W_r. Let J denote the indexing sequence of $\mathcal{A}_r$. We prove that there is a subset $W_{r+1} \subseteq W_r$ of size $U_{\mathcal{C},r}(|W_r|) \geqslant U_{\mathcal{C},r}(|W|)$ and an insulator $\mathcal{A}_{r+1}$ of height $r+1$ and cost $\text{const}(\mathcal{C}, r)$ that insulates W_{r+1}. We consider two cases, depending on whether $\mathcal{A}_r$ is orderless or ordered.

Assume first that $\mathcal{A}_r$ is orderless. We apply Lemma 5.19 to $\mathcal{A}_r$. This yields a sequence $I \subseteq J$ of length $U_t(|J|)$ such that one of the following three cases applies.

1. G contains a prepattern of order t on $\mathcal{A}_r|_I$.
2. G contains a row-extension $\mathcal{A}_r|_I$ with cost $\text{const}(k_r, t) \leqslant \text{const}(\mathcal{C}, r)$.
3. $\mathcal{A}_r|_I$ is orderable.

We set $W' := A[*, 1]$ where A is the grid of $\mathcal{A}_r|_I$. By the definition of an orderless subgrid we have $W' \subseteq W_r \subseteq W$ and

$$|W'| = |\text{tail}(I)| = |I| - 1 \geqslant U_t(|J|).$$

As $\mathcal{A}_r$ is of height r and cost k_r, the same holds for $\mathcal{A}_r|_I$. By our choice of t, the first case cannot apply. In the second case, the row extension of $\mathcal{A}_r|_I$ witnesses that W' is $(r+1, \text{const}(\mathcal{C}, r))$-insulated, and we conclude by setting $W_{r+1} := W'$. It remains to handle the third case. We apply Lemma 5.22 to the orderable insulator $\mathcal{A}_r|_I$, yielding an ordered insulator $\mathcal{B}$ of cost $k_r \cdot r$ and height r that also insulates W'. Up to renaming, we can assume that $\mathcal{B}$ is indexed by $\text{tail}(I)$, the same sequence that also indexes $\mathcal{A}_r|_I$. We apply Lemma 5.20 to $\mathcal{B}$ with $k := k_r \cdot r$ and t. This yields a sequence $K \subseteq \text{tail}(I)$ with $|K| = U_t(|J|)$ such that one of the following two cases applies.

1. G contains a prepattern of order t on $\mathcal{B}|_K$.
2. G contains a row-extension of $\mathcal{B}|_K$ with cost $\text{const}(k, t) \leqslant \text{const}(\mathcal{C}, r)$.

By our choice of t, the first case cannot apply. In the second case, by definition of an ordered subgrid, there exists a set $W_{r+1} \subseteq W' \subseteq W$ that is insulated by a row extension of $\mathcal{B}|_K$ and has size at least

$$|W_{r+1}| = |\text{tail}(K)| = |K| - 1 \geqslant U_t(|J|).$$

Now W_{r+1} is the desired $(r+1, k)$-insulated set. This concludes the case where $\mathcal{A}_r$ is orderless. If $\mathcal{A}_r$ is ordered we can directly apply Lemma 5.20, which just improves the bounds. This completes the inductive proof.

The induction can be easily turned into an algorithm. The trivial insulator from Lemma 5.17 can be computed in time $O(|V(G)|)$. The running time of the inductive step follows from the running times of Lemma 5.19, Lemma 5.22, and Lemma 5.20. □

5.4 Insulation Property Implies Flip-Breakability

We recall the definition of flip-breakability.

Definition 2.10 (Flip-Breakability). A graph class $\mathcal{C}$ is *flip-breakable* if for every radius $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $k_r \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N_r(m)$ there exist subsets $A, B \subseteq W$ with $|A|, |B| \geqslant m$ and a k_r-flip H of G such that:

$$\text{dist}_H(A, B) > r.$$

In this section we prove the following.

Proposition 5.24. *Every graph class $\mathcal{C}$ with the insulation property is flip-breakable.*

Moreover, if $\mathcal{C}$ is also prepattern-free, then there is an algorithm that, given a radius r, graph G, and set W, computes in time $O_{\mathcal{C},r}(|V(G)|^2)$ subsets $A, B \subseteq W$, partition $\mathcal{K}_\star$, and relation $F_\star \subseteq \mathcal{K}_\star^2$ witnessing flip-breakability. (This means $|A|, |B| \geqslant U_{\mathcal{C},r}(|W|)$, $|\mathcal{K}_\star| \leqslant \mathrm{const}(\mathcal{C}, r)$, and $\mathrm{dist}_H(A, B) > r$ in the flip H defined by $\mathcal{K}_\star$ and $F_\star$.)

Let W be a set of vertices in a graph G. We call W *(r, k)-flip-breakable*, if there exist two disjoint sets $A, B \subseteq W$, each of size at least $\frac{1}{3}|W|$, and a k-flip H of G, such that

$$N_r^H(A) \cap N_r^H(B) = \emptyset.$$

Proposition 5.24 will be implied by the following.

Lemma 5.25. *Let W be a $(2r+1, k)$-insulated set of at least $18r$ vertices in a graph G. Then W is $(r, 2^{48r^2k})$-flip-breakable in G.*

Moreover, there is an algorithm that, given a radius r, graph G, set W, and an insulator witnessing that W is $(2r+1, k)$-insulated, computes in time $O_{r,k}(|V(G)|)$ the subsets $A, B \subseteq W$, partition $\mathcal{K}_\star$, and relation $F_\star \subseteq \mathcal{K}_\star^2$ that witness the $(r, 2^{48r^2k})$-flip-breakability of W.

Proof. Let $\mathcal{A} = (A, \mathcal{K}, F, R)$ be the insulator of cost k and height $2r + 1$ that insulates W. Since we will only be using the insulator properties (O.3) and (O.4) and by Observation 5.3, we do not have to distinguish whether $\mathcal{A}$ is orderless or ordered. Up to renaming, we can assume that A is indexed by a sequence $I = (1, \ldots, n)$ for some $n \in \mathbb{N}$. Let $l_1, l_2, r_1, r_2 \in I$ be the indices such that $(l_1, l_1 + 1, \ldots, l_2)$ and $(r_1, r_1 + 1, \ldots, r_2)$ are the sequences containing the first and last $2r$ elements of I respectively. Since $\mathcal{A}$ insulates W, we can choose indices $m_1 > \ell_2$ and $m_2 < r_1$ such that $(m_1, m_1 + 1, \ldots, m_2)$ contains $2r$ elements and there are two disjoint sets $W_1 := W \cap M_1 \cap A[*, 1]$ and $W_2 := W \cap M_2 \cap A[*, 1]$, where

$$M_1 := \bigcup_{l_2 < i < m_1} A[i, *] \quad \text{and} \quad M_2 := \bigcup_{m_2 < i < r_1} A[i, *],$$

such that W_1 and W_2 each contain at least $\frac{1}{2}(|W| - 6r) \geqslant \frac{1}{3}|W|$ vertices. The sets W_1 and W_2 will play the role of the sets A and B in the flip-breakability statement. We define a new grid B of height h indexed by $J := (1, \ldots, 3 \cdot 2r + 2)$ on the same vertex set as A. The rows of B are the same as the rows of A. The columns of B are in order

$$\underbrace{A[l_1, *],\ \ldots,\ A[l_2, *]}_{2r \text{ columns}},\ M_1,\ \underbrace{A[m_1, *],\ \ldots,\ A[m_2, *]}_{2r \text{ columns}},\ M_2,\ \underbrace{A[r_1, *],\ \ldots,\ A[r_2, *]}_{2r \text{ columns}}.$$

See Figure 5.5 for a visualization. We observe that B coarsens the columns of A in the following sense.

Observation 5.26. Let $u \in B[j, t]$ and $v \in B[j', t']$ such that $j < j'$ for some $j, j' \in J$ and $t, t' \in [h]$. Then also $u \in A[i, t]$ and $v \in A[i', t']$ for some $i < i' \in I$.

Additionally, the construction ensures that

- B has $(3 \cdot 2r + 2) \cdot (2r + 1) \leqslant 24r^2$ cells,
- $W_1 \subseteq B[2r + 1, 1]$ and $W_2 \subseteq B[4r + 2, 1]$, and
- $\mathrm{int}(A) = \mathrm{int}(B)$.

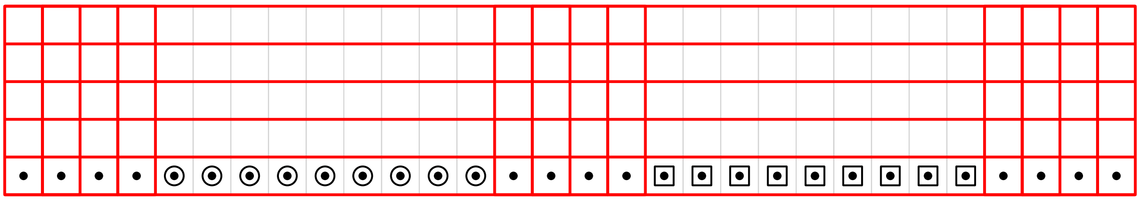

Figure 5.5: An example of a coarsening for radius $r = 2$. The coarsening B (in bold red) overlays the grid A (in gray). The sequence W is located in the bottom row. The two subsequences W_1 and W_2 are marked with circles and squares, respectively.

We build $\mathcal{K}_\star$ as a refinement of $\mathcal{K}$ by encoding into the color of every vertex $v \in V(G)$ for every color $X \in \mathcal{K}$ and for every cell $B[i, j]$ of B the information

(C.1) whether $v \in X$,

(C.2) whether $v \in B[i, j]$,

(C.3) whether v is adjacent in G to a vertex from $X \cap \text{int}(B)$.

$\mathcal{K}_\star$ has at most $2^{k \cdot 24r^2 \cdot 2}$ colors. By (C.1), we can define for every $X \in \mathcal{K}_\star$ a color $\mathcal{K}(X) \in \mathcal{K}$ such that for all $v \in X$ we have $v \in \mathcal{K}(X)$. We define $F_\star \subseteq \mathcal{K}^2$ via the following four rules.

(F.1) If $X \subseteq B[j, t]$ and $Y \subseteq B[j', t']$ and $j < j'$ and $t' \in \{t, t - 1\}$, then

$$(X, Y) \in F_\star \Leftrightarrow \big(\mathcal{K}(X), \mathcal{K}(Y)\big) \in R.$$

(F.2) If $X \subseteq B[j, t]$ and $Y \subseteq B[j', t']$ and $j > j'$ and $t' \in \{t, t + 1\}$, then

$$(X, Y) \in F_\star \Leftrightarrow \big(\mathcal{K}(Y), \mathcal{K}(X)\big) \in R.$$

(F.3) If $X \not\subseteq B$ and $Y \subseteq \text{int}(B)$, or vice versa, then

$$(X, Y) \in F_\star \Leftrightarrow \text{there is an edge between } X \text{ and } Y \text{ in } G.$$

(F.4) Otherwise, $(X, Y) \in F_\star \Leftrightarrow \big(\mathcal{K}(X), \mathcal{K}(Y)\big) \in F$.

By construction, $F_\star$ is symmetric and therefore describes a valid flip. Let $G' := G \oplus F$ and $G_\star := G \oplus F_\star$.

Claim 5.27. *Let $u \in \text{int}(B)$ and $v \in B$ be vertices from rows that are not close in B. Then u and v are non-adjacent in $G_\star$.*

Proof. Let u and v be as in the claim. Then case (F.4) applies and u and v are adjacent in $G_\star$ if and only if they are adjacent in G'. By construction of B, we have also $u \in \text{int}(A)$ and $v \in A$ and u and v are in rows that are not close in A. It follows by property (O.4.1) of $\mathcal{A}$ that u and v are non-adjacent in both G' and $G_\star$. ■

Claim 5.28. *Let $u \in \text{int}(B)$ and $v \in B$ be vertices from columns that are not close in B. Then u and v are non-adjacent in $G_\star$.*

Proof. Let $u \in B[j, t]$ and $v \in B[j', t']$ be as in the claim for some $j, j' \in J$ and $t, t' \in [h]$. If u and v are in rows that are not close, then we are done by Claim 5.27. We can therefore assume $t' \in \{t - 1, t, t + 1\}$.

Assume first that $j < j'$. Since $B[j, *]$ and $B[j', *]$ are not close, we even have $j + 1 < j'$. If $t' = t + 1$, then case (F.4) applies and u and v are adjacent in $G_\star$ if and only if they are adjacent in G'. By Observation 5.26, the adjacency between u and v in G' can be determined using the insulator property (O.4.2) of $\mathcal{A}$: u and v are non-adjacent as desired. If $t' \in \{t, t - 1\}$, then by (C.2), case (F.1) applies and the following are equivalent. Let $X := \mathcal{K}_\star(u)$ and $Y := \mathcal{K}_\star(v)$.

1. The adjacency between u and v was flipped when going from G to $G_\star$.
2. $\big(\mathcal{K}(X), \mathcal{K}(Y)\big) \in R$. (by (F.1))
3. u and v are adjacent in G. (by Observation 5.26 and (O.4.3))

The equivalence between the first and the last item establishes that u and v are non-adjacent in $G_\star$, as desired.

The proof for $j > j'$ works symmetrically. If $t' = t - 1$, then case (F.4) applies, and we argue using (O.4.2). If $t' \in \{t, t+1\}$, then case (F.2) applies, and we argue using (O.4.4). ■

Claim 5.29. *Let $u \in \mathrm{int}(B)$ and v be a vertex not in B. Then u and v are non-adjacent in $G_\star$.*

Proof. Let $X := \mathcal{K}_\star(u)$ and $Y := \mathcal{K}_\star(v)$ be the colors of u and v in $\mathcal{K}_\star$. By (C.2), we have $Y \cap A = \emptyset$ and $X \subseteq \mathrm{int}(A)$. It follows from property (O.3) of $\mathcal{A}$ that every vertex from Y is homogeneous to X in G. Moreover, by (C.3), either every or no vertex in Y has a neighbor in X in G. It follows that the connection between X and Y is homogeneous in G. Also, by (C.3), case (F.3) applies, and the following are equivalent.

1. The adjacency between u and v was flipped when going from G to $G_\star$.
2. There is an edge between X and Y in G. (by (F.3))
3. There is an edge between u and v in G. (by homogeneity of $X \ni u$ and $Y \ni v$ in G)

The equivalence between the first and the last item establishes that u and v are non-adjacent in $G_\star$, as desired. ■

Combining Claim 5.27, Claim 5.28, and Claim 5.29 yields the following crucial observation.

Observation 5.30. In $G_\star$, if vertices $u \in \mathrm{int}(B)$ and $v \in V(G)$ are adjacent, then they are in cells that are close in B.

A straightforward induction now yields that the $2r$-neighborhood of every vertex $u \in W_1 \subseteq B[2r+1, 1]$ in $G_\star$ satisfies $N_{2r}^{G_\star}(u) \subseteq B[\leqslant 4r+1, *]$. In particular, it contains no vertex from $W_2 \subseteq B[4r+2, 1]$. We therefore have $N_r^{G_\star}(W_1) \cap N_r^{G_\star}(W_2) = \emptyset$ as desired. This finishes the proof of the flip-breakability of W. To bound the running time, note that W_1, W_2, $\mathcal{K}_\star$, and $F_\star$ can all be computed in time $O_{r,k}(|V(G)|)$. □

We can now prove Proposition 5.24.

Proof of Proposition 5.24. The non-algorithmic part of the statement immediately follows from Lemma 5.25. For the algorithmic part, assume that $\mathcal{C}$ is also prepattern-free. Given any set W, we can compute an insulated subset $W' \subseteq W$ and a witnessing insulator $\mathcal{A}$ using Proposition 5.23. Applying the algorithm given by Lemma 5.25 to W' and $\mathcal{A}$ yields the desired result. □

5.5 Flip-Breakability Implies Monadic Dependence

We prove the second main implication of the section.

Proposition 5.31. *Every flip-breakable graph class is monadically dependent.*

Let $\varphi(x, y)$ be a formula over the signature of colored graphs, G^+ be a colored graph, and W be a set of vertices in G^+. We say that φ *shatters* W in G^+, if there exists vertices $(a_R)_{R \subseteq W}$ such that for all $b \in P$ and $R \subseteq W$,

$$G^+ \models \varphi(a_R, b) \Leftrightarrow b \in R.$$

Let G be an (uncolored) graph, and W be a set of vertices in G. We say that φ *monadically shatters* W in G, if there exists a coloring G^+ of G in which φ shatters W.

Fact 5.32 ([4]). *A graph class $\mathcal{C}$ is monadically dependent if and only if for every formula $\varphi(x, y)$ over the signature of colored graphs, there exists a bound m such that φ monadically shatters no set of size m in any graph of $\mathcal{C}$.*

Proposition 5.31 is implied by the following.

Lemma 5.33. *Let $\varphi(x, y)$ be a formula and $k \in \mathbb{N}$. There exist $r_\varphi, m_{\varphi,k} \in \mathbb{N}$, where r_φ depends only on φ, such that no graph G contains a set of at least $m_{\varphi,k}$ vertices W such that*

- *φ monadically shatters W in G, and*
- *W is (r_φ, k)-flip-breakable in G.*

Proof. Let q be the quantifier rank of φ. We set $r_\varphi := 2 \cdot 7^q + 1$. Let s be the number of colors used by φ and let $t := p(q, s \cdot k)$ be the number of colors from Lemma 4.10, that are needed to determine the truth value of formulas in the signature of $(s \cdot k)$-colored graphs that have the same quantifier-rank as $\varphi(x, y)$. Let $m_{\varphi,k} := 3(t + 1)$.

Assume now towards a contradiction the existence of an (r_φ, k)-flip-breakable in set W in G of size $m_{\varphi,k}$ such that φ monadically shatters W in G. Then there exists an s-coloring G^+ of G in which φ shatters W. We apply flip-breakability, which yields a k-flip H of G together with two disjoint sets $A, B \subseteq W$ each of size at least $t + 1$, such that $N_r^H(A) \cap N_r^H(B) = \emptyset$. By using k colors to encode the flip, we can rewrite φ to a formula ψ with the same quantifier-rank as φ, such that there exists an $(s \cdot k)$-coloring H^+ of H where for all $u, v \in V(G)$,

$$G^+ \models \varphi(u, v) \Leftrightarrow H^+ \models \psi(u, v).$$

In particular, ψ shatters W in H^+. Since ψ has the same quantifier-rank as φ and is a formula over the signature of $(s \cdot k)$-colored graphs, by Lemma 4.10 there exists a coloring of H^+ with t colors such that the truth of $\psi(u, v)$ in H^+ only depends on the colors of u and v for all vertices u and v with distance greater than r in H^+. Recall that A and B each have size $t + 1$. By the pigeonhole principle, there exist two distinct vertices $a_1, a_2 \in A$ that are assigned the same color and two distinct vertices $b_1, b_2 \in B$ that are assigned the same color. Since ψ shatters W in H^+, there exists a vertex $v \in V(G)$ such that

$$H^+ \models \psi(v, a_1) \wedge \neg\psi(v, a_2) \wedge \psi(v, b_1) \wedge \neg\psi(v, b_2).$$

By Lemma 4.10, v must be contained in $N_r^H(A) \cap N_r^H(B)$, as the truth of ψ is inhomogeneous among both v and $\{a_1, a_2\}$ and among v and $\{b_1, b_2\}$; a contradiction to $N_r^H(A) \cap N_r^H(B) = \emptyset$. □

Combining the results of Section 5.3 and Section 5.4 yields the desired chain of implications.

prepattern-free ⇒ insulation-property ⇒ flip-breakable ⇒ mon. dependent

In the remaining two sections of Chapter 5, we provide the deferred proofs of the two insulator growing lemmas (Lemmas 5.19 and 5.20).

5.6 Sample Sets

We work towards proving Lemmas 5.19 and 5.20, which grow the height of an insulator. In this section we show that in prepattern-free classes, given an insulator $\mathcal{A}$, we can extract a subinsulator $\mathcal{B}$ and a small *sample set* of vertices which can be used to approximately represent the connections of all the vertices in the graph towards $\mathcal{B}$. We give some notation to make this statement precise.

Definition 5.34. Given a graph G, a vertex $v \in V(G)$, and a set $A \subseteq V(G)$, we define the *atomic type* of v over A as

$$\text{atp}(v/A) := \{(R,a) : G \models R(v,a), R \in \{E, =\}, a \in A\}.$$

Observation 5.35. Let G be a graph, $u, v \in V(G)$, and $A \subseteq B \subseteq V(G)$. Then

$$\text{atp}(u/A) \neq \text{atp}(v/A) \quad \Rightarrow \quad \text{atp}(u/B) \neq \text{atp}(v/B).$$

Definition 5.36. Let G be a graph containing an insulator $\mathcal{A}$ with grid A indexed by I. Let $v, s_<, s_>$ be vertices from G, $i \in I$, and $m \in \mathbb{N}$. We say v *is* $(m, i, s_<, s_>)$*-sampled on* $\mathcal{A}$ if

$$\text{atp}(v/A[<i, *]) = \text{atp}(s_</A[<i, *]) \quad \text{and} \quad \text{atp}(v/A[\geqslant i+m, *]) = \text{atp}(s_>/A[\geqslant i+m, *]).$$

We call m the *margin*, i the *exceptional index*, $s_<$ the *left-sample*, and $s_>$ the *right-sample*.

Definition 5.37. Fix $p \in \mathbb{N}$. Let G be a graph containing an insulator $\mathcal{A}$ indexed by I and let $S \subseteq V(G)$. We say S *samples* G on $\mathcal{A}$ with margin m if there exists functions $\text{ex} : V(G) \to I$ and $s_<, s_> : V(G) \to S$ such that every $v \in V(G)$ is $\big(m, \text{ex}(v), s_<(v), s_>(v)\big)$-sampled on $\mathcal{A}$.

We are now ready to state the main result of this section.

Lemma 5.38. *Fix $t \in \mathbb{N}$. For every graph G and insulator $\mathcal{A}$ indexed by J in G, there is a subsequence $I \subseteq J$ of size $U_t(|J|)$ such that either*

- *G contains a prepattern of order t on $\mathcal{A}|_I$, or*
- *there is a set $S \subseteq V(G) \setminus \mathcal{A}|_I$ of size $\text{const}(t)$ that samples G on $\mathcal{A}|_I$ with margin 2.*

Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes I and one of the two outcomes (a prepattern or a sampling set S) in time $O_t(|V(G)|^2)$.

We will build the set S iteratively by extracting single sampling vertices one by one.

5.6.1 Extracting Single Sample Vertices

Before we show how to extract a new sample vertex, we state some auxiliary lemmas about subinsulators and sampling sets.

Lemma 5.39. *Let $\mathcal{A}$ be an insulator indexed by J in a graph G. Let I be a subsequence of J, $u, v \in V(G)$, and $i \in \text{tail}(I)$. Let A and B be the grids of $\mathcal{A}$ and $\mathcal{A}|_I$ respectively. Then*

$$\text{atp}(u/A[i, *]) \neq \text{atp}(v/A[i, *]) \quad \Rightarrow \quad \text{atp}(u/B[i, *]) \neq \text{atp}(v/B[i, *]).$$

Proof. Follows from Observation 5.35 and Observation 5.8. ■

Lemma 5.40. *Let $\mathcal{A}$ be a grid indexed by a sequence J of length at least four in a graph G and let $v \in V(G)$. There exists a subsequence $I \subseteq J$ of length $|I| \geqslant \frac{1}{2}|J|$ such that $v \notin \mathcal{A}|_I$.*
Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes I in time $O(|V(G)|)$.

Proof. Up to renaming, assume $J = (1, \dots, n)$. If $v \notin \mathcal{A}|_I$, we can just set $I := J$. Otherwise, let $i \in J$ be such that $v \in A[i, *]$ where A is the grid of $\mathcal{A}$. We choose I as the larger of the two sequences $(1, \dots, i-1)$ and $(i, \dots, n)$. Since I has length at least two, this defines a subinsulator. By Observation 5.8, $v \notin A|_I$. The bound on the running time is obvious. □

Lemma 5.41. *Fix $m \in \mathbb{N}$. Let G be a graph containing an insulator $\mathcal{A}$ indexed by J and $S \subseteq V(G)$ be a set that samples G on $\mathcal{A}$ with margin m. For every $I \subseteq J$, S also samples G on $\mathcal{A}|_I$ with margin m.*

Proof. Let $B := A|_I$ be the grid of $\mathcal{A}|_I$. As in the first item in the proof of Lemma 5.9, we observe for every $i \in I$

$$B[<i,*] \subseteq A[<i,*] \quad \text{and} \quad B[>i,*] \subseteq A[>i,*].$$

Hence, by contrapositive of Observation 5.35, for all vertices $u, s_<, s_>$,

$$\text{atp}(v/A[<i,*]) = \text{atp}(s_</A[<i,*]) \quad \Rightarrow \quad \text{atp}(v/B[<i,*]) = \text{atp}(s_</B[<i,*]),$$
$$\text{atp}(v/A[>i,*]) = \text{atp}(s_>/A[>i,*]) \quad \Rightarrow \quad \text{atp}(v/B[>i,*]) = \text{atp}(s_>/B[>i,*]).$$

The lemma now follows from Definition 5.37. □

We will now show how to extract a new sample vertex. We start with a given set of sample vertices S. In the absence of large prepatterns, we either find a subinsulator on which S samples G, or find a new vertex v by which we will later use to extend the sample set.

Lemma 5.42. *Fix $k, t \in \mathbb{N}$. For every graph G and insulator $\mathcal{A}$ indexed by J, and every vertex set S of size at most k, there is a subsequence $I \subseteq J$ of size $U_{k,t}(|J|)$ such that either*

- *G contains a bi-prepattern of order t on $\mathcal{A}|_I$, or*
- *S samples G on $\mathcal{A}|_I$ with margin 2, or*
- *there is a vertex $v \notin \mathcal{A}|_I$, such that for all $s \in S$ and every column C in $\mathcal{A}|_I$*

$$\text{atp}(v/C) \neq \text{atp}(s/C).$$

Moreover, there is an algorithm that, given G, $\mathcal{A}$, and S, computes the sequence I and one of the three outcomes (a bi-prepattern, the conclusion that S samples G, or a vertex v) in time $O_{k,t}(|V(G)|^2)$.

Proof. We first show how to construct a sequence I with the desired properties. Afterwards we analyze the running time. Getting the desired bounds will then require a small preprocessing that reduces the size of J. The proof is split into multiple paragraphs.

Notation. For vertices $v, s \in V(G)$ and a set $U \subseteq V(G)$, we say v *is s-connected to U* if

$$\text{atp}(v/U) = \text{atp}(s/U).$$

We generalize this to sets $S_\star \subseteq S$ and say v *is $S_\star$-connected to U* if

$$\{s \in S : v \text{ is } s\text{-connected to } U\} = S_\star.$$

Ramsey. We start by defining some coloring to which we will apply Ramsey's theorem. Let A be the grid of $\mathcal{A}$ and M_A be the function from Lemma 5.6. For $t' \in \{6, 4t\}$ and $S_1, \ldots, S_{t'-1} \subseteq S$, we label all t'-tuples $i_1 < \cdots < i_{t'} \in J$ with a color indicating whether

$$\exists v \bigwedge_{l=1,\ldots,t'-1} v \text{ is } S_l\text{-connected to } M_A(i_l, i_{l+1}). \tag{5.1}$$

Applying Ramsey's Theorem (Fact 4.13) to this coloring gives us a subsequence $I' \subseteq J$ such that for all $t' \in \{8, 6t\}$ and $S_1, \ldots, S_{t'-1} \subseteq S$, the above equation (5.1) either holds for all or no t'-tuples $i_1 < \cdots < i_{t'} \in I'$. Note that the number of colors is bounded by const(k, t), which guarantees $|I'| \geqslant U_{k,t}(|J|)$. We can assume without loss of generality that $I :=$ tail(I') has length at least $24t$. In order to simplify notation, we assume up to renaming that $I' = (0, \ldots, n)$ and $I = (1, \ldots, n)$. Let $B := A|_{I'}$ be the grid of $\mathcal{B} := \mathcal{A}|_{I'}$. By Lemma 5.6, we observe the following.

Observation 5.43. For all $i \in I$ we have $B[i, *] = M_A(i-1, i)$.

We say a subsequence of a sequence K is 1-*spaced* if it contains no consecutive elements from K.

Claim 5.44. *Let $s \in S$ and $\bar{w} \in \{0,1\}^{2t}$. If there is a 1-spaced subsequence $i_1 < \cdots < i_{2t}$ of I such that*

$$\exists v \bigwedge_{l=1,\ldots,2t} v \text{ is } s\text{-connected to } B[i_l, *] \text{ if and only if } w_l = 1,$$

then the above holds for all such 1-spaced subsequences of the same length.

Proof. By Observation 5.43, the claim holds for a 1-spaced subsequence $i_1 < \cdots < i_{2t}$ if and only if equation (5.1) holds for the corresponding $4t$-tuple

$$i_1 - 1 < i_1 < i_2 - 1 < i_2 < \cdots < i_{2t} - 1 < i_{2t} \in I'$$

of distinct elements and some $S_1, \ldots, S_{4t-1} \subseteq S$ with $s \in S_{2l-1} \Leftrightarrow w_l = 1$. Note that we use 1-spacedness to guarantee that all elements in the $4t$-tuple are distinct.

Assume there is a 1-spaced subsequence satisfying the claimed statement, and let the sets $S_1, \ldots, S_{4t-1} \subseteq S$ be witnesses of the truth of equation (5.1) for the corresponding $4t$-tuple. By Ramsey's theorem, equation (5.1) with $S_1, \ldots, S_{4t-1} \subseteq S$ holds for all $4t$-tuples, and thus, the claimed statement holds for all 1-spaced subsequences. ■

Claim 5.45. *If there is $i, j \in I$ with $i + 2 < j$ and sets $P_i, Q_i, P_j, Q_j \subseteq S$ with $P_i \neq Q_i$ and $P_j \neq Q_j$ such that*

$$\begin{aligned} \exists v\ \ & v \text{ is } P_i\text{-connected to } B[i, *] \\ & \wedge\, v \text{ is } Q_i\text{-connected to } B[i+1, *] \\ & \wedge\, v \text{ is } P_j\text{-connected to } B[j, *] \\ & \wedge\, v \text{ is } Q_j\text{-connected to } B[j+1, *] \end{aligned}$$

then the above holds for all $i, j \in I$ with $i + 2 < j$.

Proof. We again use Observation 5.43. Thus, the statement of the claim holds for a given $i, j \in I$ if and only if equation (5.1) holds for the corresponding 6-tuple

$$i - 1 < i < i + 1 < j - 1 < j < j + 1 \in I'$$

of distinct elements and some $S_1, \ldots, S_5 \subseteq S$ with $S_1 = P_i$ and $S_2 = Q_i$ and $S_4 = P_j$ and $S_5 = Q_j$. The rest follows as in Claim 5.44. ■

Constructing a Prepattern. We say an index $i \in I$ is an *alternation point* of a vertex v on $\mathcal{B}$ if v is P-connected to $B[i, *]$ and Q-connected to $B[i+1, *]$ for distinct sets $P \neq Q \subseteq S$.

Claim 5.46. *One of the following two conditions holds.*

1. *For every vertex $v \in V(G)$ with alternation points $i, j \in I$, we have $|i - j| \leqslant 2$.*
2. *G contains a bi-prepattern of order t on $\mathcal{B}$.*

Proof. Assume the first condition fails and let $v \in V(G)$ be a vertex with alternation points $i, j \in I$ such that $i + 2 < j$. By Claim 5.45, we can assume without loss of generality that v has alternation points $6t$ and $19t$. In the following, we find two 1-spaced subsequences of I

$$I_\star = \underbrace{i_1 < \cdots < i_t}_{\subseteq(1,\dots,4t)} \;<\; \underbrace{i_\star}_{\in\{6t,6t+1\}} \;<\; \underbrace{i_{t+1} < \cdots < i_{2t}}_{\subseteq(9t,11t)},$$
$$J_\star = \underbrace{j_1 < \cdots < j_t}_{\subseteq(13t,\dots,17t)} \;<\; \underbrace{j_\star}_{\in\{19t,19t+1\}} \;<\; \underbrace{j_{t+1} < \cdots < j_{2t}}_{\subseteq(22t,\dots,24t)}$$

and vertices $s_1, s_2 \in S$ satisfying the following property: Either

- $i_\star$ is the first index among $I_\star$, such that v is s_1-connected to $B[i_\star, *]$, or
- $i_\star$ is the first index among $I_\star$, such that v is *not* s_1-connected to $B[i_\star, *]$,

and moreover, either

- $j_\star$ is the first index among $J_\star$, such that v is s_2-connected to $B[j_\star, *]$, or
- $j_\star$ is the first index among $J_\star$, such that v is *not* s_2-connected to $B[j_\star, *]$.

We start by choosing $s_1 \in S$ to be an arbitrary vertex witnessing the alternation point $6t$ of v, that is,

$$v \text{ is } s_1\text{-connected to } B[6t, *] \quad \Leftrightarrow \quad v \text{ is not } s_2\text{-connected to } B[6t+1, *].$$

By a simple majority argument, we can choose $i_1, \dots, i_t$ to be a 1-spaced subsequence of $(1, \dots, 4t)$ such that v is s_1-connected either to none or to all of $B[i_1, *], \dots, B[i_t, *]$. In the first case (respectively last case) we set $i_\star$ to be the index from $\{6t, 6t+1\}$ such that v is s_1-connected to $B[i_\star, *]$ (respectively *not* s_1-connected to $B[i_\star, *]$). We can now choose $i_{t+1}, \dots, i_{2t}$ to be an arbitrary 1-spaced subsequence of $(9t, 11t)$. This concludes the construction of $I_\star$. We further choose $s_2 \in S$ to be an arbitrary vertex witnessing the alternation point $19t$ of v and construct $J_\star$ analogously as a 1-spaced subsequence of $(13t, \dots, 24t)$.

For every $i \in [t]$ let I_i be a subsequence of $I_\star$ of length t such that $i_\star$ is the ith element of I_i. Such a sequence exists, since $i_\star$ has both t successors and t predecessors in $I_\star$. Similarly, for every $j \in [t]$ let J_j be a subsequence of $J_\star$ of length t such that $j_\star$ is the jth element of J_j. For every $i, j \in [t]$, by concatenating I_i and J_j, we obtain a 1-spaced subsequence $I_{i,j} = (p_1, \dots, p_t, q_1, \dots, q_t)$ of I of length $2t$ such that

- the s_1-connection from v to $B[p_1, *], \dots, B[p_{i-1}, *]$ is homogeneous
  but switches at $B[p_i, *]$,
- the s_2-connection from v to $B[q_1, *], \dots, B[q_{j-1}, *]$ is homogeneous
  but switches at $B[q_j, *]$.

By Claim 5.44 and witnessed by v and all the $I_{i,j}$, we can fix an arbitrary 1-spaced subsequence $(p_1, \dots, p_t, q_1, \dots, q_t)$ of I of length $2t$ and there exist vertices $\{c_{i,j} : i, j \in [t]\}$ such that either

- $\forall i, j \in [t]$: $c_{i,j}$ is not s_1-connected to $B[p_1, *], \dots, B[p_{i-1}, *]$
  but s_1-connected to $B[p_i, *]$, or
- $\forall i, j \in [t]$: $c_{i,j}$ is s_1-connected to $B[p_1, *], \dots, B[p_{i-1}, *]$
  but not s_1-connected to $B[p_i, *]$,

and similarly either

- $\forall i, j \in [t]$: $c_{i,j}$ is not s_2-connected to $B[q_1, *], \dots, B[q_{j-1}, *]$
  but s_2-connected to $B[q_j, *]$, or
- $\forall i, j \in [t]$: $c_{i,j}$ is s_2-connected to $B[q_1, *], \dots, B[q_{j-1}, *]$
  but not s_2-connected to $B[q_j, *]$.

Let $\alpha_1(y; x, s_1)$ be the quantifier-free formula checking whether $\text{atp}(x/\{y\}) \neq \text{atp}(s_1/\{y\})$. Whenever v is not s_1-connected to a column $B[i,*]$, then this is witnessed by an element $u \in B[i,*]$ such that $G \models \alpha_1(u; v, s_1)$. If v is s-connected to $B[i,*]$, then no such element exists in $B[i,*]$. Among $p_1, \dots, p_t$, we have that either

- $\forall i, j \in [t]$: p_i is the first index such that $\alpha_1(B[p_i,*]; c_{i,j}, s_1)$ is *empty* (this happens in the case where the $c_{i,j}$ are s_1-connected to $B[p_i,*]$), or
- $\forall i, j \in [t]$: p_i is the first index such that $\alpha_1(B[p_i,*]; c_{i,j}, s_1)$ is *non-empty* (this happens in the case where the $c_{i,j}$ are *not* s_1-connected to $B[p_i,*]$).

Similarly, there is a quantifier-free formula $\alpha_2(y, x, s_2)$ checking whether

$$\text{atp}(x/\{y\}) \neq \text{atp}(s_2/\{y\}).$$

Among $q_1, \dots, q_t$, we have that either

- $\forall i, j \in [t]$: q_j is the first index such that $\alpha_2(B[q_j,*]; c_{i,j}, s_2)$ is *empty*, or
- $\forall i, j \in [t]$: q_j is the first index such that $\alpha_2(B[q_j,*]; c_{i,j}, s_2)$ is *non-empty*.

This proves that G contains a bi-prepattern of order t on $\mathcal{B}$ witnessed by the sequences $(p_1, \dots, p_t)$, $(q_1, \dots, q_t)$, the $c_{i,j}$ vertices, the parameters s_1, s_2, and the formulas α_1 and α_2. ∎

Extracting a Sample Vertex. If the second condition of Claim 5.46 holds, then G contains a bi-prepattern of order t on $\mathcal{A}|_{I'}$, so I' can play the role of the sequence I in the statement of the lemma, and we are done. We therefore assume from now on that the first condition of Claim 5.46 holds.

Claim 5.47. *One of the following two conditions holds.*

1. *For every vertex $u \in V(G)$, that is $\emptyset$-connected to columns $B[i,*]$ and $B[j,*]$ for $i, j \in I$, we have $|i - j| \leqslant 1$. In particular each vertex is $\emptyset$-connected to at most two columns in B.*
2. *There is a sequence K of length at least $\frac{1}{6}|I|$ and a vertex $v \notin \mathcal{A}|_K$, such that v is $\emptyset$-connected to every column of $\mathcal{A}|_K$.*

Proof. Assume the first condition fails. We have a vertex u that is $\emptyset$-connected to columns $B[i,*]$ and $B[j,*]$ for $i + 1 < j \in I$. In order to show the second condition, we will first find a vertex v together with a sequence $K' \subseteq I$ of length at least $\frac{1}{3}|I|$, such that v is $\emptyset$-connected to every column $B[p,*]$ with $p \in K$. As (i, j) is a 1-spaced subsequence of I, Claim 5.44 yields a vertex v that is $\emptyset$-connected to $B[m,*]$ and $B[m+2,*]$ where $m := \lfloor \frac{1}{2}|I| \rfloor$. Assume v has at least one alternation point on $\mathcal{B}$ and its earliest alternation point q satisfies $m \leqslant q$. Then u has the same connection type to every $B[p,*]$ with $1 \leqslant p \leqslant m$. As v is $\emptyset$-connected to $B[m,*]$, we can set $K' := (1, \dots, m)$. Otherwise, either v has no alternation points on $\mathcal{B}$, or the latest alternation point q of v satisfies $q \leqslant m + 1$, by Claim 5.46. Then u has the same connection type to every $B[p,*]$ with $m+2 \leqslant p \leqslant n$. As v is $\emptyset$-connected to $B[m+2,*]$, we can set $K' := (m+2, \dots, n)$. This finishes the construction of v and K'. We use Lemma 5.40 to obtain a sequence $K \subseteq K'$ of length $|K| \geqslant \frac{1}{2}|K'| \geqslant \frac{1}{6}|I|$ with $v \notin \mathcal{A}|_K$. By Lemma 5.39, v is $\emptyset$-connected to every column of $\mathcal{A}|_K$. ∎

Establishing the Sampling Property. If the second condition of Claim 5.47 holds, then there exists a sequence K that can play the role of I in the statement of the lemma, and we are done. We therefore assume from now on that the first condition of Claim 5.47 holds.

Claim 5.48. *S samples G on $\mathcal{A}|_K$ with margin 2 for $K := (2, \dots, n-2) \subseteq I$.*

Proof. We have to choose for every vertex $v \in V(G)$ an exceptional index $i \in \text{tail}(K) = (3\ldots, n-2)$ and two vertices $s_<, s_> \in S$ such that v is $s_<$-connected to $A|_K[<i, *]$ and $s_>$-connected to $A|_K[>i+1, *]$. By Observation 5.7, we have $A|_K = B|_K$. Thus, by Observation 5.8, we have

$$A|_K[<i, *] \subseteq B[\{3, \ldots, i-1\}, *] \quad \text{and} \quad A|_K[>i+1, *] \subseteq B[\{i+2, \ldots, n-2\}, *].$$

Take any vertex v. By Claim 5.46, there exist successive indices $i_1 \in I$ and $i_2 := i_1 + 1$, and sets $S_1, S_2 \subseteq S$, such that v is S_1-connected to all columns $B[<i_1, *]$ and S_2-connected to all columns $B[>i_2, *]$.

Assume $S_1 \neq \emptyset \neq S_2$. Then we can arbitrarily choose $s_< \in S_1$ and $s_> \in S_2$ and set

$$i := \begin{cases} 3 & \text{if } i_1 < 3, \\ n-2 & \text{if } i_1 > n-2, \\ i_1 & \text{otherwise.} \end{cases}$$

Now v is $s_<$-connected to $B[\{3, \ldots, i-1\}, *]$ and $s_>$-connected to $B[\{i+2, \ldots, n-2\}, *]$, proving the claim.

Assume $S_1 = \emptyset \neq S_2$. If $i_1 > 3$, then v is $\emptyset$-connected to the first three columns of B, contradicting Claim 5.47. So we have $i_1 \leqslant 3$. We set $i := 3$ and choose an arbitrary $s_> \in S_2$. As desired, v is $s_>$-connected to $B[\{i+2, \ldots, n-2\}, *]$. As $B[\{3, \ldots, i-1\}, *]$ is empty any vertex from S can take the role of $s_<$.

Assume $S_1 \neq \emptyset = S_2$. The proof is symmetric to the previous case.

Assume $S_1 = \emptyset = S_2$. Since $|I| > 8$, we either find left of i_1 or right of i_2 at least three columns of B to which v is $\emptyset$-connected. This is a contradiction to Claim 5.47. ■

We have successfully established the sampling property which proves that K (I in the statement) has the desired properties.

Running Time. In the previous paragraphs, we have proven the existence of a sequence I with the desired properties. Let us redefine $n := |V(G)|$. To show that I can be constructed in time $O_{k,t}(n^2)$, we first consider the following preprocessing routine. Let $t_\star := \max(6, 4t)$ and choose a subsequence $J_0 \subseteq J$ of size $\lfloor |J|^{1/t_\star} \rfloor$. Note that $|J_0| \leqslant n^{1/t_\star}$. By applying the construction to $\mathcal{A}|_{J_0}$ and J_0 instead of $\mathcal{A}$ and J, we obtain a subsequence $I \subseteq J_0 \subseteq J$ with the desired properties that still has size $|I| \geqslant U_{k,t}(|J_0|) \geqslant U_{k,t}(|J|)$. By this argument, and as we can build J_0 and $\mathcal{A}|_{J_0}$ in time $O_{k,t}(n)$, we can from now on assume without loss of generality that $|J| \leqslant n^{1/t_\star}$.

Towards computing the coloring needed for the Ramsey application, we first compute for each $v \in V(G)$, $S_\star \subseteq S$, and $i \in J$, whether v is $S_\star$-connected to the column $A[i, *]$ of $\mathcal{A}$. As the columns of A are disjoint and $|S| \leqslant k$, this takes a total time of $O(2^k \cdot n^2)$. Moreover, for each vertex $v \in V(G)$, subset $S_\star \subseteq S$ and pair $i < i' \in J$ we calculate whether v is $S_\star$-connected to $M_A(i, i')$. By Lemma 5.6, if $\mathcal{A}$ is orderless this amounts to checking whether v is $S_\star$-connected to $A[i', *]$, which we have already computed in the previous step. If $\mathcal{A}$ is ordered, we instead check whether v is $S_\star$-connected to each of $A[m, *]$ for $i < m \leqslant i'$. As $|J| \leqslant \sqrt{n}$ and with the data from the previous step we can do this check in time $O(\sqrt{n})$ for a single vertex v, set $S_\star$ and pair i, i'. Since $|J| \leqslant n^{1/4}$, there are at most $\sqrt{n}$ pairs i, i' that need to be checked. It follows that we can compute the desired data for all vertices v, sets $S_\star$, and pairs i, i' in total time $O(2^k \cdot n \cdot \sqrt{n} \cdot \sqrt{n}) = O(2^k \cdot n^2)$. We recall the construction of the coloring for the Ramsey application: For $t' \in \{6, 4t\}$ and $S_1, \ldots, S_{t'-1} \subseteq S$, the t'-tuples $i_1 < \cdots < i_{t'} \in J$ are labeled with a color indicating whether

$$\exists v \bigwedge_{l=1,\ldots,t'-1} v \text{ is } S_l\text{-connected to } M_A(i_l, i_{l+1}).$$

Using our precomputed information, for a single t'-tuple, we can compute its colors in time $O_{k,t}(n)$ by iterating over all vertices $v \in V(G)$. As $|J| \leqslant n^{1/t_\star}$, there are at most n many t'-tuples from J, so we can compute the coloring in time $O_{k,t}(n^2)$. Due to the size bounds on J, applying Ramsey's Theorem (Fact 4.14) to the coloring runs in time $O_{k,t}(n)$. This yields the sequence I.

Having constructed I, we obtain the insulator $\mathcal{B} = \mathcal{A}|_{I'}$. By Observation 5.43, our precomputed information can also be used to check whether a vertex $v \in V(G)$ is $S_\star$-connected to a column $B[i, *]$ of $\mathcal{B}$ for some $S_\star \subseteq S$ and $i \in I$. Let us now show how to compute one of the three outcomes: a bi-prepattern, a new sample vertex v, or a large subsequence of I on which S samples G.

- We can check for every vertex, whether it contains two alternation points with distance bigger than 2 on $\mathcal{B}$ in time $O_k(n \cdot |I|) \leqslant O_k(n^2)$. If such a vertex exists, the proof of Claim 5.46 yields that there is a bi-prepattern on $\mathcal{B}$ of size t on every 1-spaced subsequence of length $2t$ of columns of $\mathcal{B}$. We can choose any such sequence and search for the witnessing vertices in time $O_{k,t}(n)$.
- If the previous case does not apply, we can again search in time $O_k(n \cdot |I|) \leqslant O_k(n^2)$ for a new sample vertex v and a corresponding subsequence of I (cf. Claim 5.47).
- If none of the two previous cases apply, we immediately find the sampled subsequence by dropping the first and two last elements of I (cf. Claim 5.48).

We have shown that each step of the construction can be carried out in time $O_{k,t}(n^2)$. This concludes the proof of Lemma 5.42. □

5.6.2 Extracting Small Sample Sets

We use the following Ramsey-type result for set systems due to Ding, Oporowski, Oxley, and Vertigan [20, Cor. 2.4] (see also [45, Thm. 2]). Recall two distinct vertices u and v are twins in a graph G if $N(u) \setminus \{v\} = N(v) \setminus \{u\}$.

Fact 5.49 ([20, Cor. 2.4] [45, Thm. 2]). *There exists a computable function $Q : \mathbb{N} \to \mathbb{N}$ such that for every $\ell \in \mathbb{N}$ and for every bipartite graph $G = (L, R, E)$, where L has size at least $Q(\ell)$ and contains no twins, contains a matching, co-matching, or half-graph of order ℓ as an induced subgraph. Moreover, there is an algorithm that, given G, computes the induced subgraph in time $O(|V(G)|^2)$.*

The computability of Q is made explicit in [45, Thm. 2]. While the construction of [20] is algorithmic, no running time is stated for Fact 5.49. To be self-contained, we instead deduce an algorithm a posteriori.

Proof of the running time of Fact 5.49. Let Q be the function given by the non-algorithmic part of Fact 5.49. To prove an algorithmic version of the statement we weaken the bounds by demanding L to have size at least $f^{-1}(Q(\ell))$ instead, where $f(x) = \lfloor \sqrt{\log(x)/2} \rfloor$.

We first compute an induced subgraph G' of G on sets $L' \subseteq L$ and $R' \subseteq R$. To this end, choose L' as an arbitrary subset of L of size $\lfloor \sqrt{\log(|L|)/2} \rfloor$. Choose $R' \subseteq R$ of size at most $|L'|^2 \leqslant \log(|L|)/2$ by picking for each pair of distinct vertices $u, v \in L'$ a vertex from the symmetric difference of the neighborhoods of u and v in R. Since L contains no twins in G, such a vertex always exists. G' has size at most $\log(|L|)$ and can be constructed in time $O(|V(G)|^2)$.

By our choice of f, we observe that L' has size at least $Q(\ell)$. By construction, L' contains no twins in G'. By the non-algorithmic version of Fact 5.49, we know that G' contains a matching, co-matching, or half-graph of order ℓ. Since G' has at most $\log(|L|)$ vertices, we can perform a brute force search in time $O(|V(G)|^2)$. As G' is an induced subgraph of G, the computed solution also applies to G. □

We can now prove Lemma 5.38, which we restate for convenience.

Lemma 5.38. *Fix $t \in \mathbb{N}$. For every graph G and insulator $\mathcal{A}$ indexed by J in G, there is a subsequence $I \subseteq J$ of size $U_t(|J|)$ such that either*

- *G contains a prepattern of order t on $\mathcal{A}|_I$, or*
- *there is a set $S \subseteq V(G) \setminus \mathcal{A}|_I$ of size $\text{const}(t)$ that samples G on $\mathcal{A}|_I$ with margin 2.*

Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes I and one of the two outcomes (a prepattern or a sampling set S) in time $O_t(|V(G)|^2)$.

Proof. We will inductively compute subsequences $I_0, I_1, \dots$ of J and subsets $S_0, S_1, \dots$ of $V(G)$ using the following procedure. For the base case we set $I_0 := J$ and $S_0 := \emptyset$. In the inductive step we are given I_i and S_i and apply Lemma 5.42 on t, $A|_{I_i}$, and S_i. This yields the subsequence I_{i+1} of size $U_{|S_i|,t}(|I_i|)$ and the insulator $(\mathcal{A}|_{I_i})|_{I_{i+1}} = \mathcal{A}|_{I_{i+1}}$ such that either

(C.1) G contains a prepattern of order t on $\mathcal{A}|_{I_{i+1}}$, or

(C.2) S_i samples G on $\mathcal{A}|_{I_{i+1}}$ with margin 2, or

(C.3) there is a vertex $s_{i+1} \notin \mathcal{A}|_{I_{i+1}}$, such that for all $s \in S_i$ and every column C in $\mathcal{A}|_{I_{i+1}}$

$$\text{atp}(s_{i+1}/C) \neq \text{atp}(s/C).$$

In the first two cases, we stop the construction and set $I := I_{i+1}$ and $S := S_i$. In the third case, we continue the construction with $S_{i+1} := S_i \cup \{s_{i+1}\}$.

Claim 5.50. *For every i we have*

(P.1) *all vertices of S_i have a pairwise different atomic type over every column of $\mathcal{A}|_{I_i}$,*

(P.2) *no $s \in S_i$ is contained in $\mathcal{A}|_{I_i}$,*

(P.3) *$|S_i| = i$, and*

(P.4) *I_i has size $U_{i,t}(J)$.*

Proof. We prove the properties by induction on i. The base cases hold trivially. The properties (P.1) and (P.2) hold on S_i and $\mathcal{A}|_{I_i}$ by induction, on S_i and $\mathcal{A}|_{I_{i+1}}$ by Lemma 5.39, and finally on S_{i+1} and $\mathcal{A}|_{I_{i+1}}$ by the choice of s_{i+1} in (C.3). By (P.1), all elements of S_i are distinct. As we only add one element per turn, this proves (P.3). It follows that I_{i+1} has size $U_{|S_i|,t}(|I_i|) = U_{i,t}(|I_i|)$ and by induction I_{i+1} has size $U_{i,t}(J)$, which proves (P.4). ■

Let $k := Q^{3t}(t)$, where Q is the function given by Fact 5.49. We have $k = \text{const}(t)$.

Claim 5.51. *If the construction runs for k steps, then G contains a prepattern of order t on $\mathcal{A}|_{I_k}$.*

Proof. By Claim 5.50, the set S_k consists of k vertices which have pairwise different atomic types with respect to every column of the grid A of $\mathcal{A}|_{I_k}$. Since no vertex of S_k is contained in A, we know that these vertices all have the same type with regard to the equality relation and must therefore have a pairwise different type with regard to the edge relation, that is, in the graph G, the vertices of S_k have pairwise different neighborhoods in every column of A. Therefore, in the semi-induced bipartite graph between any subset $L \subseteq S_k$ and any column $R := A[i, *]$ of A, there are no twins in L and the preconditions of Fact 5.49 are met. We iterate Fact 5.49 a total number of $3t$ times between S_k and $3t$ columns of A, which we can choose arbitrarily. Finally, we apply the pigeonhole principle. This yields a size t subset $S_\star \subseteq S_k$ and columns $C_1, \dots, C_t$ of A containing subsets $R_1, \dots, R_t$ such that the semi-induced bipartite graph between $S_\star$ and R_i

- is a matching for all $i \in [t]$, or
- is a co-matching for all $i \in [t]$, or
- is a half-graph for all $i \in [t]$.

This witnesses a mono-prepattern of order t on $\mathcal{A}|_{I_k}$. ■

We can now finish the proof. By Claim 5.51, if the construction runs for k steps, we set $I := I_k$ and G contains a prepattern of order t on $A|_I$. By Claim 5.50, I_k has size $U_{k,t}(|J|)$. Since $k = \text{const}(t)$, this is equivalent to $U_t(|J|)$, as desired.

Otherwise, the construction terminates with $I := I_i$ and $S := S_i$ for some $i \leqslant k$, as either (C.1) or (C.2) holds. By the same reasoning as before, we have $|I| \geqslant U_t(|J|)$. By Claim 5.50, we have $|S| \leqslant k$ and $S \subseteq V(G) \setminus \mathcal{A}|_I$. In case (C.1), we have a prepattern of order t on $\mathcal{A}|_I$ and in case (C.2) S samples G on $\mathcal{A}|_I$ with margin 2.

As $k = \text{const}(t)$, the running time of $O_t(|V(G)|^2)$ for the construction follows easily from the running times of Lemma 5.42 and Fact 5.49. □

5.6.3 Sample Sets for Orderless Insulators

For orderless insulators we want to strengthen Lemma 5.38 by improving the guarantees given by the sampling set. For convenience, we restate the definition of a sampling set.

Definition 5.36. Let G be a graph containing an insulator $\mathcal{A}$ with grid A indexed by I. Let $v, s_<, s_>$ be vertices from G, $i \in I$, and $m \in \mathbb{N}$. We say v *is* $(m, i, s_<, s_>)$*-sampled on* $\mathcal{A}$ if

$$\text{atp}(v/A[<i, *]) = \text{atp}(s_</A[<i, *]) \quad \text{and} \quad \text{atp}(v/A[\geqslant i+m, *]) = \text{atp}(s_>/A[\geqslant i+m, *]).$$

We call m the *margin*, i the *exceptional index*, $s_<$ the *left-sample*, and $s_>$ the *right-sample*.

Definition 5.37. Fix $p \in \mathbb{N}$. Let G be a graph containing an insulator $\mathcal{A}$ indexed by I and let $S \subseteq V(G)$. We say S *samples* G on $\mathcal{A}$ with margin m if there exists functions $\text{ex} : V(G) \to I$ and $s_<, s_> : V(G) \to S$ such that every $v \in V(G)$ is $\big(m, \text{ex}(v), s_<(v), s_>(v)\big)$-sampled on $\mathcal{A}$.

We say S *symmetrically samples* G on $\mathcal{A}$, if we can choose $s_< = s_>$ in the above definition. For orderless insulators, we want to decrease the sampling margin to 1 and make the sampling symmetric.

Lemma 5.52. *Let G be a graph containing an orderless insulator $\mathcal{A}$ indexed by J, and let $S \subseteq V(G)$ be a set that samples G on $\mathcal{A}$ with margin 2. Let $I \subseteq J$ be obtained by removing every second element of J. S samples G on $\mathcal{A}|_I$ with margin 1.*

Proof. The exceptional indices of each vertex are successive. By only keeping every second column, we reduce the number of exceptional indices to at most 1, which corresponds to margin 1. □

Lemma 5.53. *Fix $t \in \mathbb{N}$. For every graph G and orderless insulator $\mathcal{A}$ indexed by J in G, we can compute a subsequence $I \subseteq J$ of size $U_t(|J|)$ and a set $S \subseteq V(G) \setminus \mathcal{A}|_I$ of size $\text{const}(t)$ such that either*

- *G contains a prepattern of order t on $\mathcal{A}|_I$, or*
- *$\mathcal{A}|_I$ is orderable, or*
- *S symmetrically samples G on $\mathcal{A}|_I$ with margin 1.*

Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes the sequence I and one of the three outcomes (a prepattern, witnesses for $\mathcal{A}|_I$ being orderable, or a set S) in time $O_t(|V(G)|^2)$.

Proof. We first apply Lemma 5.38 to $\mathcal{A}$ and J, which yields a sequence I_0 of length $U_t(J)$ and a set $S \subseteq V(G) \setminus \mathcal{A}|_{I_0}$ of size const(t) such that either

- G contains a prepattern of order t on $\mathcal{A}|_{I_0}$, or
- S samples G on $\mathcal{A}|_{I_0}$ with margin 2.

In the first case we are done by setting $I := I_0$, so assume the second case. Lemma 5.52 yields a sequence $I_1 \subseteq I_0$ of length $U_t(|I_0|) = U_t(|J|)$ such that S samples G on $\mathcal{A}|_{I_1}$ with margin 1. Let B be the grid of $\mathcal{B} := \mathcal{A}|_{I_1}$. We color every element $i \in \text{tail}(I_1)$ by a color that encodes for all $s_1, s_2 \in S$ the information whether

$$N(s_1) \cap B[i,*] \subseteq N(s_2) \cap B[i,*]. \tag{5.2}$$

This requires $|S|^2 = \text{const}(t)$ many colors. Inducing tail(I_1) on the largest color class yields a monochromatic subsequence $I_2 \subseteq \text{tail}(I_1)$ of length $U_t(|I_1|) = U_t(|I_0|) = U_t(|J|)$, where we can interpret monochromaticity as follows.

Fact 5.54. *For every $s_1, s_2 \in S$, if $N(s_1) \cap B[i,*] \subseteq N(s_2) \cap B[i,*]$ for one $i \in I_2$, then it holds for every $i \in I_2$. In particular, if two vertices $s_1, s_2 \in S$ have the same neighborhood in one column $B[i,*]$ for some $i \in I_2$ then they have the same neighborhood in every column $B[i,*]$ with $i \in I_2$.*

By definition, $\mathcal{A}|_{I_2}$ consists of the columns $\{B_{i,*} : i \in \text{tail}(I_2)\}$. By Lemma 5.41, S also samples G on $\mathcal{A}|_{I_2}$ with margin 1. Whenever we have two vertices in S that have the same neighborhood in every column of $\mathcal{A}|_{I_2}$ we can remove one of them from S while still preserving that S samples G on $\mathcal{A}|_{I_2}$ with margin 1. Thus, by Fact 5.54, for distinct vertices $s_1, s_2 \in S$ and for every $i \in I_2$ we have $N(s_1) \cap B[i,*] \neq N(s_2) \cap B[i,*]$. Together with Fact 5.54 we obtain the following.

Fact 5.55. *For distinct vertices $s_1, s_2 \in S$, either*

- *for all $i \in I_2$: $\big(N(s_1) \cap B[i,*]\big) \not\subseteq \big(N(s_2) \cap B[i,*]\big)$, or*
- *for all $i \in I_2$: $\big(N(s_1) \cap B[i,*]\big) \not\supseteq \big(N(s_2) \cap B[i,*]\big)$.*

Recall that a vertex $v \in V(G)$ is *s-connected* to a set $U \subseteq V(G)$ if atp$(s/U) = $ atp(v/U). In the next step we color every 4-element subsequence $\bar{\iota} = (\iota_1, \ldots, \iota_4) \subseteq I_2$. In order to apply Ramsey's theorem, we encode for every 4-tuple $\bar{s} = s_1 \ldots s_4 \in S^4$ the information whether

$$\exists v \bigwedge_{i \in [4]} v \text{ is } s_i\text{-connected to } B[\iota_i, *]. \tag{5.3}$$

Ramsey's Theorem (Fact 4.13) yields a monochromatic subsequence $I_3 \subseteq I_2$ of length $U_t(|J|)$, where we can interpret monochromaticity as follows.

Fact 5.56. *For every $\bar{s} \in S^4$, whenever there exists a 4-element subsequence $\bar{\iota}$ of I_3 that satisfies* (5.3)*, then every 4-element subsequence $\bar{\iota}$ of I_3 satisfies* (5.3).

By Lemma 5.41, S still samples G on $\mathcal{A}|_{I_3}$ with margin 1. Therefore, we can choose for every vertex $v \in V(G)$ two samples $s_<(v), s_>(v) \in S$ and an exceptional index ex(v) such that

- v is $s_<(v)$-connected to all $B[i,*]$ with $i \in I_3$ and $i <$ ex(v), and
- v is $s_>(v)$-connected to all $B[i,*]$ with $i \in I_3$ and $i >$ ex(v).

We remove the first and last two elements of I_3 to obtain I_4. If for every vertex v with ex$(v) \in I_4$ we have $s_<(v) = s_>(v)$, then S symmetrically samples G on $\mathcal{A}|_{I_4}$, and we can complete the proof by setting $I := I_4$. Assume therefore, there is a vertex $v_\star$ with ex$(v_\star) \in I_4$ and $s_<(v_\star) \neq s_>(v_\star)$. Let $s_1 := s_<(v_\star)$ and $s_2 := s_>(v_\star)$.

Claim 5.57. *For every $i \in I_4$ there exists a vertex c_i such that*

- *c_i is s_1-connected to all $B[j,*]$ with $j \in I_3$ and $j < i$, and*
- *c_i is s_2-connected to all $B[i,*]$ with $j \in I_3$ and $j > i$.*

Proof. Let ι_1, ι_2 and ι_3, ι_4 be the two immediate predecessors and successors of $\mathrm{ex}(v_\star)$ in I_3. Those are distinct indices, and they exist, since $\mathrm{ex}(v_\star) \in I_4$ and I_4 was obtained from I_3 by removing the first and last two elements. Therefore, $\bar{\iota} := (\iota_1, \iota_2, \iota_3, \iota_4)$ is a 4-element subsequence of I_3. Additionally, $v_\star$ is s_1-connected to $B[\iota_1,*], B[\iota_2,*]$ and s_2-connected to $B[\iota_3,*], B[\iota_4,*]$

Pick $i \in I_4$, and let ι'_1, ι'_2 and ι'_3, ι'_4 be the two immediate predecessors and successors of i in I_3. It follows by Fact 5.56 that there is a vertex c_i that is s_1-connected to $B[\iota'_1,*], B[\iota'_2,*]$ and s_2-connected to $B[\iota'_3,*], B[\iota'_4,*]$. Since S samples G on $\mathcal{A}|_{I_3}$ with margin 1, we have that $\mathrm{ex}(c_i) \in \{\iota'_2, i, \iota'_4\}$. Since c_i is s_1-connected to $B[\iota'_1,*]$, it is s_1-connected to all $B[j,*]$ with $j \in I_3$ and $j \leqslant \iota'_1$. A symmetric statement holds for s_2 and $j \geqslant \iota'_4$. This proves that c_i has the desired properties. ■

Consider Fact 5.55 for s_1 and s_2. We can assume $\forall i \in I_2 : \big(N(s_1) \cap B[i,*]\big) \not\supseteq \big(N(s_2) \cap B[i,*]\big)$, as the alternative case will follow by a symmetric argument. It follows that for every $i \in I_2$ there exists a vertex $b_i \in B[i,*]$ such that $b_i \in N(s_2) \setminus N(s_1)$. Now by Claim 5.57, we have that for all $i, j \in I_4$

- c_i is non-adjacent to b_j if $j < i$, and
- c_i is adjacent to b_j if $j > i$.

By a simple majority argument we find a subsequence $I_5 \subseteq I_4$ of length at least $\frac{1}{2}|I_4| = U_t(|J|)$ such that either

- c_i and b_j are adjacent if and only if $j \geqslant i$ for all $i, j \in I_5$, or
- c_i and b_j are adjacent if and only if $i > j$ for all $i, j \in I_5$.

By possibly dropping the first element from I_5 and shifting the indices of the c_i by one, we can always assume the first case applies. Now $(b_i)_{i \in I_5}$ and $(c_i)_{i \in I_5}$ witness that $\mathcal{A}|_{I_5}$ is orderable as desired, so we can set $I := I_5$.

Running Time. Let $n := |V(G)|$. Using the same preprocessing as in the run time analysis of Lemma 5.42, we can assume that $|J| \leqslant n^{1/4}$. We first apply Lemma 5.38, which runs in time $O_t(n^2)$. Coloring of the elements of $\mathrm{tail}(I_1)$ and building I_2 can be done in time $O_t(n)$. We then color the 4-tuples of the resulting sequence. Similarly to the proof of Lemma 5.42, this can be done in time $O_t(n^2)$. As we have ensured $|J| \leqslant n^{1/4}$, applying Ramsey to this coloring takes time $O_t(n)$. For the resulting sequence we can search for the c_i-vertices of Claim 5.57. We do this by testing for every vertex $v \in V(G)$ the role of which of the c_i it can take. For a single vertex this can be done in time $O_t(n)$ by checking its connections to each column of the insulator. An exhaustive search over all vertices in G therefore takes time $O_t(n^2)$. If we find a suitable candidate for every $i \in I_4$, we can compute I_5 in time $O_t(n)$, which yields the desired orderable insulator. If we cannot find suitable candidates for all c_i, we can conclude that S has the desired sampling property. The overall running time is $O_t(n^2)$ as desired. □

5.7 Extending Insulators

We are now ready to prove the insulator growing lemmas. We first prove the orderless and then the ordered case.

5.7.1 Extending Orderless Insulators

Lemma 5.19 (Orderless Insulator Growing). *Fix $k, t \in \mathbb{N}$. For every graph G and orderless insulator $\mathcal{A}$ indexed by J of cost k in G, we can compute a subsequence $I \subseteq J$ of length $U_t(|J|)$ such that either*

- *G contains a prepattern of order t on $\mathcal{A}|_I$,*
- *$\mathcal{A}|_I$ is orderable, or*
- *there exists a row-extension of $\mathcal{A}|_I$ of cost $\mathrm{const}(k, t)$ in G.*

Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes the sequence I and one of the three outcomes (a prepattern, witnesses for $\mathcal{A}|_I$ being orderable, or a row-extension) in time $O_{k,t}(|V(G)|^2)$.

Proof. Let $\mathcal{A} = (A, \mathcal{K}, F, F)$. Apply Lemma 5.53 to $\mathcal{A}$, which yields a subsequence $I \subseteq J$ of size $U_t(|J|)$ and a set $S \subseteq V(G) \setminus \mathcal{A}|_I$ of size const(t) such that either

- G contains a prepattern of order t on $\mathcal{A}|_I$, or
- $\mathcal{A}|_I$ is orderable, or
- S symmetrically samples G on $\mathcal{A}|_I$ with margin 1.

In the first two cases we are done, so we assume the last case: there exist functions $s : V(G) \to S$ and ex : $V(G) \to \mathrm{tail}(I)$ such that every $v \in V(G)$ is $(1, \mathrm{ex}(v), s(v), s(v))$-sampled on $\mathcal{A}|_I$. Let A and $B := A|_I$ be the grids of $\mathcal{A}$ and $\mathcal{B} := \mathcal{A}|_I$. Let $I_\star := \mathrm{tail}(I)$ be the sequence indexing B and let h be the height of A and B.

Defining the Grid. We build a row-extension C of B. By definition of a row-extension, we have $C[i, j] := B[i, j]$ for all $i \in I_\star$ and $j \in [h]$. It remains to define the row $C[*, h + 1]$. For every $i \in I_\star$, we define $C[i, h + 1]$ to contain every vertex v that

- is not contained in B, and
- disagrees with its sample in the cell below, that is, $\mathrm{atp}(v/C[i, h]) \neq \mathrm{atp}(s(v)/C[i, h])$.

As every vertex v is sampled with margin 1 in $\mathcal{B}$, v can disagree with $s(v)$ in at most one column of B, so no vertex gets assigned into multiple columns. Furthermore, we only assign vertices to $C[*, h + 1]$ which were not in B, so the cells of C are pairwise disjoint and C is a valid grid. Thus, C is a row-extension of B.

Claim 5.58. *For every $i \in I_\star$ and $v \in C[i, *]$, we have $\mathrm{ex}(v) = i$.*

Proof. As v is $(1, \mathrm{ex}(v), s(v), s(v))$-sampled on $\mathcal{B}$, we have

$$\mathrm{atp}(v/B[i, *]) \neq \mathrm{atp}(s(v)/B[i, *]) \Rightarrow \mathrm{ex}(v) = i.$$

If $v \in C[i, h + 1]$, then the premise is satisfied by construction. If $v \in C[i, \leqslant h] = B[i, *]$, then, since $s(v) \notin B[i, *]$, the premise is again satisfied as we have

$$(=, v) \in \mathrm{atp}(v/B[i, *]) \setminus \mathrm{atp}(s(v)/B[i, *]). \qquad \blacksquare$$

The rest of the proof will be devoted to constructing $\mathcal{K}_\star$ and $F_\star$ such that the row-extension $\mathcal{C} := (C, \mathcal{K}_\star, F_\star, F_\star)$ is an insulator of cost const(k, t) in G.

Defining the Insulator. We build $\mathcal{K}_\star$ as a refinement of $\mathcal{K}$ by encoding into the color of every vertex $v \in V(G)$ for every color $X \in \mathcal{K}$ and sample vertex $s \in S$ the information

(C.1) whether $v \in X$,

(C.2) whether $v \in C$,

(C.3) whether $v \in C[*, h]$,
(C.4) whether $v \in C[*, h+1]$,
(C.5) whether $v \in N(s)$,
(C.6) whether $s(v) = s$.

As $|\mathcal{K}| = k$ and $|S| = \text{const}(t)$, we have $|\mathcal{K}_\star| = \text{const}(k, t)$. By (C.6), (C.5), and (C.1), we can assign to every color $X \in \mathcal{K}_\star$

- a sample vertex $s(X) \in S$, such that $s(v) = s(X)$ for all $v \in X$,
- sample neighbors $S(X) \subseteq S$, such that $N(v) \cap S = S(X)$ for all $v \in X$, and
- a color $\mathcal{K}(X) \subseteq \mathcal{K}$, such that $\mathcal{K}(v) = \mathcal{K}(X)$ for all $v \in X$.

In order to show that $\mathcal{C}$ is an insulator in G, it remains to define the symmetric relation $F_\star \subseteq \mathcal{K}_\star^2$ such that property (U.1) is satisfied. We define $F_\star$ via the following four cases. Let $X, Y \in \mathcal{K}_\star$.

(F.1) If $X \subseteq C[*, h]$ and $Y \subseteq C[*, h]$, then $(X, Y) \in F_\star \Leftrightarrow \big(s(Y) \in S(X) \vee s(X) \in S(Y)\big)$.
(F.2) If $X \not\subseteq C[*, \leqslant h]$ and $Y \subseteq C[*, h]$, then $(X, Y) \in F_\star \Leftrightarrow s(X) \in S(Y)$.
(F.3) If $X \subseteq C[*, h]$ and $Y \not\subseteq C[*, \leqslant h]$, then $(X, Y) \in F_\star \Leftrightarrow s(Y) \in S(X)$.
(F.4) Otherwise, $(X, Y) \in F_\star \Leftrightarrow \big(\mathcal{K}(X), \mathcal{K}(Y)\big) \in F$.

By construction, $F_\star$ is symmetric and therefore describes a valid flip. Let $G' := G \oplus F$ and $G_\star := G \oplus F_\star$.

Proving Properties of Insulator. We have to show (U.1): for all $i \in I_\star$ there exists $a_i \in V(G)$ such that for all $r \in [h+1]$

$$C[i, 1] = N_0^{G_\star}[a_i] = \{a_i\} \quad \text{and} \quad C[i, \leqslant r] = N_{r-1}^{G_\star}[a_i]. \tag{5.4}$$

We first show that our flip conserves this property for $r \in [h]$, and handle $r = h+1$ later.

Claim 5.59. *For all $i \in I_\star$ and $r \in [h]$, we have*

$$C[i, \leqslant r] = B[i, \leqslant r] = N_{r-1}^{G'}[a_i] = N_{r-1}^{G_\star}[a_i].$$

Proof. The first two equalities follow by construction and property (U.1) of $\mathcal{B}$. It remains to prove $N_{r-1}^{G'}[a_i] = N_{r-1}^{G_\star}[a_i]$. We prove the claim by induction on r. The base case is trivial. For the inductive step, assume the property holds for $r \in [h-1]$ and we want to show it for $r+1$. We show that for every vertex u we have $u \in N_r^{G_\star}[a_i]$ if and only if $u \in N_r^{G'}[a_i]$. We can assume $u \notin N_{r-1}^{G_\star}[a_i] = N_{r-1}^{G'}[a_i]$, as we would be done by induction otherwise. With these prerequisites, the following are equivalent.

1. $u \in N_r^{G_\star}[a_i]$.
2. u has a neighbor in $N_{r-1}^{G_\star}[a_i]$ in $G_\star$. (as $u \notin N_{r-1}^{G_\star}[a_i]$)
3. u has a neighbor in $C[i, \leqslant r]$ in $G_\star$. (by induction)

Let v be a vertex in $C[i, \leqslant r] \subseteq C[i, <h]$. By (C.2), (C.3), and (C.4), we have $\mathcal{K}_\star(v) \subseteq C[i, <h]$ and case (F.4) applies: v has the same neighborhood in $G_\star$ as in G'. Hence, the following are equivalent to the above.

4. u has a neighbor in $C[i, \leqslant r]$ in G'.
5. u has a neighbor in $N_{r-1}^{G'}[a_i]$ in G'. (by induction)
6. $u \in N_r^{G'}[a_i]$. (as $u \notin N_{r-1}^{G'}[a_i]$)

■

Having proved Claim 5.59, in order to establish (U.1), it remains to prove

$$C[i, h+1] = N_h^{G_\star}[a_i] \setminus N_{h-1}^{G_\star}[a_i].$$

We show the equivalence of the two sets by proving containment in both directions separately.

Claim 5.60. *For all $i \in I_\star$ we have $C[i, h+1] \subseteq N_h^{G_\star}[a_i] \setminus N_{h-1}^{G_\star}[a_i]$.*

Proof. Let $u \in C[i, h+1]$. As $u \notin C[i, \leqslant h]$, we have by Claim 5.59 that $u \notin N_{h-1}^{G_\star}[a_i]$. It remains to show $u \in N_h^{G_\star}[a_i]$. By construction, we have

$$\text{atp}(u/C[i,h]) \neq \text{atp}(s(u)/C[i,h]).$$

As neither u nor $s(u)$ is contained in B, the difference in their atomic type must be witnessed by a vertex $v \in C[i, h]$ in the symmetric difference of their neighborhoods. We want to argue that u and v are adjacent in $G_\star$. Let $X := \mathcal{K}_\star(u)$ and $Y := \mathcal{K}_\star(v)$. By (C.4) and (C.3), we have $X \subseteq C[i, h+1]$ and $Y \subseteq C[i, h]$. Hence, case (F.2) from the construction of $F_\star$ applies and the following are equivalent.

1. The adjacency between u and v was flipped when going from G to $G_\star$.
2. $s(X) \in S(Y)$. (by (F.2))
3. $s(u) \in N(v) \cap S$. (by definition)
4. $s(u)$ is a neighbor of v in G. (by definition)
5. u is a non-neighbor of v in G. (v is in the sym. diff. of $N(s(u))$ and $N(u)$)

The equivalence between the first and the last item establishes that u and v are adjacent in $G_\star$. By Claim 5.59, we have $v \in N_{h-1}^{G_\star}[a_i]$, so $u \in N_h^{G_\star}[a_i]$ and the claim is proved. ■

Claim 5.61. *For all $i \in I_\star$ we have $C[i, h+1] \supseteq N_h^{G_\star}[a_i] \setminus N_{h-1}^{G_\star}[a_i]$.*

Proof. Let u be a vertex in $N_h^{G_\star}[a_i] \setminus N_{h-1}^{G_\star}[a_i]$. By Claim 5.59, this is witnessed by a vertex

$$v \in N_{h-1}^{G_\star}[a_i] = N_{h-1}^{G'}[a_i] = C[i,h]$$

that is adjacent to u in $G_\star$. We prove that $u \in C[i, h+1]$ by ruling out all other possibilities.

- Assume that $u \in C[j, \leqslant h-1]$ for some $j \in I_\star$. By Claim 5.59, we have

  $$u \in N_{h-2}^{G_\star}[a_j] = N_{h-2}^{G'}[a_j].$$

  As we assumed that $u \notin N_{h-1}^{G_\star}[a_i]$, we have that $j \neq i$. Additionally, case (F.4) applies and u has the same neighborhood in G' and $G_\star$. Hence, u is adjacent to v also in G'. Now $v \in N_{h-1}^{G'}(a_j) = C[j, \leqslant h]$, but still $v \in C[i, h]$. This is a contradiction to $C[i, h]$ and $C[j, \leqslant h]$ being disjoint.
- Assume that $u \notin C$ or $u \in C[j, h+1]$ for some $j \neq i \in I_\star$. We first show

  $$\text{atp}(u/C[i,h]) = \text{atp}(s(u)/C[i,h]). \tag{5.5}$$

  If $u \notin C$, we deduce (5.5) from the construction of C. If $u \in C[j, h+1]$, we apply Claim 5.58, which yields $j = \text{ex}(u) \neq i$. We then deduce (5.5) by the sampling property.

Let $X := \mathcal{K}_\star(u)$ and $Y := \mathcal{K}_\star(v)$. By our choice of u and v, we know their adjacency in $G_\star$ was determined by case (F.2) and the following are equivalent.

1. The adjacency between u and v was flipped when going from G to $G_\star$.
2. $s(X) \in S(Y)$. (by (F.2))
3. $s(u) \in N(v) \cap S$. (by definition)
4. v is adjacent to $s(u)$ in G. (by definition)
5. v is adjacent to u in G. (by (5.5))

The equivalence between the first and the last item establishes that u and v are non-adjacent in $G_\star$, a contradiction.

- Finally, we assume that $u \in C[j,h]$ for some $j \in I_\star$. As we know that $u \notin N^{G_\star}_{h-1}[a_i]$, Claim 5.59 yields $i \neq j$. Then Claim 5.58 applied to $u \in C[j,h]$ and $v \in C[i,h]$ yields

$$j = \text{ex}(u) \neq i = \text{ex}(v),$$

which together with the sampling property gives

$$\text{atp}(u/C[i,h]) = \text{atp}(s(u)/C[i,h]) \quad \text{and} \quad \text{atp}(v/C[j,h]) = \text{atp}(s(v)/C[j,h]). \tag{5.6}$$

Let $X := \mathcal{K}_\star(u)$ and $Y := \mathcal{K}_\star(v)$. By our choice of u and v, we know their adjacency in $G_\star$ was determined by case (F.1) and the following are equivalent.

1. The adjacency between u and v was flipped when going from G to $G_\star$.
2. $s(Y) \in S(X)$ or $s(X) \in S(Y)$. (by (F.1))
3. $s(v) \in N(u) \cap S$ or $s(u) \in N(v) \cap S$. (by definition)
4. u is a neighbor of $s(v)$ in G or v is a neighbor of $s(u)$ in G. (by definition)
5. u is a neighbor of v in G or v is a neighbor of u in G. (by (5.6))
6. u is a neighbor of v in G.

The equivalence between the first and the last item establishes that u and v are non-adjacent in $G_\star$, a contradiction.

Having exhausted all other possibilities, we conclude that $u \in C[i,h+1]$, which proves the claim. ■

The combination of Claim 5.59, Claim 5.60, and Claim 5.61 proves property (U.1). Hence, $\mathcal{C} := (C, \mathcal{K}_\star, F_\star, F_\star)$ is an insulator of cost const(k,t) in G. This proves that $\mathcal{C}$ is the desired row extension. It remains to analyze the running time.

Running Time. Let $n := |V(G)|$. The application of Lemma 5.53 runs in time $O_t(n^2)$. In the case where a sample set S is returned, we can calculate the witnessing function s in time $O_t(n^2)$ by comparing each vertex $v \in V(G)$ with each of the const(t) many vertices from S over every column of $\mathcal{A}|_I$. With the function s at hand, we can also build the row-extension C of B in time $O_t(n^2)$. The construction of $\mathcal{K}_\star$ and $F_\star$ runs in time $O_{k,t}(n)$. This yields an overall running time of $O_{k,t}(n^2)$. □

5.7.2 Extending Ordered Insulators

Lemma 5.20 (Ordered Insulator Growing). *Fix $k, t \in \mathbb{N}$. For every graph G and ordered insulator $\mathcal{A}$ with cost k, indexed by J in G, we can compute a subsequence $I \subseteq J$ of length $U_t(|J|)$ such that either*

- *G contains a prepattern of order t on $\mathcal{A}|_I$, or*
- *G contains a row-extension of $\mathcal{A}|_I$ with cost $\mathrm{const}(k,t)$.*

Moreover, there is an algorithm that, given G and $\mathcal{A}$, computes the sequence I and one of the two outcomes (a prepattern or a row-extension) in time $O_{k,t}(|V(G)|^2)$.

Proof. We first apply Lemma 5.38 to $\mathcal{A}$, which yields a subsequence $I \subseteq J$ of size $U_t(|J|)$ and a set $S \subseteq V(G) \setminus \mathcal{A}|_I$ of size $\mathrm{const}(t)$ such that either

- G contains a prepattern of order t on $\mathcal{A}|_I$, or
- S samples G on $\mathcal{A}|_I$ with margin 2.

In the first case we are done, so assume the second case. By possibly taking a subsequence and applying Lemma 5.41, we can assume the following.

Property 5.62. I does not contain the first and last two elements of J.

Let B be the grid of $\mathcal{B} := \mathcal{A}|_I$ and let $I_\star := \mathrm{tail}(I)$ be the sequence indexing B. As S samples G on $\mathcal{B}$ with margin 2, there exist functions $\mathrm{ex} : V(G) \to I_\star$ and $s_<, s_> : V(G) \to S$, such that every $v \in V(G)$ is $(2, \mathrm{ex}(v), s_<(v), s_>(v))$-sampled in $\mathcal{B}$. We can assume ex is chosen maximal in the following sense: for every vertex $v \in V(G)$ and index $i \in I_\star$, if $\mathrm{atp}(v/B[\leqslant i, *]) = \mathrm{atp}(s_<(v)/B[\leqslant i, *])$, then $\mathrm{ex}(v) \geqslant i$.

Defining the Grid. We will build a row-extension C of B. By definition of a row-extension we have $C[i,r] := B[i,r]$ for all $i \in I_\star$ and $r \in [h]$. It remains to define the row $C[*, h+1]$. For every $i \in I_\star$, we set $C[i, h+1]$ to contain every vertex v such that

- v is not contained in B, and
- $\mathrm{atp}(v/B[i,h]) \neq \mathrm{atp}(s_<(v)/B[i,h])$ and i is the minimal index in $I_\star$ with this property.

It is easy to see that cells of C are pairwise disjoint and C is a valid grid. Furthermore, we have the following property.

Claim 5.63. *For every $i \in I_\star$ and $v \in C[i,*]$, we have $\mathrm{ex}(v) \in \{i-1, i\}$.*

Proof. If $v \in C[i, \leqslant h] = B[i,*]$, then, since $B[i,*]$ contains v but neither $s_<(v)$ nor $s_>(v)$, the atomic type of v differs from the atomic types of both $s_<(v)$ and $s_>(v)$ over $B[i,*]$. By the sampling property with margin 2, we have that $i = \mathrm{ex}(v)$ or $i = \mathrm{ex}(v)+1$.

If $v \in C[i, h+1]$, then by construction the atomic type of v differs from the atomic type of $s_<(v)$ over $B[i,*]$. This yields $\mathrm{ex}(v) \leqslant i$. As we have chosen i minimal, we have

$$\mathrm{atp}(v/B[\leqslant i-1, *]) = \mathrm{atp}(s_<(v)/B[\leqslant i-1, *]).$$

As we have chosen ex maximal, we have $\mathrm{ex}(v) \geqslant i-1$. ■

The rest of the proof will be devoted to constructing $\mathcal{K}_\star$, $F_\star$, and $R_\star$ such that the row-extension $\mathcal{C} := (C, \mathcal{K}_\star, F_\star, R_\star)$ is an insulator of cost $\mathrm{const}(k,t)$ in G.

Defining the Insulator. We build $\mathcal{K}_\star$ as a refinement of $\mathcal{K}$ by encoding into the color of every vertex $v \in V(G)$ for every color $X \in \mathcal{K}$ and sample vertex $s \in S$ the information

(C.1) whether $v \in X$,

(C.2) whether $v \in C$,

(C.3) whether $v \in C[*, h]$,

(C.4) whether $v \in C[*, h+1]$,

(C.5) whether $v \in N(s)$,

(C.6) whether $v \in \text{int}(B)$,

(C.7) whether $s_<(v) = s$,

(C.8) whether $s_>(v) = s$,

(C.9) whether v has a neighbor in $X \cap \text{int}(B)$.

As $\mathcal{K}$ has size k and S has size const(t), we have $|\mathcal{K}_\star| = \text{const}(k, t)$. By (C.7), (C.8), (C.5), and (C.1), we can define for every color $X \in \mathcal{K}_\star$

- a left sample $s_<(X) \in S$, such that for all $v \in X$ we have $s_<(v) = s_<(X)$,
- a right sample $s_>(X) \in S$, such that for all $v \in X$ we have $s_>(v) = s_>(X)$,
- sample neighbors $S(X) \subseteq S$, such that for all $v \in X$ we have $N(v) \cap S = S(X)$, and
- a color $\mathcal{K}(X) \in \mathcal{K}$, such that for all $v \in X$ we have $\mathcal{K}(v) = \mathcal{K}(X)$.

We define $F_\star \subseteq \mathcal{K}_\star^2$ via the following four cases. Let $X, Y \in \mathcal{K}_\star$.

(F.1) If $X \subseteq C[*, h+1]$ and $Y \subseteq C[*, h]$, then $(X, Y) \in F_\star \Leftrightarrow s_<(X) \in S(Y)$.

(F.2) If $X \subseteq C[*, h]$ and $Y \subseteq C[*, h+1]$, then $(X, Y) \in F_\star \Leftrightarrow s_<(Y) \in S(X)$.

(F.3) If $X \subseteq C[*, h+1]$ and $Y \subseteq C[*, {<}h]$, or vice-versa, then

$$(X, Y) \in F_\star \Leftrightarrow \text{there is an edge between } X \text{ and } Y \text{ in } G.$$

(F.4) Otherwise, $(X, Y) \in F_\star \Leftrightarrow \big(\mathcal{K}(X), \mathcal{K}(Y)\big) \in F$.

In order for $F_\star$ to define a valid flip, $F_\star$ has to be symmetric. This is satisfied by our definition. The cases (F.1) and (F.2) are dual and for (F.3) and (F.4) the symmetry follows from the symmetry of their conditions and the symmetry of the edge relation and F.

We define $R_\star \subseteq \mathcal{K}_\star^2$ via the following three cases.

(R.1) If $X \subseteq C[*, h+1]$ and $Y \subseteq C[*, {\leqslant}h]$, or

(R.2) if $X \subseteq C[*, h]$ and $Y \subseteq C[*, h]$, then

$$(X, Y) \in R_\star \Leftrightarrow s_>(X) \in S(Y).$$

(R.3) Otherwise, $(X, Y) \in R_\star \Leftrightarrow \big(\mathcal{K}(X), \mathcal{K}(Y)\big) \in R$.

Proving Properties of the Insulator. Let $G' := G \oplus F$ and $G_\star := G \oplus F_\star$. We prove the required insulator properties.

Claim 5.64 ((O.1))**.** *Every two vertices in different rows of C have different colors in $\mathcal{K}_\star$.*

Proof. If none of the two vertices is in $C[*, h+1]$, then the property holds as $\mathcal{K}_\star$ is a refinement of $\mathcal{K}$, which satisfied this property in B. Otherwise, exactly one of them is contained in $C[*, h+1]$, and we can distinguish them using (C.4). ■

Claim 5.65 ((O.2))**.** *Every vertex $u \in C[i, r]$ with $r > 1$ has a neighbor in the cell $C[i, r-1]$ in $G_\star$.*

Proof. If $r \leqslant h$, then by (O.2) of $\mathcal{B}$, there is a vertex $v \in C[i, r-1]$ that is adjacent to u in G'. By (F.4), u and v are also adjacent in $G_\star$.

It remains to check the case where $u \in C[i, h+1]$. By construction there exists a vertex $v \in C[i, h]$ in the symmetric difference of $N(u)$ and $N(s_<(u))$. By (C.4) and (C.3), we have $\mathcal{K}_\star(u) \subseteq C[*, h+1]$ and $\mathcal{K}_\star(v) \subseteq C[*, h]$. Case (F.1) applies, and the following are equivalent.

1. The adjacency between u and v was flipped from G to $G_\star$.
2. $s_<(\mathcal{K}(u)) \in S(\mathcal{K}(v))$. (by (F.1))
3. $s_<(u) \in N(v) \cap S$. (by definition)
4. $s_<(u)$ is adjacent to v in G. (by definition)
5. u is non-adjacent to v in G. (v is in the sym. diff. of $N(s_<(u))$ and $N(u)$)

The equivalence between the first and the last item establishes that u and v are adjacent in $G_\star$. ■

Claim 5.66 ((O.3))**.** *For every $v \notin C$ and $X \in \mathcal{K}_\star$, v is homogeneous to $X \cap \text{int}(C)$ in G.*

Proof. Let $X_C := X \cap \text{int}(C)$. As C is an extension of B, we also have $v \notin B$. By construction of C and (C.6) and (C.3), we have either $X_C \subseteq \text{int}(B)$ or $X_C \subseteq C[*, h] = B[*, h]$. In the first case we conclude by (O.3) of $\mathcal{B}$. In the second case, since v did not get sorted into C, we know by construction

$$\text{atp}(v/C[*, h]) = \text{atp}(s_<(v)/C[*, h]).$$

By (C.5), $s_<(v)$ is homogeneous to X, and so is v. ■

In the following claims we prove (O.4). For this purpose let $u \in C[i, r]$ for some $i \in I_\star$ and $r < h+1$, and let $v \in C$. Up to renaming, we additionally assume $I_\star = (1, \dots, n)$.

Claim 5.67 ((O.4.1))**.** *If $u \in \text{int}(C)$ and u and v are in rows that are not close, then they are non-adjacent in $G_\star$.*

Proof. Assume first $v \in C[*, h+1]$. Then $u \in C[*, <h]$, as u and v are in rows that are not close. By construction of C this yields $u \in \text{int}(B)$ and $v \notin B$. Using the properties of our coloring, we conclude the following. Let $X := \mathcal{K}_\star(u)$ and $Y := \mathcal{K}_\star(v)$.

- $X \subseteq \mathcal{K}(X) \cap \text{int}(B)$. (by (C.1) and (C.6))
- No vertex of Y is contained in B. (by (C.2))
- Every or no vertex in Y has a neighbor in $\mathcal{K}(X) \cap \text{int}(B)$. (by (C.9))

Combining the above with (O.3) of $\mathcal{B}$, we know that the connection between X and Y is homogeneous in G. Additionally, $X \subseteq C[*, <h]$ and $Y \subseteq C[*, h+1]$ so case (F.3) applies and the following are equivalent.

1. The adjacency between u and v was flipped from G to $G_\star$.
2. There is an edge between X and Y in G. (by (F.3))
3. There is an edge between u and v in G. (as $X \ni u$ and $Y \ni v$ are homogeneous)

The equivalence between the first and the last item establishes that u and v are non-adjacent in $G_\star$.

If $v \notin C[*, h+1]$, then since also $u \notin C[*, h+1]$, case (F.4) applies and u and v have the same adjacency in $G_\star$ as in G'. Note that in this case, by construction and Observation 5.8, u and v are both contained in both grids A and B and in both grids, they are in rows that are not close. If $v \in C[*, h]$, then again $u \in \text{int}(B)$ and u and v non-adjacent in G' by (O.4.1) of $\mathcal{B}$. Otherwise, $v \in C[*, <h]$ and by Property 5.62 and Observation 5.8 we have $v \in \text{int}(A)$. Now u and v non-adjacent in G' by (O.4.1) of $\mathcal{A}$. ■

Claim 5.68 ((O.4.2)). *If $v \in C[<i, r-1] \cup C[>i, r+1]$, then u and v are non-adjacent in $G_\star$.*

Proof. If at least one of u and v is contained in $C[*, <h] = B[*, <h]$, then case (F.4) applies and u and v have the same adjacency in $G_\star$ as in G'. Possibly exchanging the roles of u and v, we can apply (O.4.2) of $\mathcal{B}$ to deduce that u and v are non-adjacent in G'.

In the remaining case we have $u \in C[i, h] = B[i, h]$ and $v \in C[>i, h+1]$ and by construction of C:

$$\text{atp}(v/B[i, h]) = \text{atp}(s_<(v)/B[i, h]).$$

Also, $\mathcal{K}_\star(u) \subseteq C[*, h]$ and $\mathcal{K}_\star(v) \subseteq C[*, h+1]$ by (C.3) and (C.4). Hence, case (F.2) applies, and the following are equivalent.

1. The adjacency between u and v was flipped from G to $G_\star$.
2. $s_<(\mathcal{K}_\star(v)) \in S(\mathcal{K}_\star(u))$. (by (F.2))
3. $s_<(v) \in N(u) \cap S$. (by definition)
4. u and $s_<(v)$ are adjacent in G. (by definition)
5. u and v are adjacent in G. (by $\text{atp}(v/B[i, h]) = \text{atp}(s_<(v)/B[i, h])$)

The equivalence of the first and the last item establishes that u and v are non-adjacent in $G_\star$. ■

Claim 5.69 ((O.4.3)). *If $v \in C[>i+1, \{r, r-1\}]$, then $G \models E(u, v) \Leftrightarrow (\mathcal{K}_\star(u), \mathcal{K}_\star(v)) \in R_\star$.*

Proof. If either $u \in C[*, < h]$ or $v \in C[*, < h]$, then case (R.3) applies, we have

$$(\mathcal{K}_\star(u), \mathcal{K}_\star(v)) \in R_\star \Leftrightarrow (\mathcal{K}(u), \mathcal{K}(v)) \in R,$$

and it remains to establish

$$G \models E(u, v) \Leftrightarrow (\mathcal{K}(u), \mathcal{K}(v)) \in R.$$

If $u \in C[*, <h]$, we argue using (O.4.3) of $\mathcal{B}$. Otherwise, we have $v \in C[*, <h]$ and argue using (O.4.4) of $\mathcal{B}$, where we exchange the roles of u and v.

We can now assume $u, v \in C[*, h]$. By (C.3), also $\mathcal{K}_\star(u), \mathcal{K}_\star(v) \subseteq C[*, h]$, and case (R.2) applies. By assumption, we have $v \in C[i', h]$ for $i + 1 < i'$. By Claim 5.63, we have that $\text{ex}(u) + 1 < i'$, and the following are equivalent.

1. u and v are adjacent in G.
2. $s_>(u)$ and v are adjacent in G. (by the sampling property)
3. $s_>(u) \in N(v) \cap S$. (by definition)
4. $s_>(\mathcal{K}_\star(u)) \in S(\mathcal{K}_\star(v))$. (by definition)
5. $(\mathcal{K}_\star(u), \mathcal{K}_\star(v)) \in R_\star$. (by (R.2))

■

Claim 5.70 ((O.4.4)). *If $v \in C[<i-1, \{r, r+1\}]$, then $G \models E(u, v) \Leftrightarrow (\mathcal{K}_\star(v), \mathcal{K}_\star(u)) \in R_\star$.*

Proof. If either

- one of u and v is contained in $C[*, <h]$, or
- both u and v are contained in $C[*, h]$,

then we can exchange u and v and the property follows from the already established property (O.4.3). It remains to prove the case where $u \in C[i,h]$ and $v \in C[<i-1,h+1]$. We have $\mathcal{K}_\star(u) \subseteq C[*,h]$ and $\mathcal{K}_\star(v) \subseteq C[*,h+1]$, and case (R.1) applies for $X = \mathcal{K}_\star(v)$ and $Y = \mathcal{K}_\star(v)$. By assumption, we have $v \in C[i', h+1]$ for $i'+1 < i$. By Claim 5.63, we have that $\text{ex}(v)+1 < i$, and the following are equivalent.

1. u and v are adjacent in G.
2. u and $s_>(v)$ are adjacent in G. (by the sampling property)
3. $s_>(v) \in N(u) \cap S$. (by definition)
4. $s_>(\mathcal{K}_\star(v)) \in S(\mathcal{K}_\star(u))$. (by definition)
5. $(\mathcal{K}_\star(v), \mathcal{K}_\star(u)) \in R_\star$. (by (R.1))

■

This proves property (O.4). Finally, property (O.5) only concerns the first row $C[*,1] = B[*,1]$, so its truth carries over from $\mathcal{B}$. Having proven all properties, it follows that $\mathcal{C}$ is an insulator. Its cost is $|\mathcal{K}_\star| = \text{const}(k,t)$. This proves that $\mathcal{C}$ is the desired row extension. It remains to analyze the running time.

Running Time. Let $n := |V(G)|$. Lemma 5.38 runs in time $O_t(n^2)$. Similarly, as in the proof of Lemma 5.19, we can build the row-extension C of B in time $O_t(n^2)$. The construction of $\mathcal{K}_\star$, $F_\star$, and $R_\star$ runs in time $O_{k,t}(n)$. This yields an overall running time of $O_{k,t}(n^2)$. □

Having proven the insulator growing lemmas, this concludes Chapter 5.

Chapter 6

Forbidden Induced Subgraphs

In Chapter 5, we have shown that for any graph class the following implications hold:

$$\text{prepattern-free} \Rightarrow \text{insulation-property} \Rightarrow \text{flip-breakable} \Rightarrow \text{mon. dependent}$$

In this chapter, we close the circle by showing the following two remaining implications:

$$\neg\text{prepattern-free} \Rightarrow \text{large flipped crossings/comparability grids} \Rightarrow \text{mon. independent}$$

This amounts to showing the equivalence (1) ⇔ (2) ⇔ (3) of our characterizations of monadically dependent graph classes (Theorem 2.3). The biggest part of the work will be to prove the first remaining implication as made explicit in the following proposition. The definition of crossings and comparability grids can be found in Section 2.2.

Proposition 6.1. *Let $\mathcal{C}$ be a graph class that is not prepattern-free. Then there exists $r \in \mathbb{N}$ such that for every $k \in \mathbb{N}$, $\mathcal{C}$ contains as an induced subgraph either*

- *a flipped star r-crossing of order k, or*
- *a flipped clique r-crossing of order k, or*
- *a flipped half-graph r-crossing of order k, or*
- *a comparability grid of order k.*

We prove this proposition in Section 6.1 and 6.2. In Section 6.3 we prove the second remaining implication: classes exhibiting the patterns listed in Proposition 6.1 are monadically independent. There we also summarize the results obtained so far (Proposition 6.47).

6.1 Transformers

As a first step, in Section 6.1, starting from large prepatterns, we will extract well-structured, but rather abstract objects called *transformers*. After that, in Section 6.2, transformers will be analyzed in detail, and crossings will be extracted from them.

6.1.1 Meshes and Transformers

Definition 6.2. A *mesh* in a graph G is an injective function $M: I \times J \to V(G)$, where I and J are two non-empty indexing sequences of the same length. We denote $V(M) := \{M(i,j) : i \in I, j \in J\}$. For a mesh M as above, by $M^\top$ denote the mesh $M^\top: J \times I \to V(G)$ such that $M^\top(i,j) = M(j,i)$ for all $i \in I, j \in J$. A mesh M has *order* m if $|I| = |J| = m$. For $I' \subseteq I$ and $J' \subseteq J$, by $M|_{I' \times J'}$ we denote the mesh obtained by restricting M to $I' \times J'$. We call $M|_{I' \times J'}$ a *submesh* of M.

Definition 6.3. Let $M\colon I \times J \to V(G)$ be a mesh in a (possibly colored) graph G. Then M is *vertical* in G if $|I| = |J| \leqslant 3$, or if there is a function $a\colon I \to V(G)$ such that

- $\text{atp}_G(M(i,j), a(i'))$ depends only on $\text{otp}(i,i')$ for all $i,i' \in I$ and $j \in J$, and
- $\text{atp}_G(M(i,j), a(i'))$ is *not the same for all* $i,i' \in I$ and $j \in J$.

We say that M is *horizontal* in G if M^{T} is vertical in G. Note that a mesh can be both horizontal and vertical.

Remark 6.4. We often write that some property or function $P(\bar{x})$ *"is the same for all tuples $\bar{x}$ from a given domain."* This means that for all $\bar{x}, \bar{y}$ (from a specified domain) $P(\bar{x}) = P(\bar{y})$.

Definition 6.5. Let $M, M'\colon I \times J \to V(G)$ be meshes in a (possibly colored) graph G. We say that in G the pair (M, M') is

- *regular* if $\text{atp}_G(M(i,j), M'(i',j'))$ depends only on $\text{otp}(i,i')$ and $\text{otp}(j,j')$ for all $i,i' \in I$ and $j,j' \in J$,
- *homogeneous* if $\text{atp}_G(M(i,j), M'(i',j'))$ is the same for all $i,i' \in I$ and $j,j' \in J$,
- *conducting* if either $|I| = |J| \leqslant 3$ or (M, M') is regular but not homogeneous in G.

Lemma 6.6. *Let (M, M') be a regular pair of meshes in a colored graph G. Then all vertices in M have the same color and all vertices in M' have the same color.*

Proof. Let $M, M'\colon I \times J \to V(G)$ be as in the statement. The colors of $M(i,j)$ and $M'(i,j)$ are encoded in $\text{atp}(M(i,j), M'(i,j))$ which only depends on $\text{otp}(i,i)$ and $\text{otp}(j,j)$ by regularity. As both order types are always $=$ for all $i \in I$ and $j \in J$, the lemma follows. ☐

Definition 6.7. A *conductor* of order n and length h is a sequence $M_1, \ldots, M_h\colon I \times J \to V(G)$ of meshes of order n, such that each pair (M_s, M_{s+1}) is conducting for $s = 1, \ldots, h-1$.

We now define the central notion of Section 6.1.

Definition 6.8. A *transformer* of order n and length h is a conductor $M_1, \ldots, M_h$ of order n and length h, such that M_1 is vertical, and M_h is horizontal.

The main result of Section 6.1 is the following.

Proposition 6.9. *Let G be a graph containing a prepattern of order n on an insulator of height h and cost k. Then G contains a transformer of order $U_{h,k}(n)$ and length at most $4h-1$.*

We start with some simple observations.

Observation 6.10. Let M be a mesh in a graph G. If M is horizontal/vertical in G, then this also holds for every submesh of M.

Observation 6.11. Let $M_1, M_2\colon I \times J \to V(G)$ be meshes in a graph G where the pair (M_1, M_2) is conducting. For all $I' \subseteq I$ and $J' \subseteq J$ with $|I'| = |J'|$ the pair $(M_1|_{I' \times J'}, M_2|_{I' \times J'})$ is also conducting in G.

Lemma 6.12. *Let G be a graph and G^+ be a coloring of G. If a pair of meshes (M, M') is conducting in G^+, then it is also conducting in G.*

Proof. Obviously the pair (M, M') is also regular in G. Via contrapositive, it remains to assume that (M, M') is homogeneous in G, and show that it also is homogeneous in G^+. To this end, observe that the atomic type $\text{atp}_{G^+}(M(i,j), M'(i',j'))$ depends only on $\text{atp}_G(M(i,j), M'(i',j'))$ as well as the colors of $M(i,j)$ and $M'(i',j')$. The former has the desired properties by homogeneity in G, and the latter by Lemma 6.6. ☐

Lemma 6.13. *Let G be a graph and G^+ be a coloring of G. If a mesh M is vertical in G^+ then it is also vertical in G.*

Proof. Let $a(\cdot)$ be the function witnessing that M is vertical in G^+. We argue as in the proof of Lemma 6.12. Here, instead of regularity, we use verticality to argue that all vertices in the range of a (respectively all vertices in the range of M) have the same atomic type in G^+. □

6.1.2 From Prepatterns to Transformers

In this subsection, we prove Proposition 6.9, extracting transformers from prepatterns. Let us first give an overview of the proof in the case where G contains a bi-prepattern. The mono-prepattern case later falls out as a subcase. The bi-prepattern in G is witnessed by an insulator $\mathcal{A}$ whose columns are indexed by a sequence K containing two subsequences I and J. Every pair of columns $(i, j) \in I \times J$ is "matched up" by an element $c_{i,j}$ through quantifier-free formulas (cf. Definition 5.10).

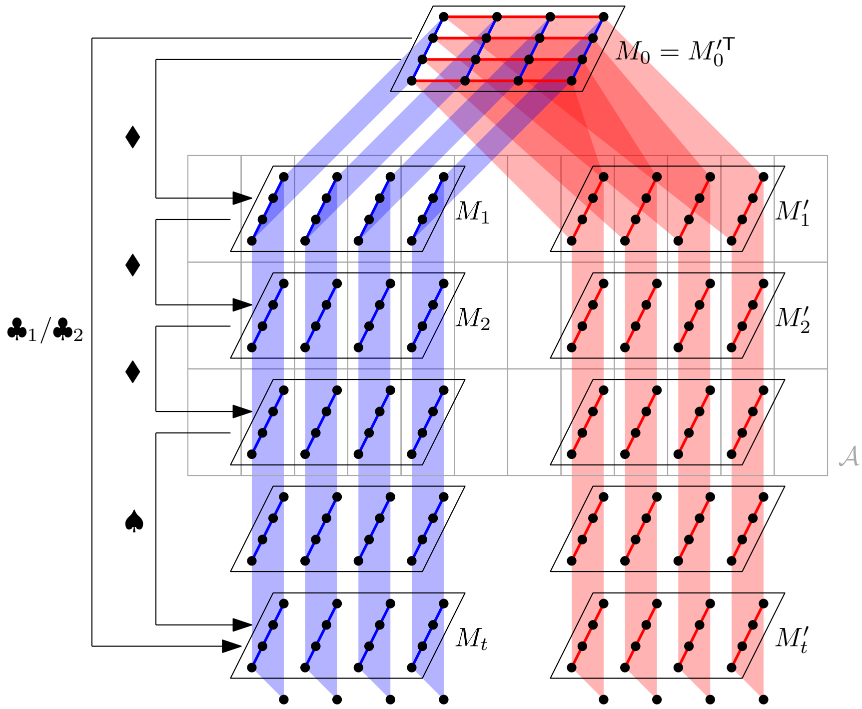

Figure 6.1: Constructing a transformer from a bi-prepattern.

Figure 6.1 is a schematic depiction of how the transformer will embed into the insulator $\mathcal{A}$. The transformer will be assembled from two conductors $C = M_0, \ldots, M_t$ and $C' = M_0', \ldots, M_t'$. The columns of the meshes of C and C' are contained in the columns of the insulator $\mathcal{A}$ indexed by I and J, respectively. The structure of $\mathcal{A}$ imposes that M_t and M_t' are both vertical. The meshes M_0 and M_0' are chosen from the vertices $c_{i,j}$ such that $M_0 = M_0'^{\mathsf{T}}$. It follows that we can glue the conductor C to the transposed meshes from C' yielding the desired transformer

$$M_t, M_{t-1}, \ldots, M_1, M_0 = M_0'^{\mathsf{T}}, M_1'^{\mathsf{T}}, \ldots, M_{t-1}'^{\mathsf{T}}, M_t'^{\mathsf{T}},$$

where M_t is vertical and $M_t'^{\mathsf{T}}$ is horizontal.

The construction of C and C' is implemented by Lemma 6.19 or Lemma 6.21 (♣$_1$ or ♣$_2$ in the picture) depending on the choice of $\sim_1$ and $\sim_2$ in the definition of a bi-prepattern. Lemma 6.19 iteratively extends the conductor by "stepping down" the insulator. A single step is performed using Lemma 6.15 (♦), which creates meshes *embedded into* descending rows of the insulator (cf. Definition 6.14). The structure of the insulator is used to establish conductivity between successive meshes/rows. This process continues until either a vertical mesh is produced or we reach the bottom row of the insulator. If we reach the bottom row, we use Lemma 6.18 (♠) to further extend the conductor to reach a vertical mesh. Here we use the special structure of the bottom row: In the unordered case, the cells how the bottom row contain only a single vertex each, which implies that the mesh in the bottom row is already vertical. In the ordered case the mesh in the bottom row is connected through short parts to a vertical mesh and we can bridge those short paths by a conductor. This finishes the sketch of Lemma 6.19 ($\clubsuit_1$). Lemma 6.21 ($\clubsuit_2$) is a rather technical case distinction, which reduces the construction of the conductor to some previously established subcase.

This concludes the overview of the proof of Proposition 6.9. We now give a detailed proof.

Definition 6.14. Let G be a graph containing an insulator $\mathcal{A}$ of height h with grid A and indexed by a sequence K. Let $M\colon I \times J \to V(G)$ be a mesh of order n in G with $I \subseteq K$. Let $r \in [h]$. We say that M *embeds into row* r *of* $\mathcal{A}$ if for all $i \in I$ and $j \in J$.

$$M(i,j) \in A[i,r].$$

Lemma 6.15 (♦ in Figure 6.1)**.** *Let G be a graph containing an insulator $\mathcal{A}$ of height h with grid A and indexed by a sequence K. Let $M\colon I \times J \to V(G)$ be a mesh of order n in G with $I \subseteq K$. Let $\alpha(x,y)$ be a quantifier-free formula in a k-coloring G^+ of G, and let $r \in [h]$ be such that for all $i \in I, j \in J$*

$$i = \min\{i' \in I : \exists v \in A[i', \leqslant r] \ : \ G^+ \models \alpha(v, M(i,j))\}. \tag{6.1}$$

Then there are meshes M_1, M_2 of order $U_{k,h}(n)$ such that

- *M_1 is a submesh of M, and*
- *the pair (M_1, M_2) is conducting in G, and*
- *M_2 is vertical in G or M_2 embeds into some row $r' \leqslant r$ of $\mathcal{A}$.*

Proof. Note that by Lemma 6.12 and Lemma 6.13, it does not matter whether we show conductivity and verticality in G or G^+. For every $i \in I$ and $j \in J$, pick a vertex $M'(i,j) \in A[i, \leqslant r]$ such that $G^+ \models \alpha(M'(i,j), M(i,j))$. Note that M is not necessarily a mesh, as it may not be injective.

Claim 6.16. *There are sequences $I' \subseteq I$ and $J' \subseteq J$ with $|I'| = |J'| \geqslant U_{k,h}(|I|) = U_{k,h}(n)$, and a row number $r' \in [r]$ such that*

(R.1) $\mathrm{atp}_{G^+}(M(i,j), M(i',j'), M'(i,j), M'(i',j'))$ *depends only on* $\mathrm{otp}(i,i')$ *and* $\mathrm{otp}(j,j')$ *for $i, i' \in I'$ and $j, j' \in J'$, and*

(R.2) $M'(i,j) \in A[i, r']$ *for all $i \in I'$.*

The proof of the claim is a straightforward application of Bipartite Ramsey (Lemma 4.15). Readers experienced in Ramsey theory are invited to skip it. However, due to the importance of Ramsey type arguments for this section, we include the details of an exemplary application here.

Proof of Claim 6.16. Up to renaming we can assume $I = J = [n]$. Let Π be the set of possible atomic types of four tuples in k-colored graphs. There exists a constant $k^* \leqslant \mathrm{const}(k,h)$ such that there is a bijection $b\colon \Pi \times [r] \to [k^*]$. We now define the coloring $c\colon [n]^2 \times [n]^2 \to [k^*]$ as

$$c\big((i,i'),(j,j')\big) := b\Big(\mathrm{atp}_{G^+}\big(M(i,j), M(i',j'), M'(i,j), M'(i',j')\big),\ r'\Big)$$

where r' is the unique row such that $M'(i,j) \in A[i,r']$. Applying Bipartite Ramsey (Lem. 4.15) to the defined coloring yields sequences $I' \subseteq I$ and $J' \subseteq J$ with $|I'| = |J'| \geqslant U_{k,h}(|I|) = U_{k,h}(n)$ such that $c\big((i,i'),(j,j')\big)$ depends only on $\mathrm{otp}(i,i')$ and $\mathrm{otp}(j,j')$ for $i,i' \in I'$ and $j,j' \in J'$. By the construction of the coloring, this proves (R.1). To see that the row containing $M'(i,j)$ is the same for all $i \in I$ and $j \in J$, notice that

$$c\big((i,i),(j,j)\big) = c\big((i',i'),(j',j')\big) \quad \text{for all } i,i' \in I' \text{ and } j,j' \in J'.$$

Therefore, $M(i,j)$ and $M(i',j')$ are in the same row, which proves (R.2). ■

Let I', J', and r' be as from the previous claim. We set $M_1 := M|_{I' \times J'}$ and $M_2 := M'|_{I' \times J'}$.

Claim 6.17. *Either M_1 is vertical, or M_2 is a mesh (that is, M_2 is injective).*

Proof. Assume $|I'| > 3$, as otherwise M_1 is vertical by definition. Assume that M_2 is not injective. Then there exist distinct pairs $(i,j),(i',j') \in I' \times J'$ such that $M_2(i,j) = M_2(i',j')$. Since $M_2(i,j)$ (equivalently $M_2(i',j')$) is from the ith (equivalently i'th) column of A, we have $i = i'$. Hence, $j \neq j'$.

By (R.1) we have $a(i) := M_2(i,j) = M_2(i,j')$ for all $i \in I'$ and $j,j' \in J'$. Now also by (R.1), $\mathrm{atp}_{G^+}(M_1(i,j), a(i'))$ depends only on $\mathrm{otp}(i,i')$ for all $i,i' \in I$ and $j \in J'$. Finally, by (6.1) and $|I'| > 3$, $\mathrm{atp}_{G^+}(M_1(i,j), a(i'))$ is not the same for all $i,i' \in I$ and $j \in J'$. In summary: $a(\cdot)$ witnesses that M_1 is vertical. ■

Assume now the mesh M_1 is vertical. By (R.1), the pair (M_1, M_1) is regular. As witnessed by the equality type, the pair is not homogeneous and therefore conducting. We can return (M_1, M_1).

Otherwise, M_2 is a mesh and embeds into the row r' of $\mathcal{A}$ by (R.2). Furthermore, the pair (M_1, M_2) is conducting: regularity follows from (R.1) and non-homogeneity from (6.1). □

Lemma 6.18 (♠ in Figure 6.1)**.** *Let G be a graph containing an insulator $\mathcal{A}$ of cost k and height h and let M be a mesh of order n that embeds into row 1 of $\mathcal{A}$. There exists a conductor $M_1, \ldots, M_t$ for $t \leqslant h$ of order $U_{k,h}(n)$ such that M_1 is a submesh of M and M_t is vertical.*

Proof. Let I, J be the sequences indexing the mesh M. Let A be the grid of $\mathcal{A}$. Assume first that $\mathcal{A}$ is orderless. Then by the insulator property (U.1), each cell in row 1 of $\mathcal{A}$ contains only a single vertex. As M embeds into this row, it has order 1. Then M is vertical and the conductor consisting only of M satisfies the conditions of the lemma.

Assume now that $\mathcal{A}$ is ordered. By (O.5), there exists a k-flip H of G and some radius $r < h$ such that the r-balls around the vertices in $A[*,1]$ are pairwise disjoint. Moreover, there are vertices $\{b(v) \in N_r^H[v] : v \in A[*,1]\}$ and $\{c_i \in V(G) : i \in I\}$ and a symbol $\sim \in \{\leqslant, \geqslant\}$ such that for all $i,j \in I$ and $v \in A[j,1]$

$$G \models E(c_i, b(v)) \text{ if and only if } i \sim j.$$

Let G^+ be the k-coloring of G in which the edge relation of the flipped graph H can be expressed by a quantifier-free formula. For each $i \in I$ and $j \in J$ let $\pi(i,j)$ be the tuple of vertices forming a shortest path from $M(i,j)$ to $b(M(i,j))$ in H. By (O.5), these paths have equal length for all i,j and consist of at most h vertices. By Bipartite Ramsey (Lemma 4.15), there exist sequences $I' \subseteq I$ and $J' \subseteq I$ of length at least $U_{k,h}(n)$ such that

$$\mathrm{atp}_{G^+}(\pi(i,j), c(i), \pi(i',j'), c(i')) \text{ depends only on } \mathrm{otp}(i,i') \text{ and } \mathrm{otp}(j,j')$$
$$\text{for all } i,i' \in I' \text{ and } j,j' \in J'.$$

For distinct pairs (i, j) and (i', j') in $I' \times J'$, $\pi(i, j)$ and $\pi(i', j')$ have no vertex in common, as they stem from two disjoint balls in H. Therefore, each of the functions $M_1, \ldots, M_h \colon I' \times J' \to V(G)$, where $M_t(i, j)$ is defined as the tth component of $\pi(i, j)$, is injective and forms a mesh. By construction, M_1 is a submesh of M. By Ramsey and (O.5), M_h is vertical. We next show that $M_1, \ldots, M_h$ is a conductor in G^+. By Ramsey, any pair of successive meshes (M_t, M_{t+1}) in the sequence is regular in G^+. It remains to show that (M_t, M_{t+1}) is not homogeneous in G^+. Consider two distinct pairs $(i, j), (i', j') \in I' \times J'$. As $\pi(i, j)$ and $\pi(i', j')$ are constructed via paths through disjoint balls in H, we have that $M_t(i, j)$ is adjacent to $M_{t+1}(i, j)$ and non-adjacent to $M_{t+1}(i', j')$ in H. It follows that (M_t, M_{t+1}) is not homogeneous in G^+. Therefore, (M_t, M_{t+1}) is conducting in G^+. Hence, $M_1, \ldots, M_h$ is a conductor in G^+, and by Lemma 6.12 and Lemma 6.13 also in G. □

Lemma 6.19 (♣$_1$ in Figure 6.1). *Let G be a graph containing an insulator $\mathcal{A}$ of cost k and height h with grid A and indexed by a sequence K. Let $M : I \times J \to V(G)$ be a mesh of order n in G with $I \subseteq K$. Let $\alpha(x, y)$ be a quantifier-free formula in a k-coloring G^+ of G such that for all $i \in I, j \in J$*

$$i = \min\{i' \in I : \exists v \in A[i', *] \ : \ G^+ \models \alpha(v, M(i, j))\}.$$

Then there is a conductor $M_1, \ldots, M_t$ of order $U_{k,h}(n)$ and length at most $2h$ in G such that M_1 is a submesh of M and M_t is vertical in G.

Proof. Note that by Lemma 6.12 and Lemma 6.13, it does not matter whether we show conductivity and verticality in G or G^+. Denote $M_0 := M$ and $r_0 := h + 1$. We inductively construct a sequence of meshes $M_1, M_2, \ldots, M_t$ where for each $s = 1, 2, \ldots, t$, the mesh M_s satisfies the following conditions:

- M_s has order $U_{s,k,h}(n)$,
- there is a row $r_s \in [r_{s-1} - 1]$ such that M_s embeds into row r_s of $\mathcal{A}$,
- there is a submesh M'_{s-1} of M_{s-1} such that the pair (M'_{s-1}, M_s) is conducting.

To construct M_1, we apply Lemma 6.15 to the mesh M_0 and the formula $\alpha(x, y)$ to obtain a conducting pair of meshes (M'_0, M_1) of order $U_{k,h}(n)$ where M'_0 is a submesh of M_0. If M_1 is vertical then M'_0, M_1 is the desired conductor, and we conclude the proof of the lemma. Otherwise, M_1 embeds into some row $r_1 \in [h]$ of $\mathcal{A}$ and satisfies the induction hypothesis.

Suppose the sequence $M_1, \ldots, M_s$ has been constructed for some $s \geqslant 1$. Below we either extend the sequence by one, or terminate the process.

Assume first $r_s > 1$. As the sequence $r_1, r_2, \ldots, r_s$ is strictly decreasing we have $s < h$. Let $\beta(x, y)$ be the formula expressing adjacency in the flip G' associated to $\mathcal{A}$. As G' is a k-flip of G, β is expressible in a k-coloring of G. Let I_s and J_s be the sequences indexing M_s. For every $i_0 < i \in I_s$ and $j \in J_s$, we have that $M_s(i, j)$ has a β-neighbor in $A[i, r_s - 1]$ but no β-neighbor in $A[i_0, r_s - 1]$. This is by the insulator property (U.1) if $\mathcal{A}$ is orderless and by (O.2) and (O.4.2) if $\mathcal{A}$ is ordered. Hence, we can apply Lemma 6.15 to the mesh M_s, the row number $r_s - 1$, and the formula $\beta(x, y)$. We obtain a conducting pair of meshes (M'_s, M_{s+1}) of order $U_{s,k,h}(n)$ where M'_s is a submesh of M_s. If M_{s+1} is vertical then we conclude the proof of the lemma returning the conductor

$$M_0|_{I' \times J'}, M_1|_{I' \times J'}, \ldots, M_s|_{I' \times J'}, M_{s+1}$$

of length at most $h + 1$, where $I' \subseteq I$ and $J' \subseteq J$ are the indexing sequences of M_{s+1}. Otherwise, M_{s+1} embeds into some row $r_{s+1} \in [r_s - 1]$ of $\mathcal{A}$ and satisfies the induction hypothesis. We continue the process.

Assume now $r_s = 1$. Note that $s \leqslant h$. Then M_s embeds into the first row of $\mathcal{A}$, and we can apply Lemma 6.18. This yields a conductor $M^\star_1, \ldots, M^\star_{t^\star}$ of length $t^\star \leqslant h$ indexed by sequences

$I' \subseteq I$ and $J' \subseteq J$ of length $U_{s,k,h}(n)$ such that $M_1^\star$ is a submesh of M_s, and $M_{t^\star}^\star$ is vertical. Now we conclude the proof of the lemma returning the following conductor of length at most $2h$:

$$M_0|_{I'\times J'}, \ldots, M_{s-1}|_{I'\times J'}, M_1^\star, \ldots, M_{t^\star}^\star. \qquad \square$$

Lemma 6.20. *Assume G contains a mono-prepattern of order n on an insulator $\mathcal{A}$ of cost k and height h. Then G contains a transformer of order $U_{k,h}(n)$ and length at most $2h$.*

Proof. Let $(c_j : j \in J)$ and $(b_{i,j} : i \in I, j \in J)$ form a mono-prepattern of order $n = |I| = |J|$ on the insulator $\mathcal{A}$. Define the function $M \colon I \times J \to V(G)$ where $M(i,j) = b_{i,j}$. By definition of a mono-prepattern, all the $b_{i,j}$ are distinct, so M is a mesh. By Bipartite Ramsey (Lemma 4.15) and Definition 5.11, there exist sequences $I' \subseteq I$ and $J' \subseteq I$ of length at least $U(n)$ such that $\mathrm{atp}_G(M(i,j), c_{j'})$ depends only on $\mathrm{otp}(j,j')$ for all $i \in I'$ and $j, j' \in J'$. Since $\mathrm{atp}_G(M(i,j), c_{j'})$ is not the same for all $i \in I'$ and $j, j' \in J'$, the submesh $M|_{I'\times J'}$ is horizontal.

We apply Lemma 6.19 to the mesh $M|_{I'\times J'}$ and formula $\alpha(x,y) := (x = y)$. This yields a conductor $M_1, \ldots, M_t$ of length at most $2h$ and order $U_{h,k}(n)$ such that M_1 is a submesh of $M|_{I'\times J'}$, and M_t is vertical. As $M|_{I'\times J'}$ is horizontal, the same holds for M_1, and thus $M_t, \ldots, M_1$ is the desired transformer. $\square$

Lemma 6.19 is complemented by the following more technical Lemma 6.21. Together, these two lemmas accommodate the two possible choices for the symbols $\sim_1, \sim_2 \in \{=, \neq\}$ in the definition of a bi-prepattern (Definition 5.10).

Lemma 6.21 (♣$_2$ in Figure 6.1)**.** *Let $\mathcal{A}$ be an insulator of cost k and height h indexed by a sequence K, let I, J be indexing sequences with $I \subseteq K$, let $M \colon I \times J \to V(G)$ be a mesh of order n, and let α be a quantifier-free formula in a k-coloring G^+ of G, such that for all $i \in I, j \in J$*

$$i = \min\{i' \in I : \neg\exists v \in A[i', *] \ : \ G^+ \models \alpha(v, M(i,j))\}.$$

Then there is a conductor $M_1, \ldots, M_t$ of order $U_{k,h}(n)$ and length at most $2h$ such that either

- *M_1 is a submesh of M and M_t is vertical in G, or*
- *$M_1, \ldots, M_t$ is a transformer in G.*

Proof. For all $i^*, i \in I$ with $i^* < i$ and $j \in J$, fix a vertex $F(i^*, i, j) \in A[i^*, *]$ such that

$$G^+ \models \alpha(F(i^*, i, j), M(i,j)).$$

By Bipartite Ramsey (Lemma 4.15), there are sets $I' \subseteq I$ and $J' \subseteq J$ with $|I'| = |J'| \geqslant U(|I|)$, such that for all $i^*, i, i' \in I'$ with $i^* < i, i'$, and for all $j, j' \in J'$,

$$\mathrm{atp}_{G^+}(F(i^*, i', j'), M(i,j)) \quad \text{depends only on } \mathrm{otp}(i,i') \text{ and } \mathrm{otp}(j,j'). \qquad (*)$$

In particular this holds for the fact whether $G^+ \models \alpha(F(i^*, i', j'), M(i,j))$. For an exemplary application of Bipartite Ramsey that illustrates how to obtain $(*)$, see the proof of Claim 6.16. Denote

- $i_{\min} := \min(I')$,
- $i_{\max} := \max(I')$,
- $i'_{\max} := \max(I' \setminus i_{\max})$,
- $j_{\min} := \min(J')$,
- $j_{\max} := \max(J')$,
- $j'_{\max} := \max(J' \setminus j_{\max})$.

For convenience, we redefine $I := I' \setminus \{i_{\min}, i_{\max}, i'_{\max}\}$, and $J := J' \setminus \{i_{\min}, i_{\max}, i'_{\max}\}$. We do so to ensure that all elements in I and J have the same order type with respect to the previously chosen extremal elements.

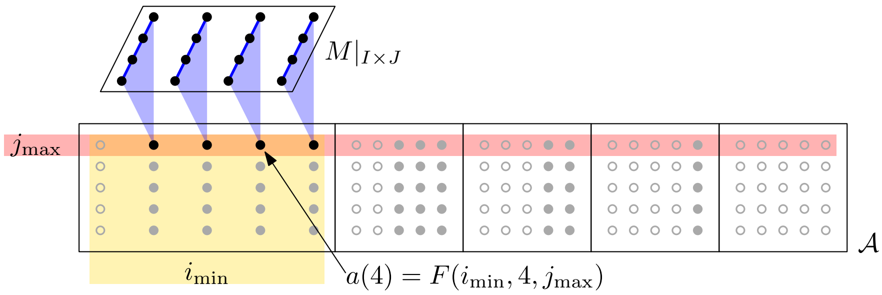

Figure 6.2: A depiction of *Case 1*. On top: the desired vertical submesh of M. On the bottom: the insulator $\mathcal{A}$, whose columns contain the $F(i^*, i, j)$ vertices (depicted as $\bullet$). Here, the outermost columns correspond to i^*-coordinates of $F(i^*, i, j)$, and within each outermost column, the inner columns and rows correspond to i- and j-coordinates, respectively. The $F(i^*, i, j)$ vertices are only defined for indices $i^* < i$, thus going rightwards the columns are filled up with placeholders (depicted as $\circ$).

Case 1: Assume $\mathrm{atp}_{G^+}(F(i^*, i', j'), M(i, j))$ is not the same for all $i^*, i, i' \in I$ and $j, j' \in J$ with $i^* < i, i'$ and $j < j'$ (here $j < j'$ is the crucial assumption). In this case, we work towards proving that M contains a large vertical submesh. The situation is depicted in Figure 6.2. We set $a(i') := F(i_{\min}, i', j_{\max})$ for all $i' \in I$. Let us now argue that $\mathrm{atp}_{G^+}(a(i'), M(i, j))$ is not the same for all $i, i' \in I$ and $j \in J$.

By assumption, there exist indices

- $i_1^*, i_1, i_1' \in I$ and $j_1, j_1' \in J$ with $i_1^* < i_1, i_1'$ and $j_1 < j_1'$, and
- $i_2^*, i_2, i_2' \in I$ and $j_2, j_2' \in J$ with $i_2^* < i_2, i_2'$ and $j_2 < j_2'$

such that

$$\mathrm{atp}_{G^+}(F(i_1^*, i_1', j_1'), M(i_1, j_1)) \neq \mathrm{atp}_{G^+}(F(i_2^*, i_2', j_2'), M(i_2, j_2)).$$

As we removed the extremal elements from I and J, we have

- $\mathrm{otp}(i_1^*, i_1, i_1') = \mathrm{otp}(i_{\min}, i_1, i_1')$ and $\mathrm{otp}(j_1, j_1') = \mathrm{otp}(j_1, j_{\max})$, as well as
- $\mathrm{otp}(i_2^*, i_2, i_2') = \mathrm{otp}(i_{\min}, i_2, i_2')$ and $\mathrm{otp}(j_2, j_2') = \mathrm{otp}(j_2, j_{\max})$.

Now applying $(*)$, we get

$$\mathrm{atp}_{G^+}(\underbrace{F(i_{\min}, i_1', j_{\max})}_{a(i_1')}, M(i_1, j_1)) \neq \mathrm{atp}_{G^+}(\underbrace{F(i_{\min}, i_2', j_{\max})}_{a(i_2')}, M(i_2, j_2)).$$

Thus, $\mathrm{atp}_{G^+}(a(i'), M(i, j))$ is not the same for all $i, i' \in I$ and $j \in J$. Additionally, as $j < j_{\max}$ for all $j \in J$, we have that $\mathrm{atp}_{G^+}(a(i'), M(i, j))$ only depends on $\mathrm{otp}(i, i')$ and no longer on j, for all $i, i' \in I$ and $j \in J$. It follows that $M_{I\times J}$ is vertical and forms the desired conductor (in this case of length one).

Case 2: Assume $\mathrm{atp}_{G^+}(F(i^*, i', j'), M(i, j))$ is not the same for all $i^*, i, i' \in I$ and $j, j' \in J$ with $i^* < i, i'$ and $j > j'$ (here $j > j'$ is the crucial assumption). We proceed as in the previous case, but with $j_{\min}$ instead of $j_{\max}$.

Case 3: Assume $\mathrm{atp}_{G^+}(F(i^*, i', j'), M(i, j))$ is not the same for all $i^*, i, i' \in I$ and $j, j' \in J$ with $i^* < i' < i$ (here $i' < i$ is the crucial assumption). Now let

$$c(j) := M(i_{\max}, j) \quad \text{and} \quad b(i^*, j') := F(i^*, i'_{\max}, j') \in A[i^*, *] \quad \text{for all } i^* \in I \text{ and } j \in J.$$

Our goal is to show that the ranges of $b(\cdot, \cdot)$ and $c(\cdot)$ form a mono-prepattern on $\mathcal{A}$. The situation is depicted in Figure 6.3. As in *Case 1*, we argue via $(*)$ that for all $i^* \in I$ and $j, j' \in J$

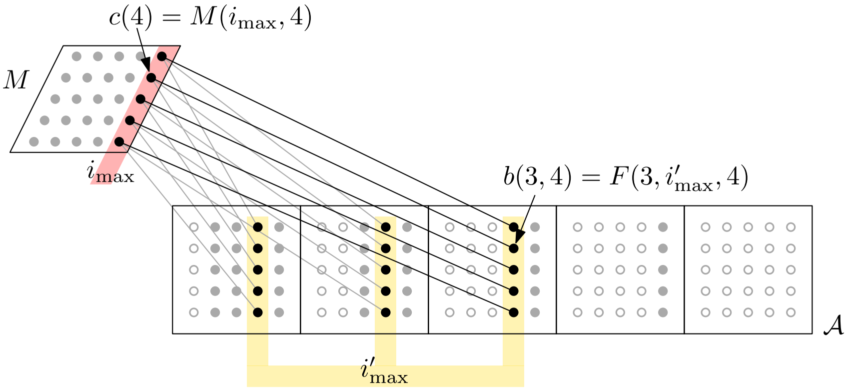

Figure 6.3: A depiction of the mono-prepattern we discover in *Case 3*. In this example, we have that $G \models E(b(i^*, j'), c(j)) \Leftrightarrow j = j'$.

(A.1) $\mathrm{atp}_{G^+}(b(i^*, j'), c(j))$ depends only on $\mathrm{otp}(j, j')$, and

(A.2) $\mathrm{atp}_{G^+}(b(i^*, j'), c(j))$ is not the same.

We claim that for all $i^* \in I$ and $j, j' \in J$ then also

(E.1) $G \models E(b(i^*, j'), c(j))$ depends only on $\mathrm{otp}(j, j')$, and

(E.2) $G \models E(b(i^*, j'), c(j))$ is not the same.

By definition of an atomic type, (A.1) implies (E.1). To show (E.2) we argue similarly to the proof of Lemma 6.12: We know that there are indices $i_1^*, i_2^* \in I$ and $j_1, j_1', j_2, j_2' \in J$ and atomic types τ_1 and τ_2 such that

$$\tau_1 = \mathrm{atp}_{G^+}(\underbrace{b(i_1^*, j_1'), c(j_1)}_{\bar{u}}) \neq \mathrm{atp}_{G^+}(\underbrace{b(i_2^*, j_2'), c(j_2)}_{\bar{v}}) = \tau_2.$$

By $(*)$, all $b(\cdot, \cdot)$ elements have the same atomic type in G^+ and all $c(\cdot)$ elements have the same atomic type in G^+. This means that the difference in the atomic types of $\bar{v}$ and $\bar{v}'$ must be caused by either a difference in their equality or adjacency type. If the difference is witnessed in the equality type, then by symmetry we can assume that

$$b(i_1^*, j_1') = c(j_1) \quad \text{and} \quad b(i_2^*, j_2') \neq c(j_2).$$

It follows that $b(i^*, j_1') = c(j_1)$ for all $i^* \in I$, since all of these pairs have the same order type $\mathrm{otp}(j_1', j_1)$. This is a contradiction to the columns of $\mathcal{A}$ being disjoint. Therefore, the difference in the types of $\bar{u}$ and $\bar{v}$ must be witnessed by their adjacency type and we have

$$G \models E(b(i_1^*, j_1'), c(j_1)) \quad \text{if and only if} \quad G \not\models E(b(i_2^*, j_2'), c(j_2)),$$

which proves (E.2).

Having proven (E.1) and (E.2) it is easily verified that the ranges of $b(\cdot, \cdot)$ and $c(\cdot)$ form a mono-prepattern on $\mathcal{A}$. Applying Lemma 6.20 to this mono-prepattern yields the desired transformer.

Case 4: Assume $\mathrm{atp}_{G^+}(F(i^*, i', j'), M(i, j))$ is not the same for all $i^*, i, i' \in I$ and $j, j' \in J$ with $i^* < i < i'$ (here $i < i'$ is the crucial assumption). We argue as in the previous case, exchanging the role of $i_{\max}$ and $i'_{\max}$. This means we define

$$c(j) := M(i'_{\max}, j) \quad \text{and} \quad b(i^*, j') := F(i^*, i_{\max}, j') \in A[i^*, *] \quad \text{for all } i^* \in I \text{ and } j \in J.$$

As in *Case 3*, the ranges of $b(\cdot, \cdot)$ and $c(\cdot)$ form a mono-prepattern on $\mathcal{A}$ and we conclude by Lemma 6.20.

Case 5: If none of the previous cases hold, then $\text{atp}_{G^+}(F(i^*, i', j'), M(i,j))$ and in particular $G^+ \models \alpha(F(i^*, i', j'), M(i,j))$ is the same for all $i^*, i, i' \in I$ and $j, j' \in J$ with $(i,j) \neq (i',j')$ and $i^* < i, i'$. By $(*)$, this also holds for the extremal elements that are in I' and J', but not in I and J. Let hence $\gamma \in \{true, false\}$ be such that

$$G^+ \models \alpha(F(i^*, i', j'), M(i,j)) \Leftrightarrow \gamma$$

for all $i^*, i, i' \in I'$ and $j, j' \in J'$ with $(i,j) \neq (i',j')$ and $i^* < i, i'$.

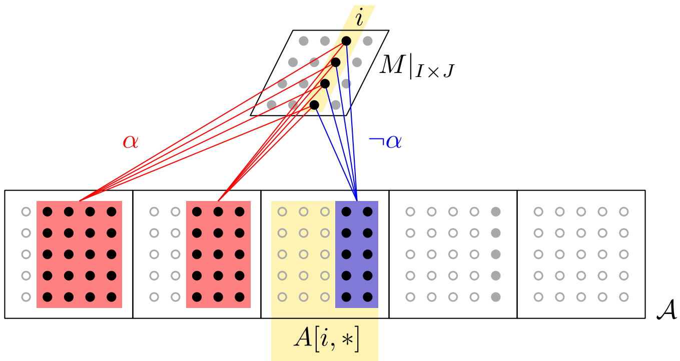

Figure 6.4: Illustration for *Case 5.1*. Depicted is the insulator $\mathcal{A}$, whose columns contain the $F(i^*, i, j)$ vertices. The $F(i^*, i, j)$ vertices are only defined for indices satisfying $i^* < i$, so going rightwards, the columns are filled up with placeholders (depicted as ○). The vertices in the ith column of the submesh $M|_{I \times J}$ are α-connected to all the $F(i^*, i, j)$ vertices in $A[<i, *]$ (marked in red), but $\neg\alpha$-connected to all the $F(i^*, i, j)$ vertices in $A[i, *]$ (marked in blue).

Case 5.1: Assume $\gamma = \mathit{true}$. At the beginning of the proof, we chose F such that $G^+ \models \alpha(F(i^*, i, j), M(i,j))$. Therefore,

$$G^+ \models \alpha(F(i^*, i', j'), M(i,j)) \quad \text{for all } i^*, i, i' \in I' \text{ with } i^* < i, i' \text{ and } j, j' \in J'. \tag{6.2}$$

Our goal is to reduce to Lemma 6.19 ($\clubsuit_1$). The situation is depicted in Figure 6.4. Let $P = \{F(i^*, i, j) : i^*, i \in I', j \in J' \text{ with } i^* < i\}$ and let G^{++} be the coloring of G^+ where the vertices of P are marked with an additional fresh color predicate. (By the definition of colored graphs given in the preliminaries, we are technically required to give every vertex *exactly* one color, and it would of course be trivial, though more cumbersome, to take this into account.) Define the quantifier-free formula $\beta(x, y) := x \in P \wedge \neg\alpha(x, y)$ in the signature of G^{++}. Let us now verify that

$$i = \min\{i' \in I : \exists v \in A[i', *] \ : \ G^{++} \models \beta(v, M(i,j))\} \tag{6.3}$$

for all $i \in I, j \in J$. As all elements in I are smaller than $i_{\max}$, for every $i \in I$, we have that the column $A[i, *]$ contains at least one element, say, $F(i, i_{\max}, j_{\max}) \in P$. By the assumption of the lemma, $M(i,j)$ is not α-connected to any element in $A[i, *]$ and therefore β-connected to $F(i, i_{\max}, j_{\max})$. By (6.2) and definition of P, $M(i,j)$ is α-connected (and therefore not β-connected) to all elements from P in columns to the left of I, as desired. Having verified (6.3), we conclude by Lemma 6.19.

Case 5.2: Assume $\gamma = \mathit{false}$. For all $i, i', i^* \in I'$ with $i^* < i, i'$ and $j, j' \in J'$, as $G^+ \models \alpha(F(i^*, i, j), M(i,j))$, it follows that

$$G^+ \models \alpha(F(i^*, i', j'), M(i,j)) \quad \text{if and only if} \quad (i,j) = (i',j').$$

Let $c(j) := M(i_{\max}, j)$ and $b(i^*, j') := F(i^*, i_{\max}, j')$ for all $i^* \in I$ and $j, j' \in J$. Then by $(*)$, and as $G^+ \models \alpha(b(i^*, j'), c(j))$ if and only if $j = j'$, we have that

- $\mathrm{atp}_{G^+}(b(i^*, j'), c(j))$ depends only on $\mathrm{otp}(j, j')$, and
- $\mathrm{atp}_{G^+}(b(i^*, j'), c(j))$ is not the same

for all $i^* \in I$ and $j, j' \in J$. As in *Case 3*, the ranges of $b(\cdot, \cdot)$ and $c(\cdot)$ form a mono-prepattern on $\mathcal{A}$ and we conclude by Lemma 6.20.

Having exhausted all cases, this proves the lemma. □

We can now prove Proposition 6.9, which we restate for convenience.

Proposition 6.9. *Let G be a graph containing a prepattern of order n on an insulator of height h and cost k. Then G contains a transformer of order $U_{h,k}(n)$ and length at most $4h-1$.*

Proof. If the prepattern is a mono-prepattern, we conclude by Lemma 6.20. Therefore, suppose that the prepattern is a bi-prepattern. Let the sequences I, J, vertices $(c_{i,j} : i \in I, j \in J)$, vertices s_1, s_2, quantifier-free formulas $\alpha_1(x; y, z_1), \alpha_2(x; y, z_1)$, and symbols $\sim_1, \sim_2 \in \{=, \neq\}$ be as in the definition of a bi-prepattern (see Definition 5.10).

Let G^+ be the expansion of G by four unary predicates, representing the sets $\{s_1\}$, $\{s_2\}$, $N_G(s_1)$, and $N_G(s_2)$. Then there are quantifier-free formulas $\beta_1(x, y)$ and $\beta_2(x, y)$ such that for all $u, v \in V(G)$ and $i \in \{1, 2\}$, we have $G^+ \models \alpha_i(u, v, s_i) \Leftrightarrow \beta_i(u, v)$.

Define the mesh $M\colon I \times J \to V(G)$ with $M(i, j) = c_{i,j}$ for $i \in I, j \in J$. Then M with the formula β_1 satisfies the assumptions of Lemma 6.19 if $\sim_1$ is $\neq$, and the assumptions of Lemma 6.21 if $\sim_1$ is $=$. We apply the appropriate lemma. In case of Lemma 6.21, this might yield a transformer and we are done. Otherwise, we obtain a conductor $M_1, \dots, M_t\colon I' \times J' \to V(G)$ with $|I'| = |J'| = U_{h,k}(n)$, $t \leqslant 2h$, where M_t is vertical, and $M_1 = M|_{I' \times J'}$.

Denote $M' := M_1^{\mathsf{T}}\colon J' \times I' \to V(G)$. Then M' with the formula β_2 satisfy the assumptions of Lemma 6.19 (with the roles of I' and J' swapped) if $\sim_2$ is $\neq$, and the assumptions of Lemma 6.21 if $\sim_2$ is $=$. We again apply the appropriate lemma. Either this yields a transformer and we are done, or we obtain a conductor $M'_1, \dots, M'_u\colon J'' \times I'' \to V(G)$ with $|J''| = |I''| = U_{h,k}(|I'|)$, $u \leqslant 2h$, where M'_u is vertical and $M'_1 = M'|_{J'' \times I''}$. Notice that

$$M'_u|_{J'' \times I''}, \dots, M'_1|_{J'' \times I''} = M_1^{\mathsf{T}}|_{J'' \times I''}, \dots, M_t^{\mathsf{T}}|_{J'' \times I''}$$

is a conductor of order $U_{h,k}(n)$ and length $u + t - 1 \leqslant 4h - 1$. Observe that M'_u and M_t are both vertical, and that M_t^{T} is horizontal. Thus, the sequence forms a transformer. □

6.1.3 Regular and Minimal Transformers

In this subsection, we normalize the transformers derived in Section 6.1.2.

Definition 6.22. A transformer $T = (M_1, \dots, M_h)$ in a graph G is *regular* if for all $s, t \in [h]$, the pair (M_s, M_t) is a regular pair of meshes (in particular also for $s = t$). We say that T is *minimal* if it is regular and for all $s, t \in [h]$ the following conditions hold:

- M_s is vertical if and only if $s = 1$,
- M_s is horizontal if and only if $s = h$,
- $M_s \neq M_t$ if $s \neq t$ (that is, no two meshes are identical),
- the pair (M_s, M_t) is conducting if $|s - t| = 1$,
- the pair (M_s, M_t) is homogeneous if $|s - t| > 1$.

In particular, in a minimal transformer $M_1, \dots, M_h$, we have that either $h = 1$ and $M_1 = M_h$ is both horizontal and vertical, or $h > 1$ and M_1 is vertical and not horizontal, M_h is horizontal and not vertical, and $M_2, \dots, M_{h-1}$ are neither horizontal nor vertical.

Lemma 6.23. *If G contains a transformer of length h and order n then G contains a regular transformer of length h and order $U_h(n)$.*

Proof. Let $M_1, \ldots, M_h$ be a transformer of length h in G. We can assume $I, J = [n]$. For all $i, j \in [n]$, let $\pi_{i,j} \in V(M)^h$ be the h-tuple

$$\pi_{i,j} := (M_1(i,j), \ldots, M_h(i,j)).$$

By Bipartite Ramsey (Lemma 4.15), there are sets $I', J' \subseteq [n]$ with $|I'| = |J'| \geqslant U_h(n)$ such that $\mathrm{atp}(\pi_{i,j}, \pi_{i',j'})$ depends only on $\mathrm{otp}(i,i')$ and $\mathrm{otp}(j,j')$, for all $i, i' \in I'$ and $j, j' \in J'$. It follows that for all $s, t \in [h]$, the meshes $M_s|_{I' \times J'}$ and $M_t|_{I' \times J'}$ form a regular mesh pair. Thus, the sequence $M_1|_{I' \times J'}, \ldots, M_h|_{I' \times J'}$ is a regular transformer. □

Lemma 6.24. *If G contains a transformer of length h and order n then G contains a minimal transformer of length at most h and order $U_h(n)$.*

Proof. By Lemma 6.23, there is a regular transformer $M_1, \ldots, M_h$ of length h and order $U_h(n)$ in G. Consider the graph $\mathcal{G}$ whose vertices are the meshes $M_1, \ldots, M_h$, and edges are pairs M_iM_j such that the pair (M_i, M_j) is conducting. Clearly, $\mathcal{G}$ contains a path of length h that starts in a vertical mesh and ends in a horizontal mesh. Let $\pi := M_{i_1}, \ldots, M_{i_p}$ be a shortest path in $\mathcal{G}$ that starts in a vertical mesh and ends in a horizontal mesh. Then the path π is an induced path of length at most h in $\mathcal{G}$, which means that a pair M_{i_s}, M_{i_t}, for distinct $s, t \in [p]$ is conducting if and only if $|s - t| = 1$. As every pair M_{i_s}, M_{i_t} is regular, it follows that $M_{i_1}, \ldots, M_{i_p}$ is a minimal transformer of length $p \leqslant h$ and order $U_h(n)$. □

6.2 Converters and Crossings

Our next goal is to analyze the structure of minimal transformers in graphs. We will arrive at a notion of a converter, which is similar to a crossing. Finally, from converters, we will obtain crossings.

6.2.1 Regular Pairs of Meshes

We study the structure of regular pairs of meshes in graphs. We introduce the following notions.

Definition 6.25. Let $M, M' \colon I \times J \to V(G)$ be two meshes in a graph G. We say that the pair (M, M') is

- *disjoint* if $V(M) \cap V(M') = \emptyset$,
- *matched* if for all $i, i' \in I$ and $j, j' \in J$,

$$G \models E(M(i,j), M'(i',j')) \quad \text{if and only if} \quad (i,j) = (i',j').$$

- *co-matched* if the pair (M, M') is matched in the complement graph $\overline{G}$,
- *non-adjacent* if $V(M)$ and $V(M')$ are non-adjacent in G,
- *fully adjacent* if $uv \in E(G)$ for all $u \in V(M)$ and $v \in V(M')$.

We prove some preliminary observations regarding regular pairs of meshes.

Lemma 6.26. *Let $M, M' \colon I \times J \to V(G)$ be a regular pair of meshes of order $n > 2$ in a graph G. Then M and M' are either identical, or disjoint.*

Proof. We show that if $M(i,j) = M'(i',j')$ for some $i,i' \in I$ and $j,j' \in J$, then $(i,j) = (i',j')$. By regularity of the pair (M, M'), this implies that $M(i,j) = M'(i,j)$ for all $i,j \in [n]$, so M and M' are identical. So suppose that $M(i,j) = M'(i',j')$ for some $i,i' \in I$ and $j,j' \in J$ with $(i,j) \neq (i',j')$. Up to reversing the order of I and up to exchanging the role of I and J, we can assume that $i < i'$. Pick $i_1 < i_2 < i_3 \in I$. Then by regularity we have that $M(i_1,j) = M'(i_2,j')$ and $M(i_1,j) = M'(i_3,j')$. Thus, $M'(i_2,j') = M'(i_3,j')$, contradicting injectivity of M'. □

Lemma 6.27. *Let $M, M' \colon I \times J \to V(G)$ be a regular pair of meshes of order n in a graph G. Suppose that $G \models E(M(i,j), M'(i',j'))$ is not the same for all $i,i' \in I$, $j,j' \in J$ with $j < j'$, or is not the same for all $i,i' \in I$, $j,j' \in J$ with $j > j'$. Then there exist subsequences $I' \subseteq I$ and $J' \subseteq J$ of order $U(n)$ such that the submeshes $M|_{I' \times J'}$ and $M'|_{I' \times J'}$ are both vertical.*

Proof. We define $j_{\min} := \min(J)$, $j_{\max} := \max(J)$, $J' := J - \{j_{\min}, j_{\max}\}$, $I' := I - \{\min(I), \max(I)\}$. We show the argument for M, while the case for M' follows by symmetry. Suppose the first case holds, that is, $G \models E(M(i,j), M'(i',j'))$ is not the same for all $i,i' \in I$, $j,j' \in J$ with $j < j'$. The other case proceeds by the same argument, exchanging the roles of $j_{\max}$ and $j_{\min}$.

Let $a(i') = M'(i', j_{\max})$ for $i' \in I'$. Then by regularity, $\mathrm{atp}(M(i,j), a(i'))$ depends only on $\mathrm{otp}(i,i')$, for $i,i' \in I'$ and $j \in J'$. Furthermore, $G \models E(M(i,j), a(i'))$ is not the same for all $i,i' \in I'$ and $j \in J'$, by the assumption and by regularity. Hence, $M|_{I' \times J'}$ is vertical. □

Lemma 6.28. *Let $M, M' \colon I \times J \to V(G)$ be a regular pair of meshes of order n in a graph G. Suppose that $G \models E(M(i,j), M'(i',j'))$ is not the same for all $i,i' \in I, j,j' \in J$ with $(i,j) \neq (i',j')$. Then there exist subsequences $I' \subseteq I$ and $J' \subseteq J$ of order $U(n)$ such that the submeshes $M|_{I' \times J'}$ and $M'|_{I' \times J'}$ are either both vertical, or both horizontal.*

Proof. It follows from the assumption that one of the following cases holds:

- $G \models E(M(i,j), M'(i',j'))$ is not the same for all $i,i' \in I$, $j,j' \in J$ with $j < j'$,
- $G \models E(M(i,j), M'(i',j'))$ is not the same for all $i,i' \in I$, $j,j' \in J$ with $j > j'$,
- $G \models E(M(i,j), M'(i',j'))$ is not the same for all $i,i' \in I$, $j,j' \in J$ with $i < i'$,
- $G \models E(M(i,j), M'(i',j'))$ is not the same for all $i,i' \in I$, $j,j' \in J$ with $i > i'$.

In the first two cases we conclude by Lemma 6.27. In the last two cases we conclude by applying the same lemma to M^{T} and M'^{T}. □

Lemma 6.29. *Let $M, M' \colon I \times J \to V(G)$ be a conducting pair of disjoint meshes of order n in G. Then there exist subsequences $I' \subseteq I$ and $J' \subseteq J$ of order $U(n)$ such that $M|_{I' \times J'}$ and $M'|_{I' \times J'}$ are either 1. both vertical, 2. both horizontal, 3. matched, or 4. co-matched.*

Proof. Since M and M' are conducting, the pair is regular but not homogeneous. For all $i,i' \in I$ and $j,j' \in J$ we have that

(1) $G \models E(M(i,j), M'(i',j'))$ depends only on $\mathrm{otp}(i,i')$ and $\mathrm{otp}(j,j')$, (by regularity)
(2) $\mathrm{atp}_G(M(i,j), M'(i',j'))$ is not always the same, (by non-homogeneity)
(3) the equality type of $M(i,j)$ and $M'(i',j')$ is always $(\neq)$, and (by disjointness)
(4) the adjacency between $M(i,j)$ and $M'(i',j')$ is not always the same. (by (2) and (3))

Assume $G \models E(M(i,j), M'(i',j'))$ is the same for all $i,i' \in I$, $j,j' \in J$ with $(i,j) \neq (i',j')$. Up to replacing G with $\overline{G}$, by (4) we have that for all $i,i' \in I$, $j,j' \in J$,

$$G \models E(M(i,j), M'(i',j')) \quad \text{if and only if } (i,j) = (i',j').$$

Thus, M and M' are matched or co-matched.

Otherwise, $G \models E(M(i,j), M'(i',j'))$ is not the same for all $i, i' \in I$, $j, j' \in J$ with $(i,j) \neq (i',j')$, and we conclude that M and M' are both vertical or both horizontal by Lemma 6.28. $\square$

Lemma 6.30. *Let G contain a minimal transformer of order n and length h. Then G contains a minimal transformer $M_1, \ldots, M_{h'}$ of order $U_h(n)$ and length $h' \leqslant h$ such that the following conditions are satisfied for $s, t \in [h']$:*

1. *If $s \neq t$, then M_s and M_t are disjoint.*
2. *If $|s-t| > 1$ then M_s and M_t are either non-adjacent or fully adjacent.*
3. *If $|s-t| = 1$ then M_s and M_t are either matched or co-matched.*

Proof. If $n \leqslant 3$, as the minimal transformer we take any transformer of order 1 and length 1, so assume $n > 3$. The first two properties hold in any minimal transformer T of order $n > 3$. Indeed, as all meshes of T are pairwise distinct and regular, by Lemma 6.26 they are pairwise disjoint. Also, any pair of non-consecutive meshes is regular and not conducting, hence (as $n > 3$) homogeneous.

We argue that we can find a minimal transformer satisfying additionally the last property. Let $T = (M_1, \ldots, M_h)$ be a minimal transformer of order n.

Applying Lemma 6.29 to every pair of consecutive meshes in T (and each time reducing the order of the transformer to $U_h(n)$) we may assume that for every pair (M_s, M_{s+1}) of consecutive meshes in T, the pair is either matched, or co-matched, or both meshes are vertical, or both are horizontal. Let $M_s, \ldots, M_t$ be a subsequence of $M_1, \ldots, M_h$ of shortest length such that M_s is vertical and M_t is horizontal. It follows that every two consecutive meshes in T' are matched or co-matched, and that $T' = (M_s, \ldots, M_t)$ is a minimal transformer. Thus, the last property in the statement is satisfied. $\square$

6.2.2 Regular Meshes

We now analyze the structure of single meshes, depending on whether they are horizontal and/or vertical. We introduce some notation.

Definition 6.31. A single mesh $M\colon I \times J \to V(G)$ is *regular* in a graph G, if the pair (M, M) is. That is, $\mathrm{atp}_G(M(i,j), M'(i',j'))$ depends only on $\mathrm{otp}(i,i')$ and $\mathrm{otp}(j,j')$ for all $i, i' \in I$ and $j, j' \in J$.

Note that every mesh in a regular/minimal transformer is regular.

Definition 6.32. A *mesh pattern* is a subset P of the four lines in the following diagram:

All 2^4 possible mesh patterns are depicted below (including the empty pattern $\cdot$).

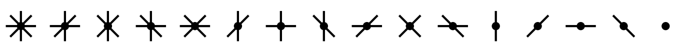

Let P be a mesh pattern. A regular mesh $M\colon I \times J \to V(G)$ with $|I|, |J| \geqslant 2$ is a P-*mesh* in a graph G if for all $(i,j), (i',j') \in I \times J$ with $i \leqslant i'$, the vertices $M(i,j)$ and $M(i',j')$ are adjacent in G if and only if one of the following conditions holds:

- $i = i'$ and $j \neq j'$ and $\vert \in P$,
- $i < i'$ and $j < j'$ and $\diagup \in P$,
- $i < i'$ and $j = j'$ and $- \in P$,
- $i < i'$ and $j > j'$ and $\diagdown \in P$.

For example, M is a $\cdot$-mesh if and only if $V(M)$ induces an independent set, and if M is a $+$-mesh, then $V(M)$ induces a rook graph in G, and if M is a $\ast$-mesh, then $V(M)$ induces a comparability grid in G. (Recall that the *comparability grid* of order n consists of vertices $\{a_{i,j} : i, j \in [n]\}$ and edges between vertices $a_{i,j}$ and $a_{i',j'}$ if and only if either $i = i'$, or $j = j'$, or $i < i' \Leftrightarrow j < j'$.)

Definition 6.33. A *generalized grid* in a graph G is a regular mesh $M \colon I \times J \to V(G)$ satisfying the following conditions:

- $G \models E(M(i,j), M(i',j'))$ does not depend only on $\mathrm{otp}(i,i')$, for $i, i' \in I$ and $j, j' \in J$ with $(i,j) \neq (i',j')$, and
- $G \models E(M(i,j), M(i',j'))$ does not depend only on $\mathrm{otp}(j,j')$, for $i, i' \in I$ and $j, j' \in J$ with $(i,j) \neq (i',j')$.

Observe that a P-mesh M is a generalized grid if and only if P is *not* among $\{\cdot,$ ⁞, ⋯$\}$ or their complements $\{\ast,$ ⋇, ⁎$\}$.

Lemma 6.34. *Let $M \colon I \times J \to V(G)$ be a regular mesh in a graph G. Then M is vertical, or is horizontal, or is a $\cdot$-mesh in G or in $\overline{G}$.*

Proof. Suppose that $G \models E(M(i,j), M(i',j'))$ is the same for all $i, i' \in I$, $j, j' \in J$ with $(i,j) \neq (i',j')$. Then $V(M)$ forms an independent set or a clique in G. Otherwise, the statement follows by Lemma 6.28, applied to $M = M'$. □

Lemma 6.35. *Let $M \colon I \times J \to V(G)$ be a regular mesh in a graph G, and assume $G \models E(M(i,j), M(i',j'))$ is the same for all $i, i' \in I$ and $j, j' \in J$ with $i < i'$. Then either in G or in $\overline{G}$, M is a* ⁞*-mesh or a $\cdot$-mesh.*

Proof. Replacing G with $\overline{G}$ if needed, we may assume that $G \models \neg E(M(i,j), M(i',j'))$ for all $i, i' \in I$ and $j, j' \in J$ with $i < i'$. By symmetry of the edge relation, we have that

$$G \models \neg E(M(i,j), M(i',j'))$$

holds for all $i, i' \in I$ and $j, j' \in J$ with $i \neq i'$. If $G \models E(M(i,j), M(i,j'))$ for some $i \in I$ and distinct $j, j' \in J$, then M is a ⁞-mesh. Otherwise, M is a $\cdot$-mesh. □

Lemma 6.36. *Let M be a regular mesh in a graph G which is not horizontal. Then either in G or in $\overline{G}$, M is a* ⁞*-mesh or a $\cdot$-mesh.*

Proof. As M^{T} is not vertical, we can apply the contrapositive of Lemma 6.27 to M^{T} and M^{T}. We conclude that $G \models E(M(i,j), M(i',j'))$ is the same for all $i, i' \in I$ and $j, j' \in J$ with $i < i'$. The conclusion follows from Lemma 6.35. □

Definition 6.37. A mesh $M \colon I \times J \to V(G)$ is *capped* if there is a function $a \colon I' \to V(G)$, where $I' = I - \{\min(I), \max(I)\}$, such that one of the following conditions holds in G or in $\overline{G}$:

(=) for all $i, i' \in I'$ and $j \in J$, $M(i,j)$ is adjacent to $a(i')$ if and only if $i = i'$, or

(⩽) for all $i, i' \in I'$ and $j \in J$, $M(i,j)$ is adjacent to $a(i')$ if and only if $i \leqslant i'$.

More precisely, a capped mesh M is α-capped, for $\alpha \in \{=, \leqslant\}$, if the above condition α holds.

Lemma 6.38. *Let $M \colon I \times J \to V(G)$ be a vertical mesh in a graph G with $|J| > 1$. Then M is capped.*

Proof. Since M is vertical, there is an $a \colon I \to V(G)$ such that, for $J' = J - \{\min(J), \max(J)\}$,

- atp$(M(i,j), a(i'))$, depends only on otp(i,i'), for all $i, i' \in I$ and $j \in J'$, and
- atp$(M(i,j), a(i'))$ is not the same for all $i, i' \in I$ and $j \in J'$.

First observe that the ranges of the functions a and M are disjoint. Assume otherwise, that is, that $a(i') = M(i,j)$ for some $i, i' \in I$ and $j \in J$. Pick $j' \in J$ distinct from j, which exists since we assume that $|J| > 1$. Then we have that $a(i') = M(i, j')$, by the first defining condition of the function a. This contradicts the fact that M is an injective function.

Let $p_<, p_=, p_> \in \{0,1\}$ be such that for each $R \in \{<, =, >\}$,

$$G \models E(M(i,j), a(i')) \quad \Leftrightarrow \quad p_R = 1 \quad \textit{for all } j \in J', i, i' \in I \textit{ with } i \, R \, i'.$$

By the assumption on a, the values $p_<, p_=, p_>$ are not all equal. Replacing G with $\overline{G}$ if needed, we can assume that $p_> = 0$. Thus, one of three cases occurs:

1. $(p_<, p_=, p_>) = (0,1,0)$,
2. $(p_<, p_=, p_>) = (1,1,0)$,
3. $(p_<, p_=, p_>) = (1,0,0)$.

Let $I' = I - \{\min(I), \max(I)\}$ and let $a|_{I'}$ be the restriction of a to the domain I'. In the first case, $(M, a|_{I'})$ is a $=$-capped mesh. In the second case, $(M, a|_{I'})$ is a $\leqslant$-capped mesh. Suppose the third case occurs, and let $b\colon I' \to V(G)$ be such that $b(i) = a(i_+)$ where i_+ is the successor of i in I. Then (M, b) is a $\leqslant$-capped mesh. □

Lemma 6.39. *Let $M\colon I \times J \to V(G)$ be a regular mesh of order $n > 1$ in a graph G. Then the following hold.*

1. *If M is not horizontal and not vertical, then M forms a ·-mesh in G or in $\overline{G}$.*
2. *If M is vertical and not horizontal, then one of two cases occurs in G or in $\overline{G}$:*
   (a) *M is a ǀ-mesh, or*
   (b) *M is a ·-mesh and capped.*
3. *If M is both horizontal and vertical, then one of four cases occurs in G or in $\overline{G}$:*
   (a) *M is a generalized grid,*
   (b) *M is a ǀ-mesh and $M^\top$ is capped,*
   (c) *M is a –-mesh and M is capped, or*
   (d) *M is a ·-mesh and both M and $M^\top$ are capped.*

Proof. The first item is by Lemma 6.34.

We prove the second item. By Lemma 6.36, either in G or in $\overline{G}$, M is a ǀ-mesh or a ·-mesh. In the first case we are done. In the second case, Lemma 6.38 yields the conclusion.

Finally, we prove the third item. Assume that M is both vertical and horizontal. Then, by Lemma 6.38, both M and $M^\top$ are capped. Suppose M is not a generalized grid, as otherwise condition (a) holds, and we are done. Then either $G \models E(M(i,j), M(i',j'))$ depends only on otp(i,i'), or it depends only on otp(j,j'), for all distinct $(i,j), (i',j') \in I \times J$. Suppose it depends only on otp(i,i'), while the other case follows by replacing M with $M^\top$. By Lemma 6.35 we have that either in G or in $\overline{G}$, M is a ǀ-mesh or a ·-mesh. If M is a ǀ-mesh, condition (b) holds. If M is a ·-mesh, condition (d) holds.

This concludes the lemma. □

Lemma 6.40. *Let G be a graph and M be a regular mesh of order n in G which is a generalized grid. Then M is a +-mesh or ×-mesh, or G contains a comparability grid of order $\lfloor \sqrt{n} \rfloor$ as induced subgraph.*

Proof. Let M be a regular mesh of order m in G, and let P be the pattern of M. Note that if P is among {·, ǀ, –} or their complements {∗, ⁎, ✱}, then M is not a generalized grid. So assume that P is not among those patterns. We can exclude $P \in \{+, \times\}$, so assume that P is among the remaining patterns, that is, $P \in$ {⟋, ⟍, ⫽, ⌿, ⫻, ⍀, ⋇, ✳}.

Up to symmetries that swap the two coordinates and invert their orders, it is enough to consider the cases $P \in$ {⟋, ⫽, ⋇}. If $P =$ ⋇, then $V(M)$ induces a comparability grid of order m in G. It remains to show that if G has a P-mesh of order $m = n^2$, for some $P \in$ {⟋, ⫽}, then G contains a comparability grid of order n.

Let $I = [n] \times [n]$, and let $\leqslant_{\text{lex}}$ denote the lexicographic order on $[n] \times [n]$. To declutter notation, below we write ij for a pair $(i, j) \in I$.

Assume G has a P-mesh of order n^2. By reindexing $([n^2], \leqslant)$ as $(I, \leqslant_{\text{lex}})$, we can view it as a P-mesh $M \colon I \times I \to V(G)$ in G. Then for all $i_1, i_2, i'_1, i'_2 \in I$ and $j_1, j_2, j'_1, j'_2 \in I$ with $(i_1 i_2, j_1 j_2) \neq (i'_1 i'_2, j'_1 j'_2)$ we have that

- if $P =$ ⟋, then $M(i_1 i_2, j_1 j_2)$ and $M(i'_1 i'_2, j'_1 j'_2)$ are adjacent if and only if

$$\underbrace{(i_1 i_2 <_{\text{lex}} i'_1 i'_2 \text{ and } j_1 j_2 <_{\text{lex}} j'_1 j'_2)}_{\text{edges from } M(i_1 i_2, j_1 j_2) \text{ to the top right}} \text{ or } \underbrace{(i'_1 i'_2 <_{\text{lex}} i_1 i_2 \text{ and } j'_1 j'_2 <_{\text{lex}} j_1 j_2)}_{\text{edges from } M(i_1 i_2, j_1 j_2) \text{ to the bottom left}}, \tag{6.4}$$

- if $P =$ ⫽, then $M(i_1 i_2, j_1 j_2)$ and $M(i'_1 i'_2, j'_1 j'_2)$ are adjacent if and only if

$$\underbrace{(i_1 i_2 = i'_1 i'_2)}_{\text{edges in the same column}} \text{ or } \underbrace{(i_1 i_2 <_{\text{lex}} i'_1 i'_2 \text{ and } j_1 j_2 <_{\text{lex}} j'_1 j'_2)}_{\text{edges from } M(i_1 i_2, j_1 j_2) \text{ to the top right}} \text{ or } \underbrace{(i'_1 i'_2 <_{\text{lex}} i_1 i_2 \text{ and } j'_1 j'_2 <_{\text{lex}} j_1 j_2)}_{\text{edges from } M(i_1 i_2, j_1 j_2) \text{ to the bottom left}}.$$

The adjacencies are depicted in Figure 6.5. Note that the cases $P =$ ⟋ and $P =$ ⫽ differ only if $(i_1 i_2 = i'_1 i'_2)$.

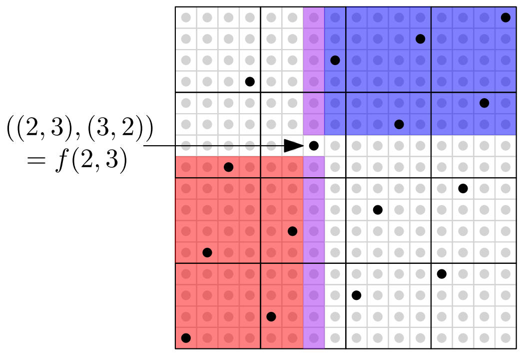

Figure 6.5: A P-mesh indexed by $([4]^2, <_{\text{lex}}) \times ([4]^2, <_{\text{lex}})$. The neighbors of $((2, 3), (3, 2))$ are (red ∪ blue) if $P =$ ⟋ and (red ∪ blue ∪ purple) if $P =$ ⫽. In solid black: the vertices in the range of f which form the comparability grid of order 4.

Consider the function $f \colon [n] \times [n] \to I \times I$, such that

$$f(i, j) = (ij, ji) \qquad \text{for } i, j \in [n].$$

The range of this function is depicted in Figure 6.5. We verify that $M' := M \circ f \colon [n] \times [n] \to V(G)$ is a ⋇-mesh. Let $(i, j), (i', j') \in [n] \times [n]$ be distinct, with $i \leqslant i'$. We need to show that $M'(i, j) = M(ij, ji)$ and $M'(i', j') = M(i'j', j'i')$ are adjacent in G if and only if $i = i'$ or $j \leqslant j'$. Note that $(ij) \neq (i'j')$, as $(i, j) \neq (i', j')$ are distinct, so the distinction between $P =$ ⟋ and $P =$ ⫽ is irrelevant. We can therefore assume $P =$ ⟋, and we will argue using (6.4) for

$$i_1 i_2 := ij,\ i'_1 i'_2 := i'j',\ j_1 j_2 := ji, \text{ and } j'_1 j'_2 := j'i'.$$

Assume first that $i = i'$. In this case we want to show that $M'(i,j)$ and $M'(i',j')$ are adjacent. As ij and $i'j'$ are distinct, we either have $j < j'$ or $j' < j$. If $j < j'$ then $ij <_{\text{lex}} i'j'$ and $ji <_{\text{lex}} j'i'$, and we conclude using the first disjunct of (6.4). If $j' < j$ then $i'j' <_{\text{lex}} ij$ and $j'i' <_{\text{lex}} ji$, and we conclude using the second disjunct of (6.4).

Assume now that $i \neq i'$, so $i < i'$. In this case we want to show that $M'(i,j)$ and $M'(i',j')$ are adjacent if and only if $j \leqslant j'$. As $ij <_{\text{lex}} i'j'$, by (6.4) we have that $M'(i,j)$ and $M'(i',j')$ are adjacent if and only if $ji <_{\text{lex}} j'i'$, which is the case if and only if $j \leqslant j'$, as desired.

Having exhausted all cases, we conclude that M' is a ∗-mesh of order n, and therefore, G contains an induced comparability grid of order n. □

6.2.3 Converters

The following notion is aimed at providing a low-level description of minimal transformers.

Definition 6.41. Fix $h, n \geqslant 1$ and a graph G. A *converter* of length h and order n in G consists of

- meshes $M_1, \ldots, M_h \colon [n] \times [n] \to V(G)$ in G,
- two functions $a, b \colon [n] \to V(G)$; we denote the range of these functions by $V(a), V(b)$,

such that the following conditions hold:

(G.1) the sets $V(M_1), \ldots, V(M_h)$ are pairwise disjoint,

(G.2) for $s, t \in \{1, \ldots, h\}$, if $|s - t| = 1$, then M_s is matched or co-matched with M_t,

(G.3) for $s, t \in \{1, \ldots, h\}$, if $|s - t| > 1$, then $V(M_s)$ is fully adjacent or non-adjacent to $V(M_t)$,

(G.4) for $s \in \{2, \ldots, h-1\}$, $V(M_s)$ is an independent set or a clique,

(G.5) Let $(M, f) \in \{(M_1, a), (M_h^{\mathsf{T}}, b)\}$. Then one of the following three conditions holds in G or in $\overline{G}$:

(C) M is a ¦-mesh, or $h = 1$ and M is a +-mesh. Moreover, $M(i,j)$ is adjacent to $f(i')$ if and only if $i = i'$, for all $i, i', j \in [n]$,

(S) M is a ··-mesh, or $h = 1$ and M is a ⁃-mesh. Moreover, $M(i,j)$ is adjacent to $f(i')$ if and only if $i = i'$, for all $i, i', j \in [n]$,

(H) M is a ··-mesh, or $h = 1$ and M is a ⁃-mesh. Moreover, $M(i,j)$ is adjacent to $f(i')$ if and only if $i' \leqslant i$, for all $i, i', j \in [n]$.

Say that (M, f) as above has *kind* $\mathbf{C}, \mathbf{S}$, or $\mathbf{H}$ respectively, if it satisfies the corresponding condition above (which stand for *clique*, *star*, and *half-graph*, respectively). A converter has *kind* (α, β), where $\alpha, \beta \in \{\mathbf{C}, \mathbf{S}, \mathbf{H}\}$, if (M_1, a) has kind α and (M_h^{T}, b) has kind β.

A converter is *proper* if the following hold:

- the sets $V(a), V(M_1), \ldots, V(M_h), V(b)$ are pairwise disjoint,
- the sets $V(a)$ and $V(M_t)$ are homogeneously connected for all $1 < t \leqslant h$,
- the sets $V(b)$ and $V(M_t)$ are homogeneously connected for all $1 \leqslant t < h$,
- the sets $V(a)$ and $V(b)$ are homogeneously connected,
- each of the sets $V(a)$ and $V(b)$ induces an independent set or a clique.

Lemma 6.42. *Let G contain a converter of length h and order n. Then, for some number $m \geqslant U_h(n)$, either G contains a proper converter of length at most h and order m, or $m = 1$.*

Proof. Let $M_1, \ldots, M_h \colon [n] \times [n] \to V(G)$ and $a, b \colon [n] \to V(G)$ form a converter. For every pair $i, j \in [n]$ let $\pi_{i,j}$ denote the sequence $(a(i), M_1(i,j), \ldots, M_h(i,j), b(j))$. Apply Bipartite Ramsey (Lemma 4.15) to get sets $I \subseteq [n]$ and $J \subseteq [n]$ of size $m = U_h(n)$, so that $\text{atp}(\pi_{i,j}\pi_{i',j'})$

only depends on $\mathrm{otp}(i, i')$ and $\mathrm{otp}(j, j')$ for $i \in I$ and $j \in J$. Observe that if $a(i) = M_s(i', j)$ for some $i, i' \in I$, $j \in J$, and $s \in [h]$, then also $a(i) = M_s(i', j')$ for some other j', which implies $M_s(i', j) = M_s(i', j')$ and contradicts injectivity of M_s, unless $|J| = m = 1$. Similarly, if $a(i) = b(j)$ for some $i \in I$ and $j \in J$, then we conclude that $a(i) = a(i')$ for all $i, i' \in I$, which implies that $|I| = m = 1$, since the conditions $\mathbf{C}, \mathbf{S}, \mathbf{H}$ imply that a is injective. A similar argument can be made for b.

The tuple $M_1' := M_1|_{I \times J}, \ldots, M_h' := M_h|_{I \times J}, a' := a|_I, b' := b|_J$, with I and J both reindexed to $[m]$, forms a converter of order m. As discussed above and by (G.1), we observe the sets $V(a'), V(b'), V(M_1'), \ldots, V(M_h')$ to be pairwise disjoint. Further note that $V(a')$ and $V(b')$ each induce an independent set or a clique by construction. Lastly, to have a *proper* converter, we have to ensure that $V(M_t')$ and $V(b')$ are homogeneous for all $1 \leqslant t < h$, and that $V(M_t')$ and $V(a')$ are homogeneous for all $1 < t \leqslant h$.

We only sketch this last argument. Suppose for example that $V(b')$ and $V(M_t')$ are not homogeneous, for some $1 \leqslant t < h$. Then M_t' is either horizontal, or is vertical, as witnessed by b'. We can therefore obtain a converter of the same order and smaller length. Thus, by replacing the converter $a', M_1', \ldots, M_h', b'$ by a shorter one if needed, we arrive at a proper converter of order m. □

Lemma 6.43. *Let G contain a transformer of length h and order n. Then there is a number $m \geqslant U_h(n)$ such that, G contains a converter of length at most h and order m, or contains a comparability grid of order m as an induced subgraph.*

Proof. By Lemma 6.24 and Lemma 6.30, G contains a minimal transformer $M_1, \ldots, M_{h'} \colon I \times J \to V(G)$ of order $U_h(n)$ and length $h' \leqslant h$, such that:

- the sets $V(M_1), \ldots, V(M_h')$ are pairwise disjoint,
- for $s, t \in [h']$ with $|s - t| = 1$, M_s and M_t are matched or co-matched,
- for $s, t \in [h']$ with $|s - t| > 1$, M_s and M_t are fully adjacent or fully non-adjacent.

Let $I' = I - \{\min(I), \max(I)\}$ and $J' = J - \{\min(J), \max(J)\}$.

Suppose first that $h' = 1$. Then M_1 is both vertical and horizontal. By Lemma 6.39, one of the following holds for $M := M_1$ in either G or $\overline{G}$:

(a) M is a generalized grid,
(b) M is a ǀ-mesh and $M^\top$ is α-capped for some $\alpha \in \{=, \leqslant\}$,
(c) M is a –-mesh and M is α-capped for some $\alpha \in \{=, \leqslant\}$, or
(d) M is a ·-mesh and there are some $\alpha_1, \alpha_2 \in \{=, \leqslant\}$ such that M is α_1-capped and $M^\top$ is α_2-capped.

In case (a), M is by definition a generalized grid in both $\overline{G}$ and G. By Lemma 6.40, M is a P-mesh with $P \in \{+, \times\}$ in G, or G contains or a comparability grid of order $U_h(n)$ as an induced graph. If M is a P-mesh, then $M|_{I' \times J'}$ together with a, b has kind $(\mathbf{C}, \mathbf{C})$, where $a \colon I' \to V(G)$ is defined by $a(i) = M(i, \min(J))$ for $i \in I'$, and where $b \colon J' \to V(G)$ is defined by $b(j) = M(\min(I), j)$ for $j \in J'$. In either case, the statement holds.

Suppose we are in case (b). Let $b \colon J' \to V(G)$ witness that $M^\top$ is α-capped. Let $a \colon I' \to V(G)$ be defined by $a(i) = M(i, \min(J))$. Then $M|_{I' \times J'}$, a and b form a converter of kind $(\mathbf{C}, \mathbf{S})$ if α is $=$, and a converter of kind $(\mathbf{C}, \mathbf{H})$ if α is $\leqslant$. The case (c) is symmetric.

Finally, suppose we are in case (d). Let $a \colon I' \to V(G)$ witness that M is α_1-capped, and $b \colon J' \to V(G)$ witness that $M^\top$ is α_2-capped. Then $M|_{I' \times J'}$, a, and b form a converter of kind (τ_1, τ_2), where $\tau_i = \mathbf{S}$ if α_i is $=$, and $\tau_i = \mathbf{H}$ if α_i is $\leqslant$, for $i = 1, 2$.

This concludes the case of length $h' = 1$. Suppose now that $h' > 1$. Then, by the minimality assumption (Definition 6.22),

- M_1 is vertical and not horizontal,
- $M_{h'}$ is horizontal and not vertical, and
- $M_2, \dots, M_{h'-1}$ are neither vertical nor horizontal.

By Lemma 6.39, we conclude that each of $M_2, \dots, M_{h'-1}$ is a $\cdot$-mesh in either G or $\overline{G}$. By Lemma 6.39 applied to M_1, one of two cases holds in either G or $\overline{G}$:

(a) M_1 is a $\wr$-mesh, or

(b) M_1 is a $\cdot$-mesh and is α-capped for some $\alpha \in \{=, \leqslant\}$.

In the first case, $M_1|_{I' \times J'}$ together with a has kind $\mathbf{C}$, where $a \colon I' \to V(G)$ is defined by $a(i) = M(i, \min(J))$ for $i \in I'$. In the second case, let $a \colon I' \to V(G)$ witness that M_1 is α-capped. Then $M_1|_{I' \times J'}$ together with a has kind $\mathbf{S}$ if α is $=$, and has kind $\mathbf{H}$ if α is $\leqslant$.

Similarly, by Lemma 6.39 applied to $M_{h'}^{\mathsf{T}}$, we conclude that $M_{h'}^{\mathsf{T}}|_{J' \times I'}$ together with some $b \colon J' \to V(G)$ has kind $\mathbf{C}, \mathbf{S}$ or $\mathbf{H}$. We thus conclude that $M_1, \dots, M_{h'}$ induced on $I' \times J'$, together with $a \colon I' \to V(G)$ and $b \colon J' \to V(G)$, form a converter of length h'. □

6.2.4 Crossings

The last step is to go from converters to crossings, whose definition we recall for convenience. For $r \geqslant 1$, the *star r-crossing* of order n is the r-subdivision of $K_{n,n}$. More precisely, it consists of *roots* $a_1, \dots, a_n$ and $b_1, \dots, b_n$ together with r-vertex paths $\{\pi_{i,j} : i, j \in [n]\}$ that are pairwise vertex-disjoint (see Figure 6.6). We denote the two endpoints of a path $\pi_{i,j}$ by $\text{start}(\pi_{i,j})$ and $\text{end}(\pi_{i,j})$. We require that roots appear on no path, that each root a_i is adjacent to $\{\text{start}(\pi_{i,j}) : j \in [n]\}$, and that each root b_j is adjacent to $\{\text{end}(\pi_{i,j}) : i \in [n]\}$. The *clique r-crossing* of order n is the graph obtained from the star r-crossing of order n by turning the neighborhood of each root into a clique. Moreover, we define the *half-graph r-crossing* of order n similarly to the star r-crossing of order n, where each root a_i is instead adjacent to $\{\text{start}(\pi_{i',j}) : i', j \in [n], i \leqslant i'\}$, and each root b_j is instead adjacent to $\{\text{end}(\pi_{i,j'}) : i, j' \in [n], j \leqslant j'\}$. Each of the three r-crossings contains no edges other than the ones described.

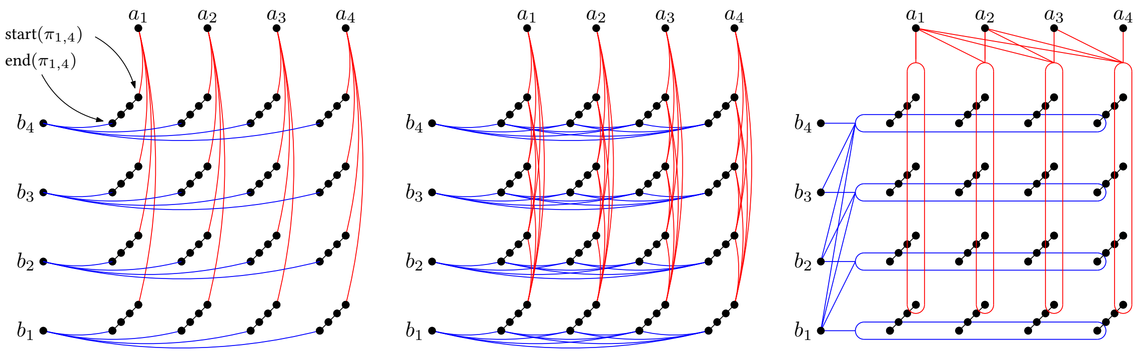

Figure 6.6: (*i*) star 4-crossing of order 4. (*ii*) clique 4-crossing of order 4. (*iii*) half-graph 4-crossing of order 4. In (*iii*), the roots are adjacent to all vertices in the encircled area.

We partition the vertex sets of the r-crossings into *layers*: The 0th layer consists of the vertices $\{a_1, \dots, a_n\}$. The lth layer, for $l \in [r]$, consists of the lth vertices of the paths $\{\pi_{i,j} : i, j \in [n]\}$ (that is, the 1st and rth layer, respectively, are $\{\text{start}(\pi_{i,j}) : i, j \in [n]\}$ and $\{\text{end}(\pi_{i,j}) : i, j \in [n]\}$). Finally, the $(r+1)$th layer consists of the vertices $\{b_1, \dots, b_n\}$. A *flipped* star/clique/half-graph r-crossing is a graph obtained from a star/clique/half-graph r-crossing by performing a flip where the parts of the specifying partition are the layers of the r-crossing.

In Figure 6.6, we present the crossings in a grid-like layout, as opposed to the biclique-like layout shown in Chapter 2, to stress the similarity to converters. We observe that a flipped

r-crossing is the same as a proper converter of kind (α, α) for some $\alpha \in \{\mathbf{C}, \mathbf{S}, \mathbf{H}\}$. So the goal is to show that from a converter of kind (α, β) we can extract a converter with $\alpha = \beta$. This is achieved in the next lemma.

Lemma 6.44. *Let G be a graph containing proper converter C of length h and order n. Then G contains as an induced subgraph either*

- *a flipped star r-crossing, or*
- *a flipped clique r-crossing, or*
- *a flipped half-graph r-crossing,*

of order $\lfloor \sqrt{n} \rfloor$, for some $1 \leqslant r \leqslant 2h + 1$.

Proof. Suppose the meshes $M_1, \ldots, M_h \colon [n] \times [n] \to V(G)$ and functions $a, b \colon [n] \to V(G)$ form a proper converter of kind (α, β) and order n in G. If $(\alpha, \beta) \neq (\mathbf{H}, \mathbf{H})$, then either $\alpha \in \{\mathbf{C}, \mathbf{S}\}$ or $\beta \in \{\mathbf{C}, \mathbf{S}\}$. By replacing $M_1, \ldots, M_h$ with the converter $M_h^\mathsf{T}, \ldots, M_1^\mathsf{T}$ of kind (β, α) if needed, we may assume that either $\beta \in \{\mathbf{C}, \mathbf{S}\}$, or $\alpha = \beta = \mathbf{H}$.

Let $\mathcal{P} = \{V(a), V(M_1), \ldots, V(M_h), V(b)\}$. By taking an induced subgraph if needed, we may assume that $V(G) = \bigcup \mathcal{P}$. It follows from the definition of a proper converter that there is a unique $\mathcal{P}$-flip G' of G with the following properties:

- the meshes M_s and M_t are non-adjacent if $|t - s| > 1$ and are matched if $|t - s| = 1$,
- for $1 < s \leqslant h$, $V(M_s)$ is non-adjacent to $V(a)$,
- for $1 \leqslant s < h$, $V(M_s)$ is non-adjacent to $V(b)$,
- $V(a)$ and $V(b)$ are non-adjacent,

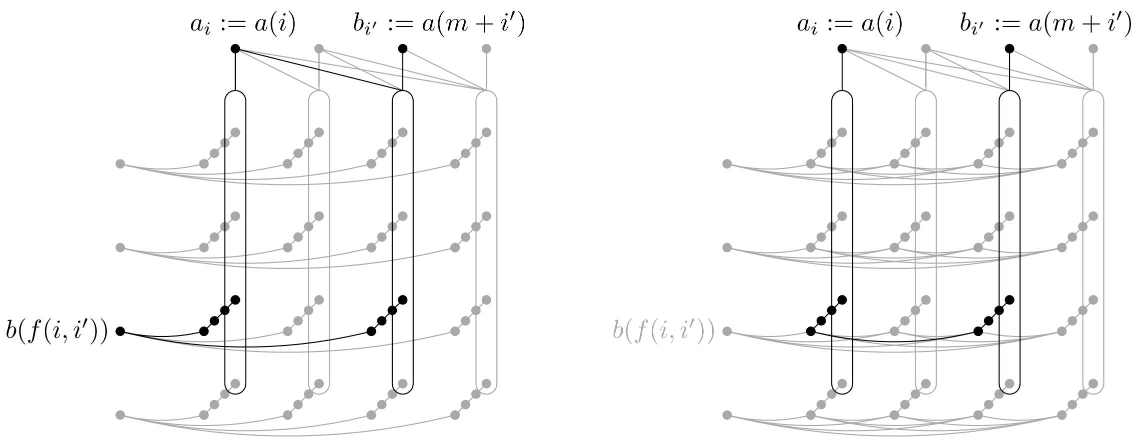

Figure 6.7: Construction of the path $\pi_{i,i'}$ between a_i and $b_{i'}$ in a proper converter of kind $(\mathbf{H}, \beta)$. On the left: the case where $\beta = \mathbf{S}$. On the right: the case where $\beta = \mathbf{C}$.

In the case $\alpha = \beta = \mathbf{H}$ we construct a flipped half-graph h-crossing of order n in G. For this, we choose the roots $a_1 := a(1), \ldots, a_n := a(n)$ on one side, and $b_1 := b(1), \ldots, b_n := b(n)$ on the other side, and connect them via the paths $\pi_{i,j}$ with $\pi_{i,j}[s] = M_1(i, j)$ for $s = 1, \ldots, h$. It follows that G contains a flipped half-graph h-crossing of order n as an induced subgraph.

Consider now the case where $\beta \in \{\mathbf{C}, \mathbf{S}\}$. We observe that

- for all $i, j, j' \in [n]$, $M_h(i, j)$ is adjacent to $b(j')$ if and only if $j = j'$.

Set $\ell := 2h$ if $\beta = \mathbf{C}$ and $\ell := 2h + 1$ if $\beta = \mathbf{S}$. We construct an ℓ-crossing. For all $i, i', j \in [n]$ with $i \neq i'$ consider the unique ℓ-vertex path $\sigma_{i,i',j}$ in G' from $M_1(i, j)$ to $M_1(i', j)$ with the following properties:

- $\sigma_{i,i',j}$ has length ℓ,
- $\sigma_{i,i',j}[d] = M_d(i,j)$ for $d = 1, \dots, h$,
- $\sigma_{i,i',j}[\ell+1-d] = M_d(i',j)$ for $d = 1, \dots, h$,
- if $\beta = \mathbf{S}$ then $\sigma_{i,i',j}[h+1] = b(j)$.

Note that $\sigma_{i,j,j'}$ is an induced path in G'. Let $m = \lfloor\sqrt{n}\rfloor$ and pick any injective function $f\colon [m] \times [m] \to [n]$. For $i, i' \in [m]$, let $\pi_{i,i'} := \sigma_{i,i'+m,f(i,i')}$. Observe that there are no edges in G between any pair of paths $\pi_{i,i'}$, $i, i' \in [m]$. We construct a flipped star/clique/half-graph ℓ-crossing in G, corresponding to the cases $\alpha = \mathbf{S}, \mathbf{C}, \mathbf{H}$, respectively. For this, we choose the roots $a_1 := a(1), \dots, a_m := a(m)$ on one side, and $b_1 := a(m+1), \dots, b_m := a(2m)$ on the other side and connect them via the paths $\pi_{i,i'}$ for $i, i' \in [m]$. Note that if $\alpha = \mathbf{H}$, then $\{a_i : i \in [m]\}$ is fully connected to $\{\mathrm{end}(\pi_{i,i'}) : i, i' \in [m]\}$. The construction is illustrated in Figure 6.7. It follows that G contains a flipped ℓ-crossing of order m as an induced subgraph. □

6.2.5 From Prepatterns to Flipped Crossings and Comparability Grids

We finally prove the main non-structure result.

Proposition 6.45. *Let G be a graph containing a prepattern of order n on an insulator of height h and cost k. Then G contains as an induced subgraph either*

- *a flipped star r-crossing of order m, or*
- *a flipped clique r-crossing of order m, or*
- *a flipped half-graph r-crossing of order m, or*
- *the comparability grid of order m,*

for some $m \geqslant U_{h,k}(n)$ and $1 \leqslant r < 8h$.

Proof. By Proposition 6.9, G contains a transformer of order $U_{k,h}(n)$ and length at most $4h-1$. By Lemma 6.43, G contains a comparability grid of order $U_{k,h}(n)$ as an induced subgraph, or a converter of length at most $4h-1$ and order $U_{k,h}(n)$. In the latter case, we can make the converter proper by Lemma 6.42. By Lemma 6.44, we obtain the required flipped r-crossing of order $U_{k,h}(n)$ for some $r \leqslant 2 \cdot (4h-1) + 1 = 8h-1$. □

As a corollary, we obtain the desired Proposition 6.1 which we restate for convenience.

Proposition 6.1. *Let $\mathcal{C}$ be a graph class that is not prepattern-free. Then there exists $r \in \mathbb{N}$ such that for every $k \in \mathbb{N}$, $\mathcal{C}$ contains as an induced subgraph either*

- *a flipped star r-crossing of order k, or*
- *a flipped clique r-crossing of order k, or*
- *a flipped half-graph r-crossing of order k, or*
- *a comparability grid of order k.*

6.3 Monadic Independence

So far, we have established that each graph class satisfies the following chain of implications.

$$\begin{aligned} &\text{no large flipped crossings/comparability grids} \\ \Rightarrow\ &\text{prepattern-free} \\ \Rightarrow\ &\text{insulation property} \\ \Rightarrow\ &\text{flip-breakable} \\ \Rightarrow\ &\text{monadically dependent} \end{aligned}$$

This corresponds to the implications (3) ⇒ (2) ⇒ (1) of Theorem 2.3, with prepattern-freeness and the insulation property added as extra steps. We now close the circle of implications by showing ¬(3) ⇒ ¬(1).

Proposition 6.46. *Let $\mathcal{C}$ be a graph class and $r \geqslant 1$, such that for every $k \in \mathbb{N}$, $\mathcal{C}$ contains as an induced subgraph*

- *a flipped star r-crossing of order k, or*
- *a flipped clique r-crossing of order k, or*
- *a flipped half-graph r-crossing of order k, or*
- *a comparability grid of order k.*

Then $\mathcal{C}$ is monadically independent.

Proof. Our goal is to show that every graph class exhibiting the listed patterns transduces the class of all graphs. We will be brief, as a more detailed proof of a stronger statement is later given in Proposition 7.3, where we show that the hereditary closure of $\mathcal{C}$ *interprets* the class of all graphs. We use the fact that transductions are transitive: if $\mathcal{C}$ transduces $\mathcal{D}$ and $\mathcal{D}$ transduces $\mathcal{E}$, then already $\mathcal{C}$ transduces $\mathcal{E}$. We transduce the class of all graphs from $\mathcal{C}$ by concatenating multiple simpler transductions, each of which depends only on r (and not on $\mathcal{C}$ or k).

As there is a transduction that produces from $\mathcal{C}$ all its induced subgraphs, we can assume that $\mathcal{C}$ is hereditary. Additionally, for every fixed r there is a fixed transduction that turns flipped r-crossings into their non-flipped versions (or more generally: a transduction that maps $\mathcal{C}$ to all of its $(r+2)$-flips). Now by the pigeonhole principle, we can assume that $\mathcal{C}$ is either the class of all (non-flipped) star/clique/half-graph r-crossings or the class of all comparability grids.

- *Star crossings:* To show that the class of all star r-crossings transduces the class of all bipartite graphs, we describe a transduction that, when given a star r-crossing with roots A and B, creates a bipartite graph (A, B, E) with arbitrary edges E. The transduction colors the roots in A and B with colors C_A and C_B respectively, and the vertices on the paths between $a \in A$ and $b \in B$ with color C_+ if $\{a, b\} \in E$ and color C_- otherwise. It is then trivial to connect $a \in A$ with $b \in B$ by a first-order formula checking if there is a path of color C_+ between them.
- *Clique crossings:* We reduce to the case of star crossings by showing that there is a transduction that turns each clique r-crossing with roots A and B into a star r-crossing of the same order. This is easy to do: For each root $a \in A \cup B$, it suffices to turn its open neighborhood $N(a)$ into an independent set.
- *Half-graph crossings:* We reduce to the case of star crossings by showing that there is a transduction that turns each half-graph r-crossing with roots A and B into a star r-crossing of the same order. Focusing on the side $A = \{a_1, \dots, a_n\}$ first, we observe for $a_i, a_j \in A$ that $i \leqslant j$ if and only if $N(a_i) \supseteq N(a_j)$. The latter condition is expressible in first-order

logic. Therefore, a transduction can remove all edges from a vertex $a_i \in A$ to $\bigcup_{i<j} N(a_j)$. Proceeding similarly for B, yields the desired star crossing.

- *Comparability grids:* We reduce to the case of half-graph crossings. We obtain the half-graph 1-crossings of order $n-1$ with roots A and B from the comparability grid of order n on vertex set $V := \{a_{i,j} : i, j \in [n]\}$ by a transduction as follows.
  1. Delete the vertex $a_{1,1}$.
  2. Clear all the edges with both endpoints in $A := \{a_{i,1} : 1 < i \leqslant n\}$.
  3. Clear all the edges with both endpoints in $B := \{a_{1,i} : 1 < i \leqslant n\}$.
  4. Clear all the edges with both endpoints in $V \setminus (A \cup B)$.

We have shown that $\mathcal{C}$ transduces the class of all bipartite graphs. This class easily transduces the class of all graphs, so $\mathcal{C}$ is monadically independent. □

We summarize the result obtained so far. The following proposition corresponds to the equivalence (1) ⇔ (2) ⇔ (3) of Theorem 2.3.

Proposition 6.47. *Let $\mathcal{C}$ be a graph class. The following are equivalent.*

(1) $\mathcal{C}$ is monadically dependent.

(2) For every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs
  - *all flipped star r-crossings of order k, and*
  - *all flipped clique r-crossings of order k, and*
  - *all flipped half-graph r-crossings of order k, and*
  - *the comparability grid of order k.*

(3) $\mathcal{C}$ is prepattern-free.

(4) $\mathcal{C}$ has the insulation property.

(5) $\mathcal{C}$ is flip-breakable.

Proof. We have already proven all the necessary implications:

- $(1) \Rightarrow (2)$: Proposition 6.46
- $(2) \Rightarrow (3)$: Proposition 6.1
- $(3) \Rightarrow (4)$: Proposition 5.23
- $(4) \Rightarrow (5)$: Proposition 5.24
- $(5) \Rightarrow (1)$: Proposition 5.31 □

We obtain the following algorithmic version of flip-breakability by combining Proposition 6.47 and Proposition 5.24.

Theorem 4.16. *For every monadically dependent graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$, there exists an unbounded function $f_r : \mathbb{N} \to \mathbb{N}$, a constant $k_r \in \mathbb{N}$, and an algorithm that, given a graph $G \in \mathcal{C}$ and $W \subseteq V(G)$, computes in time $O_{\mathcal{C},r}(|V(G)|^2)$ two subsets $A, B \subseteq W$ with $|A|, |B| \geqslant f_r(|W|)$ and a k_r-flip H of G such that:*

$$\mathrm{dist}_H(A, B) > r.$$

Chapter 7

Model Checking Hardness

In this chapter we prove the following hardness result.

Theorem 2.4. *The first-order model checking problem is* AW$[*]$*-hard on every hereditary, monadically independent graph class.*

The above theorem will follow easily, once we have completed our characterization theorem for monadically dependent graph classes, which we restate here for convenience.

Theorem 2.3. *Let $\mathcal{C}$ be a graph class. Then the following are equivalent.*

(1) $\mathcal{C}$ is monadically dependent.

(2) $\mathcal{C}$ is flip-breakable.

(3) For every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs

- *all flipped star r-crossings of order k, and*
- *all flipped clique r-crossings of order k, and*
- *all flipped half-graph r-crossings of order k, and*
- *the comparability grid of order k.*

(4) The hereditary closure of $\mathcal{C}$ does not efficiently interpret the class of all graphs.

In Proposition 6.47 we have already shown the equivalences (1) $\Leftrightarrow$ (2) $\Leftrightarrow$ (3). By Corollary 4.6, the remaining equivalence (1) $\Leftrightarrow$ (4) of Theorem 2.3 implies the hardness result Theorem 2.4. The direction $\neg$(4) $\Rightarrow$ $\neg$(1) is easy to prove.

Lemma 7.1. *Let $\mathcal{C}$ be a graph class such that the hereditary closure of $\mathcal{C}$ efficiently interprets the class of all graphs. Then $\mathcal{C}$ is not monadically dependent.*

Proof. Every class transduces its own hereditary closure. As transductions are more expressive than interpretations, the hereditary closure of $\mathcal{C}$ in particular transduces the class of all graphs. The lemma now follows by transitivity of transductions. $\square$

In particular, if a graph class $\mathcal{C}$ interprets the class of all graphs, then it is not even dependent: the formula that witnesses the interpretation has the independence property on $\mathcal{C}$. We therefore obtain the following result by Braunfeld and Laskowski [9] as a corollary of Theorem 2.3.

Corollary 7.2. *A hereditary graph class is monadically dependent if and only if it is dependent.*

We spend the remainder of this chapter showing $\neg$(3) $\Rightarrow$ $\neg$(4) which completes Theorem 2.3.

Proposition 7.3. *Let $\mathcal{C}$ be a hereditary graph class and $r \geqslant 1$, such that for all $k \in \mathbb{N}$, $\mathcal{C}$ contains*

- *a comparability grid of order k, or*
- *a flipped star r-crossing of order k, or*
- *a flipped clique r-crossing of order k, or*
- *a flipped half-graph r-crossing of order k.*

Then $\mathcal{C}$ efficiently interprets the class of all graphs.

Note that it is rather straightforward to prove that, under the assumptions of Proposition 7.3, some class $\mathcal{C}^+$ of *colored* graphs from $\mathcal{C}$ efficiently interprets the class of all graphs, by analyzing the proof of Proposition 6.46. This, however, would yield a significantly weaker statement than Theorem 2.4: that the first-order model checking problem on *colored* graphs from $\mathcal{C}$ is AW[$*$]-hard, for every hereditary, monadically independent class $\mathcal{C}$.

7.1 Crossings and Comparability Grids

Since we explicitly refer to the individual vertices of our r-crossing patterns in this section, let us restate their definition in greater detail and explicitly name their vertex sets.

Definition 7.4 (r-crossings)**.** For every radius $r \geqslant 1$ we define the *star, clique, and half-graph r-crossing* of order n as the graph whose vertex set

$$\{a_i : i \in [n]\} \cup \{b_i : i \in [n]\} \cup \{p_{i,j,t} : i, j \in [n], t \in [r]\}$$

is partitioned into $l := r + 2$ *layers* $\mathcal{L} := \{L_0, \ldots, L_{r+1}\}$ with

- $L_0 := \{a_i : i \in [n]\}$,
- $L_t := \{p_{i,j,t} : i, j \in [n]\}$ for all $t \in [r]$,
- $L_{r+1} := \{b_j : j \in [n]\}$,

and whose edges are defined as follows. The vertices $(a_i, p_{i,j,1}, \ldots, p_{i,j,r}, b_j)$ form a path for all $i, j \in [n]$. Each of the three types enforces separate additional edges.

- The star r-crossing contains no additional edges.
- For the clique r-crossing,
  - $p_{i,j,1}$ and $p_{i,j',1}$ are adjacent for all $j \neq j' \in [n]$, and
  - $p_{i,j,r}$ and $p_{i',j,r}$ are adjacent for all $i \neq i' \in [n]$.
- For the half-graph r-crossing,
  - a_i is adjacent to $p_{i',j,1}$ for all $i \leqslant i' \leqslant n$ and for all $j \in [n]$, and
  - b_j is adjacent to $p_{i,j',r}$ for all $i \in [n]$ and for all $j \leqslant j' \leqslant n$.

For every $r \geqslant 1$, let $\mathcal{B}_r^\star$, $\mathcal{B}_r^\bullet$, $\mathcal{B}_r^\blacktriangleleft$ be the hereditary closure of the class of all star, clique, and half-graph r-crossings, respectively. The definition of layers naturally generalizes to the induced subgraphs contained in $\mathcal{B}_r^\star$, $\mathcal{B}_r^\bullet$, $\mathcal{B}_r^\blacktriangleleft$. As we have explicitly named the vertices, their partition into layers is unique.

Lastly, recall that the *comparability grid* of order n consists of vertices $\{a_{i,j} : i, j \in [n]\}$ and edges between vertices $a_{i,j}$ and $a_{i',j'}$ if and only if either $i = i'$, or $j = j'$, or $i < i' \Leftrightarrow j < j'$. Denote by $\mathcal{B}^\blacksquare$ the hereditary closure of the class of all comparability grids.

7.2 Twins

Two vertices u and v are *twins* in a graph G, if $N^G(u) \setminus \{u, v\} = N^G(v) \setminus \{u, v\}$. This relation is transitive and definable in first-order logic:

$$\text{twins}(x, y) := \forall z : (z \neq x \wedge z \neq y) \to (E(z, x) \leftrightarrow E(z, y)).$$

The *twin classes* of a graph G are the equivalence classes of the twin-relation of G. For every fixed $k \in \mathbb{N}$, the formulas

$$\#\text{twins}_{\geqslant k}(x) := \exists z_1, \dots, z_k : \bigwedge_{(i,j) \in \binom{k}{2}} z_i \neq z_j \wedge z_i \neq x \wedge \text{twins}(z_i, x),$$
$$\#\text{twins}_{=k}(x) := \#\text{twins}_{\geqslant k}(x) \wedge \neg(\#\text{twins}_{\geqslant k+1}(x))$$

express that x has at least or exactly k twins, respectively (equivalently, the twin class containing x has at least or exactly $k + 1$ elements).

7.3 Reversing Flips

Combining hereditariness and the pigeonhole principle, we observe the following.

Observation 7.5. Let $\mathcal{C}$ be a hereditary graph class and $r \geqslant 1$, such that for all $k \in \mathbb{N}$, $\mathcal{C}$ contains

- a comparability grid of order k, or
- a flipped star r-crossing of order k, or
- a flipped clique r-crossing of order k, or
- a flipped half-graph r-crossing of order k.

Then $\mathcal{C}$ contains either $\mathcal{B}^{\blacksquare}$ or there is $\mathcal{B} \in \{\mathcal{B}_r^{\star}, \mathcal{B}_r^{\bullet}, \mathcal{B}_r^{\blacktriangleleft}\}$, such that $\mathcal{C}$ contains a *layer-wise* flip of each graph in $\mathcal{B}$.

Recall that the class $\mathcal{B}_r^{\star}$ contains all induced subgraphs of all star r-crossings. A *layer-wise* flip of a graph G in $\mathcal{B}_r^{\star}$ is an $\mathcal{L}$-flip of G: the flip respects the layered structure of the class $\mathcal{B}_r^{\star}$ (cf. Definition 7.4). We define layer-wise flips of graphs in $\mathcal{B}_r^{\bullet}$ and $\mathcal{B}_r^{\blacktriangleleft}$ in the same way.

In this subsection we use interpretations to undo the flips and recover the graphs from $\mathcal{B}$. However, we only recover graphs without twins and without isolated vertices. Let $\mathcal{T}$ be the class of all graphs containing twins. Let $\mathcal{I}$ be the class of all graphs containing isolated vertices.

Lemma 7.6. *Fix $r \geqslant 1$ and $\mathcal{B} \in \{\mathcal{B}_r^{\star}, \mathcal{B}_r^{\bullet}, \mathcal{B}_r^{\blacktriangleleft}\}$. Let $\mathcal{C}$ be a hereditary class containing a layer-wise flip of each graph from $\mathcal{B}$. Then $\mathcal{C}$ efficiently interprets $\mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$.*

Proof. The class $\mathcal{C}$ contains a layer-wise flip of each graph from $\mathcal{B}$. By hereditariness of $\mathcal{C}$ and the pigeonhole principle, we can assume that the layers in all those graphs are flipped "in the same way". Formally, there exists an integer $k \leqslant r + 2$, a mapping $\text{lc} : \{0, \dots, r + 1\} \to [k]$ (for **l**ayer **c**olor), and a symmetric relation $R \subseteq [k]^2$ with the following property.

For every $G \in \mathcal{B}$ the graph $\text{flip}(G) := G \oplus R|_G$ is contained in $\mathcal{C}$.

Here the relation $R|_G$ is defined as follows. Let $\mathcal{L} = \{L_0, \dots, L_{r+1}\}$ be the layers of G as in Definition 7.4. Let $\mathcal{K}_G = \{C_1, \dots, C_k\}$ be the unique k-coloring of $V(G)$ satisfying $L_i \subseteq C_{\text{lc}(i)}$ for all $L_i \in \mathcal{L}$. Each layer is monochromatic, but some layers may have the same color. We define $R|_G \subseteq \mathcal{K}_G^2$ such that $(C_i, C_j) \in R|_G$ if and only if $(i, j) \in R$ for all $i, j \in [k]$. Without loss of generality, we can assume k is chosen minimal in the following sense.

- Every color is used: the map lc is surjective.
- No two colors can be merged: for all $i \neq j \in [k]$ there exists $d \in [k]$ such that

$$(i, d) \in R \Leftrightarrow (j, d) \notin R.$$

From now on, we assume that each graph $G \in \mathcal{B}$ is implicitly k-colored according to the coloring $\mathcal{K}_G$ defined above. We next use twin classes to uniquely quantify representative vertices of each color class. The representatives are added as follows. For any k-colored graph G, let $\operatorname{prep}(G)$ be the graph obtained from G by adding, for each color $i \in [k]$, $(i+1)$ many isolated vertices $s_{i,1}, \ldots, s_{i,(i+1)}$ of color i to G.

Claim 7.7. *If $G \in \mathcal{B}$, then also* $\operatorname{prep}(G) \in \mathcal{B}$.

Proof. For every $n \in \mathbb{N}$ let B_n denote the star/clique/half-graph r-crossing of order n. As $G \in \mathcal{B}$, there exists an embedding f from G to B_n, for some $n \in \mathbb{N}$. Let $l := r + 2$ be the number of layers in an r-crossing. We define $c := r \cdot (l+1)$ and $m := n + 2c$. We show how to embed $\operatorname{prep}(G)$ into B_m. First note that the function

$$g(\cdot) := \bigcup_{i,j \in [n], t \in [r]} \{a_i \mapsto a_{i+c}, \quad b_j \mapsto b_{j+c}, \quad p_{i,j,t} \mapsto p_{i+c,j+c,t}\}$$

is an embedding from B_n to B_m. It follows that $h(\cdot) := g \circ f$ is an embedding from G to B_m. Let $h(G) := \{h(v) : v \in V(G)\}$. Importantly, every vertex a_i, b_j, or $p_{i,j,t}$ in $h(G)$ satisfies

$$i, j \in \{c+1, \ldots, c+n\} \quad \text{and} \quad t \in [k].$$

We next choose a set $S(L)$ of $(l+1)$ vertices from each layer $L \in \{L_0, \ldots, L_{r+1}\}$ as follows.

- $S(L_0) := \{a_{c+n+i} : i \in [l+1]\}$,
- $S(L_t) := \{p_{i'+i,i'+i,t} : i \in [l+1]\}$ for every $t \in [r]$ and $i' := (t-1) \cdot (l+1)$,
- $S(L_{r+1}) := \{b_{c+n+j} : j \in [l+1]\}$.

See Figure 7.1 for a visualization.

We define $I := S(L_0) \cup \ldots \cup S(L_{r+1})$. By construction, I and $h(G)$ are disjoint. Let us now argue that every vertex from I is isolated in the induced subgraph $B'_m := B_m[h(G) \cup I]$.

- Let $a_i \in I$. Then $i > c + n$. All the neighbors of a_i in B_m are of the form $p_{i',j,t}$ for some $i' \geqslant i$. All vertices $p_{i',j,t}$ in B'_m satisfy $i' \leqslant c + n$, so they are non-adjacent to a_i.
- Let $b_j \in I$. The same reasoning as in the previous case applies.
- Let $p_{i,j,t} \in I$. Then $i \leqslant c$ and $j \leqslant c$. By the same reasoning as before, $p_{i,j,t}$ is non-adjacent to all vertices of the form $a_{i'}$ or $b_{j'}$ in B'_m. Furthermore, any neighbor of $p_{i,j,t}$ of the form $p_{i',j',t'}$ in B_m must satisfy $i' = i$ or $j' = j$. By construction, B'_m contains no vertex $p_{i',j',t'}$ satisfying $i' = i$ or $j' = j$, apart from $p_{i,j,t}$ itself.

This proves that the set I is indeed isolated in B'_m. We want to stress that our argument works for each of the three classes $\mathcal{B} \in \{\mathcal{B}^{\star}_r, \mathcal{B}^{\bullet}_r, \mathcal{B}^{\blacktriangleleft}_r\}$.

We finally embed $\operatorname{prep}(G)$ into B'_m. The vertices of G are embedded using h. The additional vertices $s_{i,1}, \ldots, s_{i,(i+1)}$ that are added to $\operatorname{prep}(G)$ for each color $i \in [k]$ can now be mapped to distinct vertices from a set $S(L)$ to which the layer coloring lc assigns color i. As desired, they have color i and are isolated. ∎

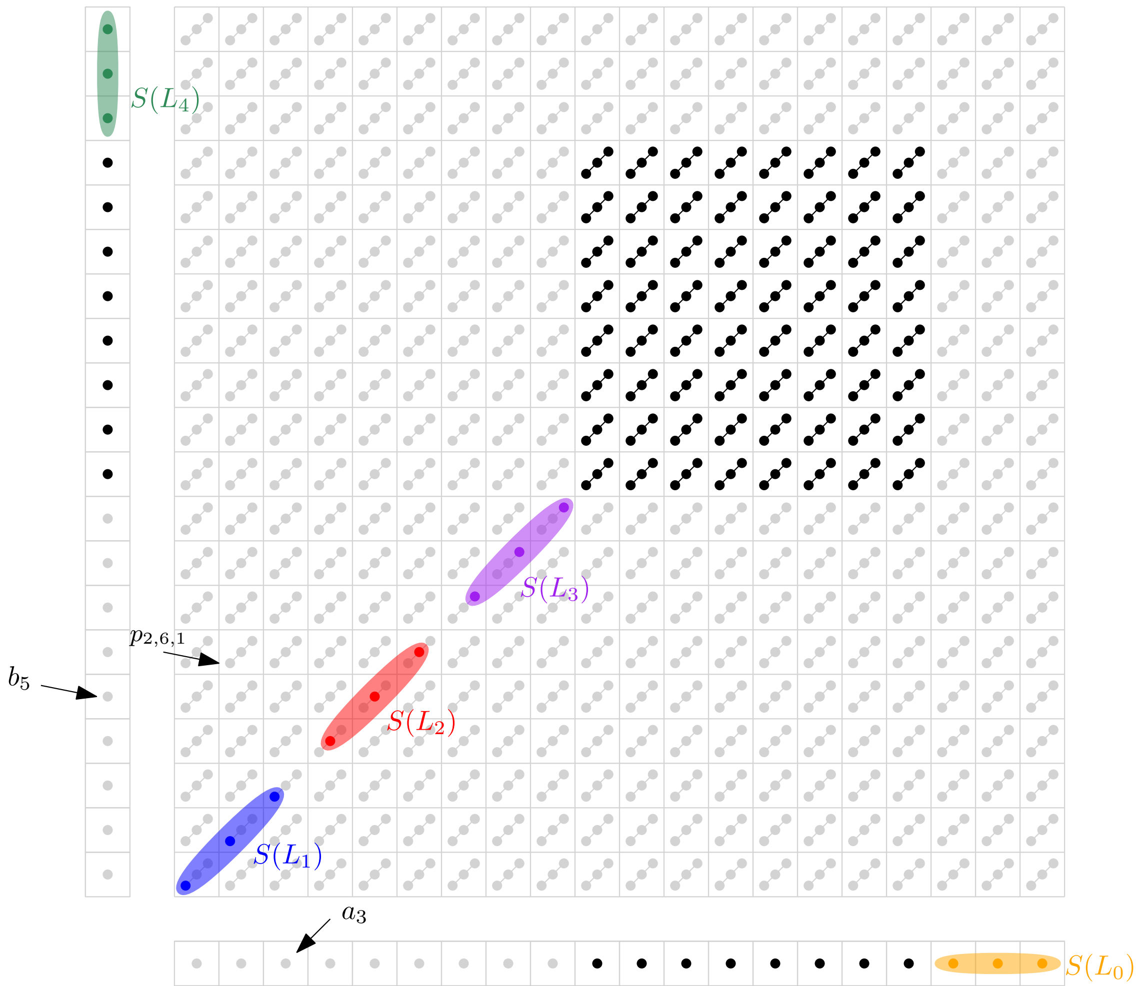

Figure 7.1: A visualization of how the sets $S(L)$ are embedded into B_m for the case $r = 3$. To preserve readability in the visualization, each set $S(L)$ has size 3 instead of $l + 1 = 6$. The black vertices correspond to the embedding of B_n into B_m.

Claim 7.8. *Let $G \in \mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$. The twin classes of* flip(prep(G)) *consist exactly of*

- *the singleton twin class $\{v\}$ for every vertex $v \in V(G)$, and*
- *the twin class $T_i := \{s_{i,1}, \dots, s_{i,i+1}\}$ for every color $i \in [k]$.*

Proof. The following is easy to see.

$$\begin{gathered}\text{For all vertices } u, v \text{ with the same color:}\\ u \text{ and } v \text{ are twins in prep}(G) \text{ if and only if they are twins in flip(prep}(G)).\end{gathered} \tag{7.1}$$

Additionally, we argue the following.

$$\text{For all vertices } u, v \text{ with different colors: } u \text{ and } v \text{ are not twins in flip(prep}(G)). \tag{7.2}$$

Let $i \neq j$ be the colors of u and v. By the assumed minimality of the coloring, there exist a color $d \in [k]$ such that $(i, d) \in R \Leftrightarrow (j, d) \notin R$. There exists at least one vertex $s_d \in \{s_{d,1}, s_{d,2}\}$ that has color d and is non-adjacent and non-equal to both u and v in prep(G). It follows that in flip(prep(G)) exactly one of u and v will be adjacent to s_d. Thus, u and v are no twins in flip(prep(G)).

Combining (7.1) and (7.2), we have that every two vertices u and v which are no twins in prep(G) are also no twins in flip(prep(G)). Since $G \notin \mathcal{T} \cup \mathcal{I}$, in prep($G$) the vertices of $V(G)$ neither have twins among $V(G)$ nor among the isolated vertices added to build prep(G) from G. It follows that each vertex from $V(G)$ is contained in a singleton twin class of flip(prep(G)) as desired. Finally, by (7.1), for every color $i \in [k]$ there is a twin class T_i containing the set of isolated vertices $\{s_{i,1}, \dots, s_{i,i+1}\}$. As argued before, T_i contains no vertices from $V(G)$. By (7.2), T_i is disjoint from T_j for every other color $j \neq i$. Then T_i is exactly $\{s_{i,1}, \dots, s_{i,i+1}\}$, as desired. ■

Claim 7.9. *For every color $i \in [k]$ there exists a formula* $\text{col}_i(x)$ *such that for every $G \in \mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$ and every vertex v in* prep(G) *we have*

$$v \textit{ has color } i \quad \Leftrightarrow \quad \text{flip(prep}(G)) \models \text{col}_i(x).$$

Proof. We argue that the following formula does the job.

$$\text{col}_i(x) := \exists z_1, \dots, z_k : \bigwedge_{j \in k} x \neq z_j \wedge \#\text{twins}_{=j}(z_j) \wedge \big(E(x, z_j) \leftrightarrow (i, j) \in R\big)$$

The formula quantifies vertices $\bar{z} = z_1 \dots z_k$ containing for each color $j \in [k]$ a vertex z_j such that

- z_j is not equal to x,
- z_j is from a twin class of size exactly $j + 1$,
- z_j is adjacent to x if and only if $(i, j) \in R$.

Let v be a vertex in prep(G). To prove the forwards direction of the claim, assume v has color i. We can choose a satisfying valuation $\bar{w}$ of $\bar{z}$ as follows. By Claim 7.8, for each color $j \in [k]$ the twin class T_j has size exactly $j + 1$ and all its vertices are isolated in prep(G) and have color j. As $|T_j| \geqslant 2$, we can pick a vertex $w_j \in T_j$ that is not equal to v. As v and w_j are non-adjacent in prep(G) and of color i and j respectively, we have flip(prep(G)) $\models E(v, w_j) \Leftrightarrow (i, j) \in R$ as desired.

For the backwards direction, assume towards contradiction that v has color $i' \neq i$ and there exists a satisfying valuation $\bar{w}$ of $\bar{z}$. By the assumed minimality of the coloring of prep(G), there exist a color $d \in [k]$ such that

$$(i', d) \in R \Leftrightarrow (i, d) \notin R.$$

Again w_d has color d and is non-adjacent to v in prep(G). By definition of flip(prep(G)) we have

$$\text{flip}(\text{prep}(G)) \models E(v, w_d) \Leftrightarrow (i', d) \in R.$$

However, for $\bar{w}$ to be a satisfying valuation of $\bar{z}$ we must have

$$\text{flip}(\text{prep}(G)) \models E(v, w_d) \Leftrightarrow (i, d) \in R.$$

Combining the three equivalences gives the desired contradiction. ■

Claim 7.10. *For every formula $\varphi(\bar{x})$ we can compute a formula* flip$(\varphi)(\bar{x})$ *such that for every graph $G \in \mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$ and every tuple $\bar{a} \in V(G)^{|x|}$,*

$$\text{prep}(G) \models \varphi(\bar{a}) \Leftrightarrow \text{flip}(\text{prep}(G)) \models \text{flip}(\varphi)(\bar{a}).$$

Proof. Using Claim 7.9, it is easy to see that for all graphs $G \in \mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$ and vertices u and v in prep(G),

$$\text{prep}(G) \models E(u, v) \quad \Leftrightarrow \quad \text{flip}(\text{prep}(G)) \models E(u, v) \ \text{ XOR } \bigvee_{i,j \in [k]} \text{col}_i(x) \wedge \text{col}_j(y) \wedge (i, j) \in R.$$

For every $\varphi(\bar{x})$, let flip$(\varphi)(\bar{x})$ be the formula obtained by replacing every occurrence of $E(x, y)$ with the formula on the right side of the above equivalence. It now easily follows by structural induction that flip$(\varphi)(\bar{x})$ has the desired properties. ■

Let $\delta(x) :=$ flip(hasNeighbor)(x) and $\varphi(x, y) :=$ flip$(E)(x, y)$, where hasNeighbor(x) is the formula checking that x is not an isolated vertex. Using Claim 7.10, we have

$$I_{\delta,\varphi}\big(\text{flip}(\text{prep}(G))\big) = G$$

for every graph $G \in \mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$. As flip(prep)$(G)$ is contained in $\mathcal{C}$ and can be computed in polynomial time from G, we have that $\mathcal{C}$ efficiently interprets $\mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$. □

7.4 Encoding Bipartite Graphs

Having undone the flips, we interpret all bipartite graphs (without isolated vertices) from our intermediate classes $\{\mathcal{B}_r^{\star}, \mathcal{B}_r^{\bullet}, \mathcal{B}_r^{\blacktriangleleft}, \mathcal{B}^{\blacksquare}\}$. Targeting this class of bipartite graphs does not restrict the general case, as the following lemma shows.

Lemma 7.11. *The class of all bipartite graphs without isolated vertices efficiently interprets the class of all graphs.*

Proof. Let $\delta(x)$ be the formula stating that x has degree at least three and $\varphi(x, y)$ be the formula stating that x and y are at distance exactly two. For every graph G, we build the graph B_G as follows. For every vertex $v \in V(G)$ we create a star with three leaves and center c_v. For every edge $(u, v) \in E(G)$, we add a new vertex adjacent to both c_u and c_v. It is easy to see that B_G is bipartite, without isolated vertices, and $I_{\delta,\varphi}(B_G) = G$. □

The following notation will be convenient. For every bipartite graph H there exists at least one *bipartite representation of* H, that is, a tuple

$$H' = (U' \subseteq \mathbb{N}, V' \subseteq \mathbb{N}, E(H') \subseteq U' \times V'),$$

such that there exist

- a bipartition of $V(H)$ into two independent sets U and V, and
- two bijections $f : U \to U'$ and $g : V \to V'$,

such that for all $u \in U$ and $v \in V$: $(f(u), f(v)) \in E(H') \Leftrightarrow (u, v) \in E(H)$. Note that U' and V' do not have to be disjoint and $E(H')$ is not necessarily symmetric.

Encoding Bipartite Graphs in Star and Clique r-Crossings

Lemma 7.12. *For every* $r \geqslant 1$ *and* $\mathcal{B} \in \{\mathcal{B}_r^{\star}, \mathcal{B}_r^{\bullet}\}$,

$$\mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I}) \text{ efficiently interprets the class of all bipartite graphs.}$$

Proof. First assume $\mathcal{B} = \mathcal{B}_r^{\star}$. Let $\delta(x)$ be the formula checking whether x has degree at least three, and let $\varphi(x, y)$ be the formula checking whether the distance between x and y is exactly $r + 1$. To prove the lemma, we show that for every bipartite graph H, we can construct a graph $B_H \in \mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$ such that $I_{\delta,\varphi}(B_H) = H$. Let

$$H' = ([n], [m], E(H') \subseteq [n] \times [m])$$

be a bipartite representation of H for some $n, m \in \mathbb{N}$. We build the graph B_H as follows. For every $i \in [n]$ we create a 1-subdivided star with three leaves consisting of: a center c_i and for every $s \in \{0, 1, 2\}$ a subdivision vertex $c_{i,s,1}$ and a leaf $c_{i,s,2}$ such that $(c_i, c_{i,s,1}, c_{i,s,2})$ form a path. We do the same for every $j \in [m]$, giving us vertices $d_j, d_{j,s,1}, d_{j,s,2}$ for every $s \in \{0, 1, 2\}$. Finally, for every edge $(i, j) \in E(H')$ we add vertices $\{q_{i,j,t} : t \in [r]\}$ and connect $(c_i, q_{i,j,1}, \ldots, q_{i,j,r}, d_j)$ to form a path of length $r + 1$. It is easy to see that $I_{\delta,\varphi}(B_H) = H$ and that B_H contains neither twins nor isolated vertices. It remains to show that B_H is an induced subgraph of a star r-crossing.

Let $N := 3(n + m)$ and B_N be the star r-crossing of order N. We give an embedding $h : V(B_H) \to V(B_N)$ of B_H into B_N. Let $f(i) := 3i - 2$ and $g(j) := 3n + 3j - 2$. We define h as follows for all $i \in [n]$, $j \in [m]$, $s \in \{0, 1, 2\}$, and $t \in [r]$.

- $h(c_i) := a_{f(i)}$
- $h(c_{i,s,1}) := p_{f(i),f(i)+s,1}$
- $h(c_{i,s,2}) := \begin{cases} p_{f(i),f(i)+s,2} & \text{if } r > 1 \\ b_{f(i)+s} & \text{if } r = 1 \end{cases}$
- $h(q_{i,j,t}) := p_{f(i),g(j),t}$
- $h(d_j) := b_{g(j)}$
- $h(d_{j,s,1}) := p_{g(j)+s,g(j),r}$
- $h(d_{j,s,2}) := \begin{cases} p_{g(j)+s,g(j),r-1} & \text{if } r > 1 \\ a_{g(j)+s} & \text{if } r = 1 \end{cases}$

See Figure 7.2 for a visualization of the embedding. It is easily checked that h is an embedding. This finishes the case $\mathcal{B} = \mathcal{B}_r^{\star}$.

For the case $\mathcal{B} = \mathcal{B}_r^{\bullet}$, we take the same edge formula φ and update the domain formula to state

$$\delta(x) := \text{“}x \text{ has degree at least three and the neighborhood of } x \text{ is a clique”}.$$

We build B_H as in the previous case but add additional edges. For every $i \in [n]$ we turn the set

$$\bigcup \big\{ \{c_{i,s,1}, c_{i,s,2}, q_{i,j,1}\} : j \in [m], (i, j) \in E(H'), s \in \{0, 1, 2\} \big\}$$

that contains exactly the neighborhood of c_i into a clique. Symmetrically, we do the same for every $j \in [m]$ with the set

$$\bigcup \big\{ \{d_{j,s,1}, d_{j,s,2}, q_{i,j,r}\} : i \in [n], (i, j) \in E(H'), s \in \{0, 1, 2\} \big\}$$

of all neighbors of d_j. Again $I_{\delta,\varphi}(B_H) = H$, B_H contains neither twins nor isolated vertices, and h is an embedding of B_H into the clique r-crossing of order $3(n + m)$. □

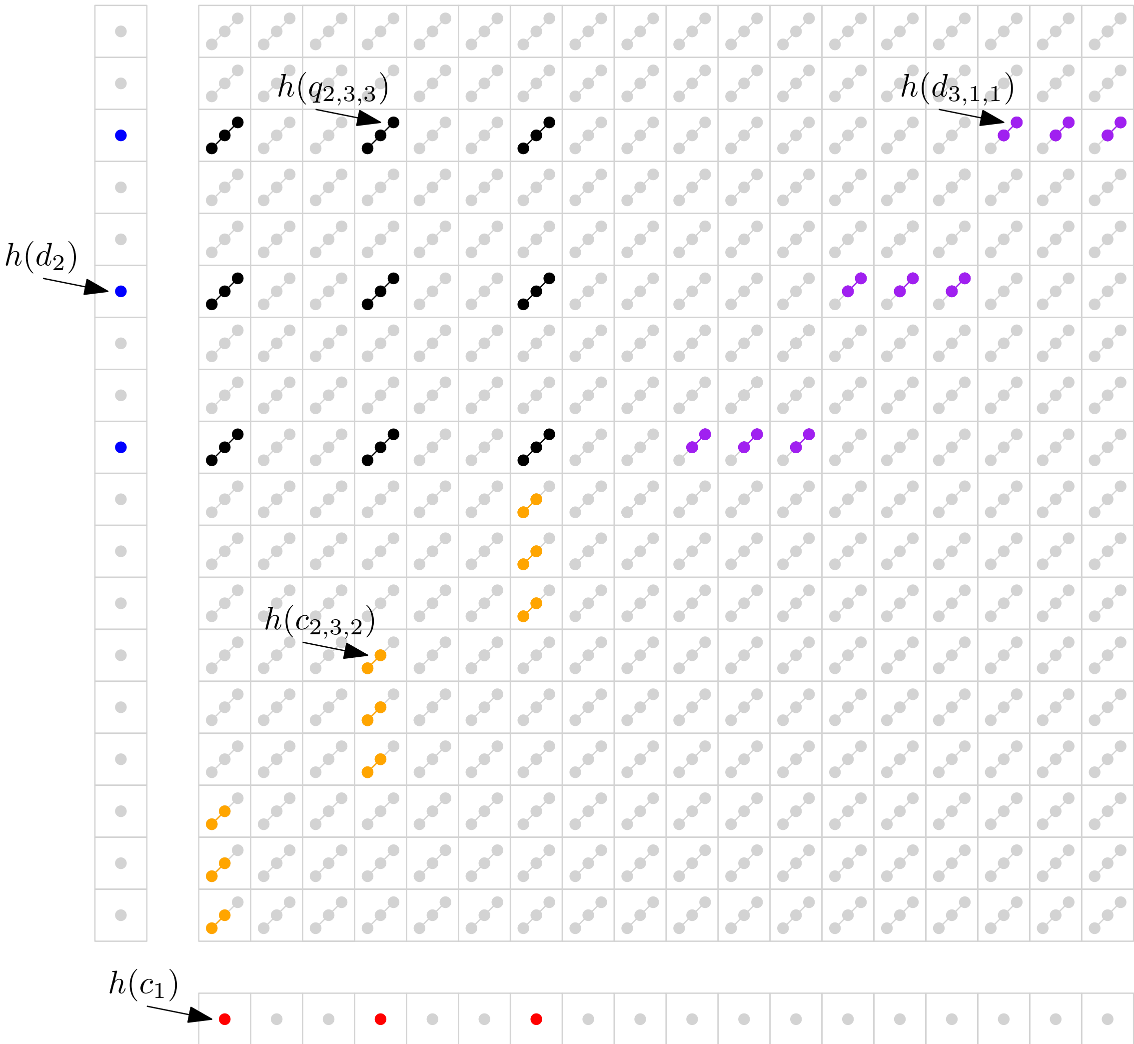

Figure 7.2: A visualization of how B_H embeds into B_N for the case where $r = 3$ and H is the biclique of order 3.

Encoding Bipartite Graphs in Half-Graph r-Crossings and in Comparability Grids

Lemma 7.13. *For every $r \geqslant 1$ and $\mathcal{B} \in \{\mathcal{B}_r^{\blacktriangleleft}, \mathcal{B}^{\blacksquare}\}$,*

$\mathcal{B} \setminus (\mathcal{T} \cup \mathcal{I})$ *efficiently interprets the class of all bipartite graphs without isolated vertices.*

Proof. We first prove the statement for $\mathcal{B} = \mathcal{B}_r^{\blacktriangleleft}$. For every $n \in \mathbb{N}$, we define $[\![n]\!] := [n] \setminus \{1, n\} = \{2, \ldots, n-1\}$. Let H be an arbitrary bipartite graph without isolated vertices and let

$$H' = ([\![n]\!], [\![m]\!], E(H') \subseteq [\![n]\!] \times [\![m]\!])$$

be a bipartite representation of H for some $n, m \in \mathbb{N}$. We define B_H to be the subgraph of the half-graph r-crossing of order $\max(n, m)$ induced by the vertices

- $A := \{a_i : i \in [\![n]\!]\}$ corresponding to the left vertices of H,
- $B := \{b_j : j \in [\![m]\!]\}$ corresponding to the right vertices of H,
- $P := \{p_{i,j,t} : (i, j) \in E(H'), t \in [r]\}$ corresponding to the edges of H,
- $\{a_n, p_{n,1,1}\}$ and $\{b_m, p_{1,m,r}\}$ which we use as auxiliary vertices.

See the left side of Figure 7.3 for a visualization.

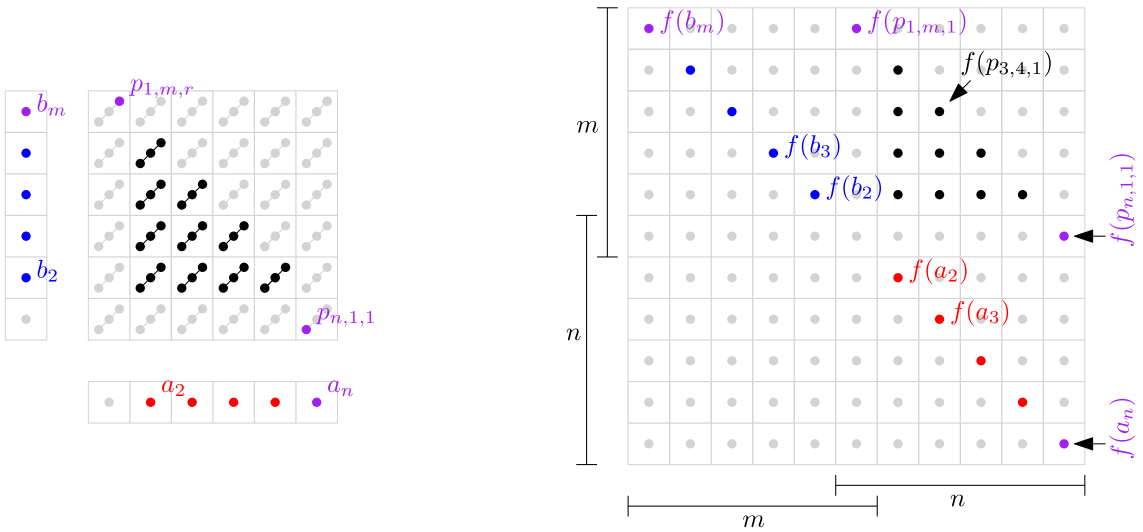

Figure 7.3: On the left: a visualization of the vertices of B_H where H is the half-graph of order 4 and $r = 3$. The vertices in A, B, and P are colored red, blue, and black respectively. On the right: a visualization of the embedding of $B_H^{\blacksquare}$ into a comparability grid, where again H is the half-graph of order 4.

Our goal is to interpret H from B_H. Let

$$A_\star := A \cup \{a_n\}, \quad B_\star := B \cup \{b_m\}, \quad P_\star := P \cup \{p_{n,1,1}, p_{1,m,r}\}.$$

Claim 7.14. $N_{B_H}(p_{n,1,1}) = A_\star$ *and* $N_{B_H}(p_{1,m,r}) = B_\star$.

Proof. The claim follows by the definition of the adjacencies in half-graph crossings. $A_\star$ is included in the neighborhood of $p_{n,1,1}$ because $i \leqslant n$ for all $i \in [\![n]\!] \cup \{n\}$. All neighbors of $B_\star$ are of the form $p_{i,j,t}$ for some $j \geqslant 2$, so no vertex of $B_\star$ is included in the neighborhood of $p_{n,1,1}$. The vertex $p_{n,1,2}$ is not contained in $P_\star$, so no vertex of $P_\star$ is included in the neighborhood of $p_{n,1,1}$. We argue symmetrically to determine the neighborhood of $p_{1,m,r}$. ■

Observation 7.15. $N_{B_H}(a_n) = \{p_{n,1,1}\}$ and $N_{B_H}(b_m) = \{p_{1,m,r}\}$.

Claim 7.16. *Except for a_n and b_m, all vertices have degree at least two.*

Proof. All vertices in A are adjacent to $p_{n,1,1}$ and at least one other $p_{i,j,1}$ vertex since H contains no isolated vertices. A symmetric statement holds for B and $p_{1,m,r}$. The vertices in P are inner vertices of paths from A to B, so they have degree at least two. The vertices $p_{n,1,1}$ and $p_{1,m,r}$ have high degree by the previous claim. ∎

Observation 7.17. B_H contains no isolated vertices.

Claim 7.18. *B_H contains no twins.*

Proof. The vertices in $A_\star$ can be distinguished from the vertices in $B_\star$ and $P_\star$ by their adjacency to $p_{n,1,1}$. Two vertices a_i and $a_{i'}$ from $A_\star$ with $i < i' \leqslant n$ can be differentiated by a vertex $p_{i,j,1}$ for some $j \in [\![m]\!]$ which is adjacent to a_i but not to $a_{i'}$. This j exist since no vertex in H is isolated. It follows that $A_\star$ contains no twins and by a symmetric argument, neither does $B_\star$.

It remains to distinguish the vertices inside $P_\star$. The auxiliary vertices $p_{n,1,1}$ and $p_{1,m,r}$ each have a private neighbor in a_n and b_m, so we can focus our attention on the set P. Let $p_{i,j,t}$ and $p_{i',j',t'}$ be two distinct vertices from P. By symmetry, we can assume that either $i < i'$ or $j < j'$ or $t < t'$. If $i < i'$ or $t < t'$ then the vertex

$$d := \begin{cases} p_{i,j,t-1} & \text{if } t > 1, \\ a_i & \text{if } t = 1, \end{cases}$$

is adjacent to $p_{i,j,t}$ but non-adjacent to $p_{i',j',t'}$. If $j < j'$ we argue symmetrically using either $p_{i,j,t+1}$ or b_j if $t = r$. We want to stress that the argument works also for the case of $r = 1$. ∎

We have proven that $B_H \in \mathcal{B}_r^{◂} \setminus (\mathcal{T} \cup \mathcal{I})$. Let us continue to show that B_H interprets H. Combining Observation 7.15, Claim 7.14, Claim 7.16, we have that the formula

$$\delta(x) := \text{“}x \text{ has degree at least 2 and is at distance exactly 2 from a degree 1 vertex”}$$

is true exactly on the vertices from $A \cup B$. It can therefore act as the domain formula of our interpretation. It remains to define the edge relation. First notice that the formula

$$\text{sameSide}(x, y) := \exists z : \text{“}z \text{ has degree 1 and is at distance exactly 2 from both } x \text{ and } y\text{”}$$

distinguishes $A_\star$ and $B_\star$: for all vertices u and v

$$B_H \models \text{sameSide}(u, v) \quad \Leftrightarrow \quad \{u, v\} \in A_\star \vee \{u, v\} \in B_\star.$$

We next construct a formula that resolves the half-graphs between A and $P_\star$ and between B and $P_\star$ in the following sense.

Claim 7.19. *There exists a formula $E_0(x, y)$ such that for all $i \in [\![n]\!]$,*

$$\{v \in V(B_H) : B_H \models E_0(a_i, v)\} = \{p_{i,j,1} : (i, j) \in E(H'), j \in [\![m]\!]\},$$

and for all $j \in [\![m]\!]$,

$$\{v \in V(B_H) : B_H \models E_0(b_j, v)\} = \{p_{i,j,r} : (i, j) \in E(H'), i \in [\![n]\!]\}.$$

Proof. The formula

$$x \prec y := \text{sameSide}(x, y) \wedge \text{“}N(y) \text{ is a strict subset of } N(x)\text{”}$$

orders $A_\star$ and $B_\star$ respectively: for all $a_i \in A_\star$ and $b_j \in B_\star$ we have

- $\big(B_H \models a_i \prec a\big) \Leftrightarrow \big(a \in \{a_{i'} : i < i' \leqslant n\}\big)$, and
- $\big(B_H \models b_j \prec b\big) \Leftrightarrow \big(b \in \{b_{j'} : j < j' \leqslant m\}\big)$.

Now the following formula

$$E_0(x, y) := E(x, y) \land \neg\big(\exists x' : x \prec x' \land E(x', y)\big)$$

has the desired properties. ■

We can finally construct the formula interpreting the edges of H

$$\varphi(x, y) := \neg\text{sameSide}(x, y) \land \exists z_1, \ldots, z_r : E_0(x, z_1) \land E_0(y, z_r) \land \bigwedge_{t<r} E(z_t, z_{t+1}),$$

which states that x and y are from different sides and connected by a path containing $r + 1$ edges and whose first and last edge are E_0-edges. The formula defines a bipartite graph with sides A and B: it is symmetric, and we have $B_H \not\models \varphi(u, v)$ for all vertices u and v such that $\{u, v\} \subseteq A$ or $\{u, v\} \subseteq B$.

Claim 7.20. *For all $i \in [\![n]\!]$ and $j \in [\![m]\!]$: $B_H \models \varphi(a_i, b_j) \Leftrightarrow (i, j) \in E(H')$.*

Proof. Assume $(i, j) \in E(H')$. Then $p_{i,j,1}, \ldots, p_{i,j,r}$ is a valuation of $z_1, \ldots, z_r$ witnessing that $B_H \models \varphi(a_i, b_j)$. For the backwards direction assume a satisfying valuation $p_1, \ldots, p_r$ of $z_1, \ldots, z_r$. By Claim 7.19 we have $p_1 = p_{i,j',1}$ and $p_r = p_{i',j,r}$ for some $j' \in [\![m]\!] \setminus \{j\}$ and $i' \in [\![n]\!] \setminus \{i\}$. The existence of $p_2, \ldots, p_{r-1}$ implies $i = i'$ and $j = j'$. Hence, $(i, j) \in E(H')$. ■

It follows that $I_{\delta,\varphi}(B_H) = H$. Since the definition of δ and φ does not depend on H, and H was chosen to be an arbitrary bipartite graph without isolated vertices, we have that this interpretation interprets all such H from their corresponding preimage $B_H \in \mathcal{B}_r^{◂} \setminus (\mathcal{T} \cup \mathcal{I})$. Furthermore, for every H the preimage B_H can be computed in polynomial time, so the interpretation is efficient. This finishes the case of $\mathcal{B} = \mathcal{B}_r^{◂}$.

We next show that the interpretation that we constructed for the previous case, where $r = 1$ and $\mathcal{B} = \mathcal{B}_1^{◂}$, also works for the case $\mathcal{B} = \mathcal{B}^{▪}$. In the previous case, we constructed for every bipartite graph without isolated vertices H, a graph $B_H \in \mathcal{B}_1^{◂}$ such that $I_{\delta,\varphi}(B_H) = H$. In this case, since $r = 1$, the formula φ collapses to

$$\varphi(x, y) := \neg\text{sameSide}(x, y) \land \exists z_1 : E_0(x, z_1) \land E_0(y, z_1)$$

and the graph B_H consists of vertices

$$\underbrace{\{a_i : i \in [\![n]\!]\}}_{=A} \cup \underbrace{\{b_j : j \in [\![m]\!]\}}_{=B} \cup \underbrace{\{p_{i,j,1} : (i, j) \in E(H')\}}_{=P} \cup \{a_n, p_{n,1,1}\} \cup \{b_m, p_{1,m,1}\}$$

where H' is again the bipartite representation of H. We now build a new graph $B_H^{▪}$ from B_H by adding additional edges as follows. We connect each vertex $p_{i,j,1} \in P$ with all the vertices $p_{i',j',1} \in P$ such that $i \leqslant i'$ and $j \leqslant j'$ (but not with itself). It is easily checked that all the previous claims for B_H, still hold true for $B_H^{▪}$:

- We only modified adjacencies inside P, so Claim 7.14 and Observation 7.15 also hold in $B_H^{▪}$.
- The degree of vertices in P only increased, so Claim 7.16 and Observation 7.17 also hold in $B_H^{▪}$.
- As $r = 1$, P is pairwise distinguished using only A and B, so $B_H^{▪}$ contains no twins (Claim 7.18).

- The construction of the formulas $\text{sameSide}(x, y)$ and $E_0(x, y)$ only depends on the previous claims and the neighborhoods of A and B, so the formulas still work as intended in $B_H^{\blacksquare}$.

It follows that $B_H^{\blacksquare}$ contains neither isolated vertices nor twins and that

$$I_{\delta,\varphi}(B_H^{\blacksquare}) = I_{\delta,\varphi}(B_H) = H.$$

It remains to show that $B_H^{\blacksquare} \in \mathcal{B}^{\blacksquare}$. We do this by showing that $B_H^{\blacksquare}$ is an induced subgraph of the comparability grid G_N of order $N := (n + m - 1)$ whose vertex set is $\{a_{i,j} : i, j \in [N]\}$. We witness this fact by constructing an embedding f of $B_H^{\blacksquare}$ into G_N as follows:

- $f(a_i) := a_{m+i-1,n-i+1}$ for all $a_i \in A_\star$,
- $f(b_j) := a_{m-j+1,n+j-1}$ for all $b_j \in B_\star$,
- $f(p_{i,j,1}) := a_{m+i-1,n+j-1}$ for all $p_{i,j,1} \in P_\star$.

See the right side of Figure 7.3 for a visualization of the embedding. Using this visualization, it is easy to check that f is indeed an embedding of $B_H^{\blacksquare}$ into G_N. This finishes the case for $\mathcal{B} = \mathcal{B}^{\blacksquare}$ and concludes the proof. $\square$

7.5 Proof of Proposition 7.3

We can finally prove Proposition 7.3, which completes the proof of Theorem 2.3.

Proposition 7.3. *Let $\mathcal{C}$ be a hereditary graph class and $r \geqslant 1$, such that for all $k \in \mathbb{N}$, $\mathcal{C}$ contains*

- *a comparability grid of order k, or*
- *a flipped star r-crossing of order k, or*
- *a flipped clique r-crossing of order k, or*
- *a flipped half-graph r-crossing of order k.*

Then $\mathcal{C}$ efficiently interprets the class of all graphs.

Proof. We first show that $\mathcal{C}$ efficiently interprets the class of all bipartite graphs without isolated vertices. By Observation 7.5, we distinguish three cases.

- Either $\mathcal{C}$ contains a layer-wise flip of each graph from $\mathcal{B}$, for $\mathcal{B} \in \{\mathcal{B}_r^\star, \mathcal{B}_r^\bullet\}$, and we can apply Lemma 7.6 and Lemma 7.12,
- or $\mathcal{C}$ contains a layer-wise flip of each graph from $\mathcal{B}_r^{\blacktriangleleft}$, and we can apply Lemma 7.6 and Lemma 7.13,
- or $\mathcal{C}$ contains $\mathcal{B}^{\blacksquare}$ and we can apply Lemma 7.13.

From there, Lemma 7.11 brings us to the class of all graphs. $\square$

Part III

Monadic Stability

Outline Part III

In this part we prove several characterizations of monadically stable graph classes and establish tractable first-order model checking.

Theorem 2.1. *For every graph class $\mathcal{C}$, the following are equivalent.*

(1) $\mathcal{C}$ is monadically stable.

(2) $\mathcal{C}$ is flip-flat.

(3) For every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs

- *all flipped star r-crossings of order k, and*
- *all flipped clique r-crossings of order k, and*
- *all flipped half-graphs of order k.*

(4) For every $r \in \mathbb{N}$ there exists $\ell \in \mathbb{N}$ such that Flipper wins the radius-r budget-2 Flipper game in at most ℓ rounds on every graph from $\mathcal{C}$.

The first two characterizations (1) ⇔ (2) ⇔ (3) are shown in Chapters 8 and 9 and mostly follow easily from our characterizations for monadically dependent class from Part II. The Flipper game characterization (1) ⇔ (4) builds on (a strengthening of) flip-flatness and is presented in Chapter 10. The remainder of this part is devoted to model checking.

Theorem 2.2. *There is an algorithm that, given a graph G and a first-order sentence φ, decides whether $G \models \varphi$, and has the following property.*

For every monadically stable class $\mathcal{C}$, there exists a function $f : \mathbb{N} \times \mathbb{R} \to \mathbb{N}$ such that on any n-vertex graph $G \in \mathcal{C}$ and sentence φ the algorithm runs in time $f(|\varphi|, \varepsilon) \cdot n^{6+\varepsilon}$ for every $\varepsilon > 0$.

The model checking algorithm utilizes the game-tree of the Flipper game as a bounded depth decomposition of the r-neighborhoods of the input graph. We use locality of first-order logic to show that the r-neighborhoods preserve sufficient information to evaluate first-order sentences. As another important ingredient of the algorithm, we prove that monadically stable classes admit *sparse neighborhood covers* that can be used to cluster neighborhoods and thereby keep the size of the game-tree small. The existence and computability of suitable neighborhood covers is shown in Chapter 11 and the model checking algorithm is presented in Chapter 12.

Chapter 8

Flip-Flatness

In this chapter we characterize monadically stable graph classes by the combinatorial property dubbed *flip-flatness*, to obtain the equivalence (1) $\Leftrightarrow$ (2) of Theorem 2.1. We recall the definition of flip-flatness.

Definition 2.8 (Flip-Flatness)**.** A graph class $\mathcal{C}$ is *flip-flat* if for every radius $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $k_r \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N_r(m)$ there exist a subset $A \subseteq W$ with $|A| \geqslant m$ and a k_r-flip H of G such that for every two distinct vertices $u, v \in A$:

$$\text{dist}_H(u, v) > r.$$

Definition 8.1. A graph class $\mathcal{C}$ is *edge-stable*, if there is a bound k, such that $\mathcal{C}$ contains no half-graph of order k as a semi-induced subgraph.

The above condition is equivalent to the edge relation $E(x, y)$ being stable on $\mathcal{C}$.

Observation 8.2. Every monadically stable graph class is edge-stable.

Lemma 8.3. *Every monadically dependent, edge-stable graph class $\mathcal{C}$ is flip-flat.*

Proof. Fix a radius $r \in \mathbb{N}$. Let G be a graph from $\mathcal{C}$ and let $W \subseteq V(G)$. To prove the proposition we have to find a $\text{const}(\mathcal{C}, r)$-flip of G in which a size $U_{\mathcal{C},r}(|W|)$ subset A of W is distance-r independent.

By Proposition 6.47, $\mathcal{C}$ has the insulation property (cf. Definition 5.15). This means there is a subset $A \subseteq W$ of size $U_{\mathcal{C},r}(|W|)$ that is $(r+1, \text{const}(\mathcal{C}, r))$-insulated in G: there is an insulator $\mathcal{A}$ of height $r + 1$ and cost $\text{const}(\mathcal{C}, r)$ that insulates A, and each element of A is in a different cell of the bottom row of A.

Assume towards a contradiction that $\mathcal{A}$ is an ordered insulator. By the insulator property (O.5), G contains a semi-induced half-graph of order at least $|A|$. As A has unbounded size, this yields a contradiction to the stability of $\mathcal{C}$.

Then $\mathcal{A}$ must be an orderless insulator. By the insulator property (U.1), this means there is a $\text{const}(\mathcal{C}, r)$-flip of G in which the r-neighborhoods of the vertices in A are disjoint. In particular, A is distance-r independent, as desired. $\square$

Lemma 8.4. *Every flip-flat graph class is monadically stable.*

Proof. Assume towards a contradiction that there exists a class $\mathcal{C}$ that is not monadically stable but flip-flat. By definition of monadic stability there exists a formula $\sigma(x, y)$ defining arbitrarily large orders in a coloring of $\mathcal{C}$, that is, for every $n \in \mathbb{N}$ there exists a graph $G \in \mathcal{C}$ and a coloring G^+

such that we find a sequence $(a_1, \dots, a_n)$ in G^+ with $G^+ \models \sigma(a_i, a_j)$ if and only if $i < j$ for all $i, j \in [n]$.

Let q be the quantifier rank of σ. We set $r := 2 \cdot 7^q + 1$. Let s be the number of colors used by σ. Let N_r and k_r be the size function and flip-budget we obtain from $\mathcal{C}$ being flip-flat with radius r. As stated in Lemma 4.10, let $t_r := p(q, s \cdot k_r)$ be the number of colors needed to determine the truth value of formulas in the signature of $(s \cdot k_r)$-colored graphs that have the same quantifier-rank as $\sigma(x, y)$. Let $n := N_r(t_r + 1)$ and fix a graph $G \in \mathcal{C}$ such that in G^+ we find a sequence I of length n ordered by σ.

We apply flip-flatness to I and find a subsequence $J \subseteq I$ of length $t_r + 1$ together with a k_r-flip H of G in which J is r-independent in H. By using k_r colors to encode the flip, we can rewrite σ to a formula σ_r with the same quantifier-rank as σ, such that there exists a $(s \cdot k_r)$-coloring H^+ of H where for all $u, v \in V(G)$,

$$G^+ \models \sigma(u, v) \Leftrightarrow H^+ \models \sigma_r(u, v).$$

Hence, σ_r still orders J in H^+. As σ_r has the same quantifier-rank as σ and is a formula over the signature of $(s \cdot k_r)$-colored graphs, by Lemma 4.10 there exists a t_r-coloring of H^+ such that the truth of $\sigma_r(u, v)$ only depends on the colors of u and v for all $u, v \in J$. By the pigeonhole principle there exist two distinct vertices $u, v \in J$ that are assigned the same color. We therefore have $H^+ \models \sigma_r(u, v) \Leftrightarrow \sigma_r(v, u)$, which is a contradiction to σ_r ordering J in H^+. □

Combining Lemma 8.3 and Lemma 8.4 yields the following.

monadically dependent and edge stable $\Rightarrow$ flip-flat $\Rightarrow$ monadically stable

We close the circle of implications by the easy observation that every monadically stable class is monadically dependent and edge stable. We obtain the following proposition, which proves the equivalence (1) ⇔ (2) of Theorem 2.1. It also reproves the result from [64] that monadically dependent edge-stable graphs are monadically stable.

Proposition 8.5. *For every graph class $\mathcal{C}$ the following are equivalent.*

- *$\mathcal{C}$ is monadically stable.*
- *$\mathcal{C}$ is flip-flat.*
- *$\mathcal{C}$ is monadically dependent and edge-stable.*

As in the dependent case, we obtain the following result by Braunfeld and Laskowski [9] as a corollary of Theorem 2.3.

Corollary 8.6. *A hereditary graph class is monadically stable if and only if it is stable.*

Proof. Every monadically stable class is also stable. Let $\mathcal{C}$ be a hereditary graph class that is not monadically stable. By Proposition 8.5, it is either not monadically dependent or not edge-stable. If $\mathcal{C}$ is not monadically dependent then, by Corollary 7.2, $\mathcal{C}$ is not dependent so in particular not stable. If $\mathcal{C}$ is not edge-stable then the edge relation has the order property on $\mathcal{C}$, so again $\mathcal{C}$ is not stable. □

Since we can extract insulators efficiently (Proposition 5.23), the proof of Lemma 8.3 yields the following algorithmic version of flip-flatness as a corollary.

Corollary 8.7. *For every monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$, there exists an unbounded function $f_r : \mathbb{N} \to \mathbb{N}$, a constant $k_r \in \mathbb{N}$, and an algorithm that, given a graph $G \in \mathcal{C}$ and $W \subseteq V(G)$, computes in time $O_{\mathcal{C},r}(|V(G)|^2)$ a subset $A \subseteq W$ with $|A| \geqslant f_r(|W|)$ and a k_r-flip H of G such that for all distinct $u, v \in A$:*

$$\mathrm{dist}_H(u, v) > r.$$

Chapter 9

Forbidden Induced Subgraphs

In this chapter we characterize monadically stable graph classes by forbidden induced subgraphs. The following proposition proves the equivalence (1) $\Leftrightarrow$ (3) of Theorem 2.1. The definitions of the crossings can be found in Section 2.2.

Proposition 9.1. *A graph class $\mathcal{C}$ is monadically stable if and only if for every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs*

- *all flipped star r-crossings of order k, and*
- *all flipped clique r-crossings of order k, and*
- *all flipped half-graphs of order k.*

Proof. For the forward direction, we assume that $\mathcal{C}$ is monadically stable and conclude as follows.

1. As $\mathcal{C}$ is monadically dependent, the characterization of monadically dependent classes via excluded induced subgraphs (Proposition 6.47) applies and $\mathcal{C}$ excludes the flipped star and clique crossings.
2. As $\mathcal{C}$ is edge-stable, it excludes all flipped half-graphs of sufficiently large order.

For the backwards direction, assume $\mathcal{C}$ excludes the induced subgraphs as stated. Now notice the following.

1. Every comparability grid of order k contains a flipped half-graph of order k as a subgraph induced by any two rows of its rows.
2. For every $r \geqslant 1$, every flipped half-graph crossing of order k contains a flipped half-graph of order k as a subgraph induced between the roots $\{a_i : i \in [k]\}$ and $\{\text{start}(\pi_{i,k}) : i \in [k]\}$.

Hence, $\mathcal{C}$ excludes also large comparability grids and flipped half-graph crossings. Then by Proposition 6.47, $\mathcal{C}$ is monadically dependent. As $\mathcal{C}$ excludes flipped half-graphs, it is also edge-stable and Proposition 8.5 yields that $\mathcal{C}$ is monadically stable, as desired. □

Chapter 10

Flipper Game

In this chapter we characterize monadically stable graph classes via the Flipper game. This corresponds to the last remaining equivalence (1) ⇔ (4) of Theorem 2.1. We formally define the rules of the game.

Definition 10.1 (Flipper game)**.** Fix $r, k \in \mathbb{N}$. The *radius-r budget-k Flipper game* is played by two players, *Flipper* and *Localizer*, on a graph G as follows. At the beginning, set $G_0 := G$. In the ith round, for $i > 0$, the game proceeds as follows.

- If $|G_{i-1}| = 1$, then Flipper wins.
- Localizer chooses G^{loc}_{i-1} as a subgraph of G_{i-1} induced by a (non-empty) subset of an r-neighborhood in G_{i-1}.
- Flipper chooses G_i as a k-flip of G^{loc}_{i-1}.

The goal is to prove that in every monadically stable graph classes $\mathcal{C}$, Flipper has a strategy to win the game in a constant number of rounds, independent of the size of the graph: for every radius r there exists a constant $\ell \leqslant \mathrm{const}(\mathcal{C}, r)$ such that Flipper wins the radius-r budget-2 Flipper game in at most ℓ rounds on every graph $G \in \mathcal{C}$. Moreover, we want Flippers winning moves to be computable in polynomial time. We introduce the necessary definitions to make a formal statement.

10.1 Preliminaries

Strategies are commonly represented by functions mapping the history of the game to a new (played) position. In our context, it will be convenient to use the following equivalent abstraction, which will fit better to our algorithmic perspective. Fix radius $r \in \mathbb{N}$. Graphs considered in consecutive rounds of the Flipper game will often be called *arenas*, for brevity. A radius-r *Localizer strategy* is a function

$$\mathsf{loc}\colon (G_i) \mapsto (G^{\mathrm{loc}}_i)$$

that maps the arena G_i at round i to Localizer's next move: a graph G^{loc}_i that is an induced subgraph of the r-ball around some vertex v in G_i.

A budget-k *Flipper strategy* flip is an algorithm that computes a function

$$\mathsf{flip}\colon (G^{\mathrm{loc}}_i, \mathcal{I}_i) \mapsto (G_{i+1}, \mathcal{I}_{i+1})$$

that maps the graph G^{loc}_i obtained from Localizer's move to Flipper's response: a k-flip of G^{loc}_i. Additionally, we allow Flipper to keep an auxiliary memory: the strategy takes, as the second argument, an *internal state* $\mathcal{I}_i$ from the previous round, and outputs an updated internal state $\mathcal{I}_{i+1}$.

The initial state $\mathcal{I}_0 := G$ consists just of the initial graph at the beginning of the game. The internal states will be used as memory and to precompute flips for future turns, which makes them convenient from an algorithmic point of view. Strategies operating with game histories instead of internal states can simulate the latter in the following sense: knowing the game history, Flipper can compute the current internal state by replaying the entire game up to the current round. Note that since we are interested in Flipper strategies that work against *any* behavior of Localizer, it is not necessary to equip Localizer strategies with memory as well.

Let loc and flip be Localizer and Flipper strategies, and let G be a graph. We define the *run* $\mathcal{R}(\mathsf{loc}, \mathsf{flip}, G)$ to be the infinite sequence of *positions*

$$\mathcal{R}(\mathsf{loc}, \mathsf{flip}, G) := (G_0, \mathcal{I}_0), (G_1, \mathcal{I}_1), (G_2, \mathcal{I}_2), (G_3, \mathcal{I}_3), \dots$$

such that $G_0 = \mathcal{I}_0 = G$, and for all $i \geqslant 0$ we have $(G_{i+1}, \mathcal{I}_{i+1}) = \mathsf{flip}(\mathsf{loc}(G_i), \mathcal{I}_i)$.

A *winning position* is a tuple $(G_i, \mathcal{I}_i)$ where G_i contains only a single vertex. A Flipper strategy flip is *ℓ-winning in the radius-r game* on a graph class $\mathcal{C}$, if for every $G \in \mathcal{C}$ and for every radius-r Localizer strategy loc, the ℓth position of $\mathcal{R}(\mathsf{loc}, \mathsf{flip}, G)$ is a winning position. Note that while $\mathcal{R}(\mathsf{loc}, \mathsf{flip}, G)$ is an infinite sequence, once a winning position is reached, it is only followed by winning positions.

Definition 10.2 (Runtime of a Flipper strategy)**.** Let flip be a Flipper strategy. For a function $f : \mathbb{N} \to \mathbb{N}$, a radius $r \in \mathbb{N}$, and a graph G we say that

flip has *runtime f in the radius-r game on G*

if for every radius-r Localizer strategy loc and $i \in \mathbb{N}$

the computation of $\mathsf{flip}(\mathsf{loc}(G_i), \mathcal{I}_i)$ takes time at most $f(|V(G)|)$

where $(G_i, \mathcal{I}_i)$ is the ith position of the run $\mathcal{R}(\mathsf{loc}, \mathsf{flip}, G)$. For a graph class $\mathcal{C}$, flip has *runtime f in the radius-r game on $\mathcal{C}$* if it has runtime f in the radius-r game on G for every $G \in \mathcal{C}$.

We remark that the time complexity is allowed to depend on the original graph G, which is possibly much larger than the current arena G_i.

The following definition will be useful to reproducibly pick representatives from vertex sets.

Definition 10.3. A *well-ordered* graph is a graph equipped with a well-order of its vertex set. Given a well-ordered graph G, for every subset $A \subseteq V(G)$, we denote by $\min_G A$ the smallest element of A according to the well-order on $V(G)$. For a graph class $\mathcal{C}$, we write $\sigma(\mathcal{C})$ for the class of all well-ordered graphs, whose underlying (unordered) graph is from $\mathcal{C}$.

10.2 Flipper's Strategy

Having established the necessary definitions, we can now state the main result of this chapter.

Theorem 10.4. *There is a budget-2 Flipper strategy* $\mathsf{flip}^\star$ *with the following property. For every monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$ there is $\ell \in \mathbb{N}$ such that* $\mathsf{flip}^\star$ *is ℓ-winning and has runtime $O_{\mathcal{C},r}(n^2)$ in the radius-r game on $\mathcal{C}$.*

The proof of Theorem 10.4 will rely on a strengthening of the flip-flatness characterization, which we call *predictable flip-flatness*. We first state a simplified version.

Proposition 10.5 (Predictable flip-flatness, simplified). *For every monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$, there is a bound $k_{\mathcal{C},r} \in \mathbb{N}$ and functions* $\mathrm{Flip}_{\mathcal{C},r}$, $\mathrm{Flat}_{\mathcal{C},r}$ *and* $\mathrm{Predict}_{\mathcal{C},r}$, *such that for all well-ordered graphs $G \in \sigma(\mathcal{C})$ and sets $X, Z \subseteq V(G)$ we have*

- $\mathrm{Flip}_{\mathcal{C},r}(G, X)$ *is a $k_{\mathcal{C},r}$-flip of G,*
- $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *is a size $U_{\mathcal{C},r}(|X|)$ subset of X,*
- $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *is distance-r independent in* $\mathrm{Flip}_{\mathcal{C},r}(G, X)$,
- $\mathrm{Predict}_{\mathcal{C},r}(G, Z)$ *is a $k_{\mathcal{C},r}$-flip of G, and*
- *if Z is a size 5 subset of* $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *then* $\mathrm{Predict}_{\mathcal{C},r}(G, Z) = \mathrm{Flip}_{\mathcal{C},r}(G, X)$.

The properties of the functions $\mathrm{Flip}_{\mathcal{C},r}$ and $\mathrm{Flat}_{\mathcal{C},r}$ mirror the guarantees given by the flip-flatness property of monadically stable graph classes that we have established in Chapter 8. The new ingredient is the function $\mathrm{Predict}_{\mathcal{C},r}$. This function allows to "predict" the flip in which the set $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ is r-independent, just by knowing a size 5 subset Z of $\mathrm{Flat}_{\mathcal{C},r}(G, X)$. It is instructive to observe the following corollary.

Corollary 10.6. *For every monadically stable class $\mathcal{C}$, $r \in \mathbb{N}$, $G \in \sigma(\mathcal{C})$, and $X_1, X_2 \subseteq V(G)$:*

$$|\mathrm{Flat}_{\mathcal{C},r}(G, X_1) \cap \mathrm{Flat}_{\mathcal{C},r}(G, X_2)| \geqslant 5 \quad \Rightarrow \quad \mathrm{Flip}_{\mathcal{C},r}(G, X_1) = \mathrm{Flip}_{\mathcal{C},r}(G, X_2).$$

Proof. We have $\mathrm{Predict}_{\mathcal{C},r}(G, Z) = \mathrm{Flip}_{\mathcal{C},r}(G, X_1) = \mathrm{Flip}_{\mathcal{C},r}(G, X_2)$, where Z is any size 5 subset of $\mathrm{Flat}_{\mathcal{C},r}(G, X_1) \cap \mathrm{Flat}_{\mathcal{C},r}(G, X_2)$. □

Intuitively, the flip $\mathrm{Predict}_{\mathcal{C},r}$ can be used to "flatten" many different sets. We will use this property to win the Flipper game: Flipper will play flips predicted by $\mathrm{Predict}_{\mathcal{C},r}$ that will help to win the game regardless of how Localizer reacts to them.

As we have mentioned earlier, Proposition 10.5 presents predictable flip-flatness in a slightly simplified form. To construct a Flipper strategy with an efficient runtime, we additionally want to be able to efficiently compute the values of $\mathrm{Predict}_{\mathcal{C},r}$. While the proof of $\mathrm{Predict}_{\mathcal{C},r}$ is effective, the following issue arises during the computation. $\mathrm{Predict}_{\mathcal{C},r}$ needs to return a flip of small budget (i.e. a $k_{\mathcal{C},r}$-flip) for all inputs Z, even for those sets Z which are not subsets of some flat set X. In order to always return $k_{\mathcal{C},r}$-flips, the algorithm for $\mathrm{Predict}_{\mathcal{C},r}$ needs to know the bound $k_{\mathcal{C},r}$, so that it can detect when a computed flip exceeds the budget, and instead return a trivial flip of small budget. (Exceeding the budget only happens if Z is not a subset of some flat set X; a case where returning a trivial flip is ok.) As for some monadically stable classes, the bound $k_{\mathcal{C},r}$ is not computable, it has to be hardwired into the algorithm, creating a different algorithm for each class $\mathcal{C}$ and radius r. This is problematic as we strive to obtain a single Flipper strategy that works for all classes and radii, so that we can ultimately derive a single model checking algorithm which runs fast on every monadically stable graph class $\mathcal{C}$ and for every formula φ (whose quantifier rank determines the radius of the Flipper game).

To sidestep this issue and get a single uniform algorithm, we give an additional input $k \in \mathbb{N}$ to the predictable flip-flatness algorithm that estimates $k_{\mathcal{C},r}$. With k as input, it is easy for the algorithm to only produce k-flips. We then only demand the prediction to be accurate if a sufficiently large $k \geqslant k_{\mathcal{C},r}$ was given as input. Even though we do not know the appropriate value of $k_{\mathcal{C},r}$, a parameterizable algorithm will enable us to "guess" the value by dovetailing. We next state the algorithmic predictable flip-flatness statement.

Proposition 10.7 (Predictable flip-flatness)**.** *There is an algorithm that takes as input $r, k \in \mathbb{N}$, a well-ordered graph G, and a size five set $Z \subseteq V(G)$, and computes in time $O_{r,k}(|V(G)|^2)$ a k-flip* $\mathrm{Predict}(r, k, G, Z)$ *of G with the following properties:*

For every monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$ there is a bound $k_{\mathcal{C},r} \leqslant \mathrm{const}(\mathcal{C}, r)$ and functions $\mathrm{Flip}_{\mathcal{C},r}$ and $\mathrm{Flat}_{\mathcal{C},r}$ such that for all well-ordered graphs $G \in \mathcal{C}$, sets $X, Z \subseteq V(G)$ and integers $k \geqslant k_{\mathcal{C},r}$ we have

- $\mathrm{Flip}_{\mathcal{C},r}(G, X)$ *is a $k_{\mathcal{C},r}$-flip of G,*
- $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *is a size $U_{\mathcal{C},r}(|X|)$ subset of X,*
- $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *is distance-r independent in* $\mathrm{Flip}_{\mathcal{C},r}(G, X)$*, and*
- *if Z is a size 5 subset of* $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *then* $\mathrm{Predict}(r, k, G, Z) = \mathrm{Flip}_{\mathcal{C},r}(G, X)$.

The proof of Proposition 10.7 is deferred to Section 10.4, and we continue with our proof of Theorem 10.4. We first present a Flipper strategy, where Flipper is allowed to use an arbitrary but finite budget k in every round and later improve the budget to 2 a posteriori.

Lemma 10.8. *There is a parameterizable algorithm* $\mathsf{flip}[\cdot, \cdot]$ *with the following properties.*

1. *For all parameters $r \in \mathbb{N}$ and $k \geqslant 3$:*
   - $\mathsf{flip}[r, k]$ *is a budget-k Flipper strategy, and*
   - $\mathsf{flip}[r, k]$ *has runtime $O_{r,k}(n^2)$ in the radius-r game on the class of all graphs.*
2. *For each monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$, there is a constant $k^*_{\mathcal{C},r} \in \mathbb{N}$, such that for each $k \in \mathbb{N}$ with $k \geqslant k^*_{\mathcal{C},r}$:*
   - $\mathsf{flip}[r, k]$ *is* $\mathrm{const}(\mathcal{C}, r)$*-winning in the radius-$r$ game on $\mathcal{C}$.*

Again the algorithm is parameterized to later allow for a uniform algorithm that works on all classes and radii. The non-uniform variant is simpler to state and follows as a corollary.

Corollary 10.9. *For every monadically stable graph class $\mathcal{C}$ and $r \in \mathbb{N}$, there are $\ell, k \in \mathbb{N}$ and an ℓ-winning budget-k Flipper strategy with runtime $O_{\mathcal{C},r}(n^2)$ in the radius-r game on $\mathcal{C}$.*

Proof. The strategy $\mathsf{flip}[r, k^*_{\mathcal{C},r}]$ has the desired properties. ☐

We continue with the proof of Lemma 10.8

Proof of Lemma 10.8. We first explain the strategy $\mathsf{flip}[\cdot, \cdot]$ in natural language and prove its properties afterwards. The easy translation to the formal layer of strategies with internal states, is left to the reader.

Description of the Strategy. Fix any $r, k \in \mathbb{N}$. The radius-r Flipper game is played on an n-vertex graph G on which we fix an arbitrary well-order.

At the micro level, Flipper will always play his moves in *move pairs*. A move pair is *defined by* a k-flip H of G as follows. We assume the current arena G_i is an induced subgraph of G.

1. Localizer localizes yielding an induced subgraph G_i^{loc} of G_i (and also of G).
2. Flipper plays the k-flip $G_{i+1} := H[V(G_i^{\mathrm{loc}})]$ of G_i^{loc}.
3. Localizer localizes yielding an induced subgraph G_{i+1}^{loc} of G_{i+1} (and also of H).
4. Flipper reverts his flip setting $G_{i+2} := G_i^{\mathrm{loc}}[V(G_{i+1}^{\mathrm{loc}})] = G[V(G_{i+1}^{\mathrm{loc}})]$.

It is easy to see that, assuming we start with an arena G_i that is an induced subgraph of G, Flippers moves are valid k-flips and we end up with an induced subgraph of G again.

At the macro level, Flipper proceeds in a sequence of *eras*, each consisting of multiple move pairs. Along the way, he keeps track of a growing chain of vertex subsets

$$X_0 \subsetneq X_1 \subsetneq X_2 \subsetneq X_3 \subsetneq \ldots,$$

where $X_0 = \emptyset$ and X_i is obtained from X_{i-1} by adding one vertex that was removed from the arena during era i. Up until Flipper wins the game, we will ensure that such a vertex always exists, and therefore $|X_i| = i$ for every $i \in \mathbb{N}$ until the game concludes.

We now describe Flipper's moves in era i ($i = 1, 2, 3, \ldots$). For every $Z \subseteq X_{i-1}$ with $|Z| = 5$, we use the algorithm from Proposition 10.7 to compute the k-flip $H_Z := \text{Predict}(2r, k, G, Z)$ of G. Note that, instead of the current arena, the original graph G is used to compute H_Z. First, Flipper performs $\binom{|X_{i-1}|}{5}$ move pairs, each defined by H_Z for a different Z as above. To conclude the era, Flipper picks an arbitrary vertex x that is still in the arena. Let H_x be the subgraph of G obtained by removing all edges incident to x. This graph is a 3-flip of G and the reason we demand $k \geqslant 3$. Flipper plays the move pair defined by H_x. Let B be the (subset of a) r-ball played by Localizer after Flipper played (an induced subgraph of) H_x during the concluding move pair defined by H_x. If B contains the vertex x, then, since x is isolated in H_x, we have $B = \{x\}$ and Localizer looses at the beginning of the next round. Otherwise, we have $x \notin B$ and x is removed from the arena during era i. We set $X_i := X_{i-1} \cup \{x\}$ and proceed to the next era. This concludes the description of $\mathsf{flip}[r, k]$.

Budget Bound. Fix any $r, k \in \mathbb{N}$ and graph G. Flipper only plays move pairs defined by k-flips of G. Moreover, after every move pair, the current arena is an induced subgraph of G. It follows by induction that each subsequent arena is an induced subgraph of a k-flip of the previous arena (and of G). Therefore, $\mathsf{flip}[r, k]$ is a valid budget-k Flipper strategy.

Runtime Bound. Fix any $r, k \in \mathbb{N}$ and graph G. In era i, $\mathsf{flip}[r, k]$ plays $\binom{|X_{i-1}|}{5}$ many move pairs, each of which is computed by the algorithm of Proposition 10.7 in time $O_{r,k}(n^2)$. Note that Flipper does not need to compute all move pairs at once. Instead, Flipper can compute a single move pair, execute it during the next rounds, and continue computing the next move pair afterwards. This way, with negligible bookkeeping overhead, Flipper needs to call the algorithm of Proposition 10.7 at most once per round. Since we have imposed no restrictions on G, this means $\mathsf{flip}[r, k]$ has runtime $O_{r,k}(n^2)$ on the class of all graphs.

Duration Bound. Fix any monadically stable class $\mathcal{C}$ and radius $r \in \mathbb{N}$. We choose $k^*_{\mathcal{C},r} := k_{\mathcal{C},2r}$ to be the bound obtained for the class $\mathcal{C}$ and radius $2r$ from Proposition 10.7. Fix any $k \geqslant k^*_{\mathcal{C},r}$. We want to show that it is ℓ-winning in the radius-r game on $\mathcal{C}$ for some $\ell \leqslant \text{const}(\mathcal{C}, r)$ which we will define later.

Let $\text{Flip}_{\mathcal{C},2r}$ and $\text{Flat}_{\mathcal{C},2r}$ be the functions provided by Proposition 10.7 for the class $\mathcal{C}$. Let t be the least integer such that $|\text{Flat}_{\mathcal{C},2r}(G, X)| \geqslant 7$ for every $G \in \sigma(\mathcal{C})$ and every set $X \subseteq V(G)$ of size at least t. This value depends on $\mathcal{C}$ and r.

Claim 10.10. *If Flipper follows* $\mathsf{flip}[r, k]$*, the game concludes within at most* t *eras.*

Proof. For contradiction, suppose the game enters era $t+1$ without termination. Denote $X := X_t$ with $|X| = t$ and let $Y := \text{Flat}_{\mathcal{C},2r}(G, X)$. By our choice of t and by Proposition 10.7, we have $|Y| \geqslant 7$ and Y is distance-$2r$ independent in $\text{Flip}_{\mathcal{C},2r}(G, X)$. Let $y_1, \ldots, y_7$ be any seven distinct vertices of Y, where y_i was added earlier to X than y_j for all $i < j$. Let $Z := \{y_1, \ldots, y_5\}$ and let s be the index of the era that concluded with adding y_6 to X. (That is, we have $X_s = X_{s-1} \cup \{y_6\}$ and in particular $X_{s-1} \cap \{y_6, y_7\} = \emptyset$.) Note that $Z \subseteq X_{s-1}$, hence within era s, Flipper played the move pair defined by Predict$(2r, k, G, Z)$. This means that during the execution of that move pair, Localizer restricted the arena to (an induced subgraph of) some r-ball $B_\star$ in Predict$(2r, k, G, Z)$. As $k \geqslant k^*_{\mathcal{C},r}$, by Proposition 10.7, we have Predict$(2r, k, G, Z) = \text{Flip}_{\mathcal{C},2r}(G, X)$, which means that y_6 and y_7 have distance greater than $2r$ in Predict$(2r, k, G, Z)$. It follows that $B_\star$ contains at most one of y_6 and y_7. This contradicts the assumption that both y_6 and y_7 were later added to X, which requires them to both be contained in the arena at the end of era s. ■

Note that in era i, Flipper applies exactly $\binom{|X_{i-1}|}{5} + 1 = \binom{i-1}{5} + 1$ move pairs. Hence, by Claim 10.10, the game terminates within at most

$$\ell := \sum_{i=1}^{t} 2 \cdot \left(\binom{i-1}{5} + 1\right) \leqslant \text{const}(\mathcal{C}, r) \quad \text{rounds.}$$

We conclude that $\mathsf{flip}[r, k]$ is ℓ-winning in the radius-r game on $\mathcal{C}$, as promised. This concludes the proof of the lemma. □

Next we show how to simulate a k-flip by a series of 2-flips.

Lemma 10.11. *Given a graph $G = G_0$ and a k-flip H of G, we can compute in time $O_k(|V(G)|^2)$ a sequence of graphs $G_1, \ldots, G_t$ with $G_t = H$, where $t \leqslant 3k^2$ and G_i is a 2-flip of G_{i-1} for $i \in [t]$.*

Proof. We compute a partition $\mathcal{P} \subseteq 2^{V(G)}$ and symmetric relation $F \in \mathcal{P}^2$ witnessing that H is a k-flip of G in time $O_k(|V(G)|^2)$ using Lemma 4.1.

We simulate flipping a single pair of parts $(P, Q) \in F$ in G by the sequence of 2-flips

$$G_1 := G \oplus \{(P \cup Q, P \cup Q)\}, \quad G_2 := G_1 \oplus \{(P, P)\}, \quad G_3 := G_2 \oplus \{(Q, Q)\}.$$

Chaining this procedure for all the at most k^2 pairs in F proves the lemma. □

Lemma 10.12. *There is a budget-2 Flipper strategy $\mathsf{flip}_2[\cdot, \cdot]$ with the following properties.*

1. *For all parameters $r \in \mathbb{N}$ and $k \in \mathbb{N}$:*
   - *$\mathsf{flip}_2[r, k]$ has runtime $O_{r,k}(n^2)$ in the radius-r game on the class of all graphs.*
2. *For each monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$, there is a constant $k^*_{\mathcal{C},r} \in \mathbb{N}$, such that for each $k \in \mathbb{N}$ with $k \geqslant k^*_{\mathcal{C},r}$:*
   - *$\mathsf{flip}_2[r, k]$ is const$(\mathcal{C}, r, k)$-winning in the radius-r game on $\mathcal{C}$.*

Proof. We rename the input parameter k to k' to free up the name and set $k := \max(k', 3)$. We describe $\mathsf{flip}_2[r, k']$ by showing how to wrap the budget-k Flipper strategy $\mathsf{flip}[r, k]$ into a budget-2 strategy, while maintaining its good properties. In round i, using $\mathsf{flip}[r, k]$, we calculate a k-flip G_i of G^{loc}_{i-1}. We simulate playing G_i by playing (at most) $3k^2$ many rounds of 2-flips as follows. Applying Lemma 10.11 to G^{loc}_{i-1} and G_i yields a sequence

$$H_i, H_{i+1}, \ldots, H_t$$

of $(t - i + 1) \leqslant 3k^2$ many 2-flips where $H_t = G_i$. In round j (with $i \leqslant j \leqslant t$), Flipper plays $G_j := H_j[V(G^{\text{loc}}_{j-1})]$. Using the guarantees of Lemma 10.11, it is easy to check that the described moves form a valid play for Flipper in the budget-2 game. Moreover, Localizers move G^{loc}_t is an induced subgraph of an r-ball in G_i. It follows by induction that, if $\mathsf{flip}[r, k]$ is ℓ-winning in the radius-r game on a class $\mathcal{C}$, then $\mathsf{flip}_2[r, k]$ is $(3k^2\ell)$-winning on the same class. In particular for every monadically stable class $\mathcal{C}$ and $k' \geqslant k^*_{\mathcal{C},r}$, we have also $k \geqslant k^*_{\mathcal{C},r}$ so $\mathsf{flip}_2[r, k]$ is $(3k^2 \cdot \text{const}(\mathcal{C}, r))$-winning on $\mathcal{C}$. Here $k^*_{\mathcal{C},r}$ has the same value as in Lemma 10.8. □

Again we obtain the following corollary by considering the strategy $\mathsf{flip}_2[r, k^*_{\mathcal{C},r}]$.

Corollary 10.13. *For every monadically stable graph class $\mathcal{C}$ and $r \in \mathbb{N}$, there is $\ell \in \mathbb{N}$ and an ℓ-winning budget-2 Flipper strategy with runtime $O_{\mathcal{C},r}(n^2)$ in the radius-r game on $\mathcal{C}$.*

The corollary already proves the implication (1) ⇒ (4) of Theorem 2.1, but it gives a different strategy for each class $\mathcal{C}$ and radius r, as we do not know the value of $k^*_{\mathcal{C},r}$. We finally use dovetailing to guess $k^*_{\mathcal{C},r}$ and obtain a single strategy that works for every monadically stable class and every radius.

Theorem 10.4. *There is a budget-2 Flipper strategy* $\mathsf{flip}^\star$ *with the following property. For every monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$ there is $\ell \in \mathbb{N}$ such that* $\mathsf{flip}^\star$ *is ℓ-winning and has runtime $O_{\mathcal{C},r}(n^2)$ in the radius-r game on $\mathcal{C}$.*

Proof. We fix an enumeration N of $\mathbb{N}^3_+$ by joining enumerations of the sets

$$[1]^3, [2]^3 \setminus [1]^3, [3]^3 \setminus [2]^3, [4]^3 \setminus [3]^3, \ldots$$

Crucially, every triple $(r, k, \ell) \in \mathbb{N}^3_+$ appears among the first $\max(r, k, \ell)^3$ elements of N. The Flipper strategy $\mathsf{flip}^\star$ again proceeds in *eras*. Let G_i be the current arena at the start of era i. We will maintain the invariant that G_i is an induced subgraph of G, which holds for $G_1 = G$. Let $(r, k, \ell) \in \mathbb{N}^3_+$ be the tuple at the ith position of the enumeration N. In era i, Flipper plays 2ℓ many rounds. Flipper first plays ℓ many rounds according to the strategy $\mathsf{flip}_2[r, k]$ given by Lemma 10.12, starting on the graph G_i. After those ℓ rounds, Flipper plays the same ℓ flips again in reverse order (and restricted to the current arena). This reverses the flips: after the 2ℓ rounds the arena is again an induced subgraph of G_i and therefore also of G. Flipper continues with the era $i + 1$. This finishes the description of $\mathsf{flip}^\star$.

As Flipper only plays flips generated by the budget-2 strategy $\mathsf{flip}_2[\cdot, \cdot]$, it is clear that $\mathsf{flip}^\star$ is itself a budget-2 strategy.

Now fix a monadically stable class $\mathcal{C}$ and a radius $r \in \mathbb{N}$. We first argue that there is a bound $\ell \leqslant \mathrm{const}(\mathcal{C}, r)$ such that $\mathsf{flip}^\star$ is ℓ-winning in the radius-r game on $\mathcal{C}$. By Lemma 10.12, there is a bound $k^*_{\mathcal{C},r} \leqslant \mathrm{const}(\mathcal{C}, r)$ such that $\mathsf{flip}_2[r, k^*_{\mathcal{C},r}]$ is ℓ^*-winning in the radius-r game on $\mathcal{C}$ for some $\ell^* \leqslant \mathrm{const}(\mathcal{C}, r)$. Let i be the position of the triple $(r, k^*_{\mathcal{C},r}, \ell^*)$ in N. By construction of N, we have $i \leqslant \max(r, k^*_{\mathcal{C},r}, \ell^*)^3 \leqslant \mathrm{const}(\mathcal{C}, r)$. Let $G \in \mathcal{C}$ be the graph on which the game is played. In era i, Flipper played according to $\mathsf{flip}_2[r, k^*_{\mathcal{C},r}]$ for ℓ^* rounds on an induced subgraph G_i of G. Without loss of generality we can assume $\mathcal{C}$ is hereditary and $G_i \in \mathcal{C}$. Since $\mathsf{flip}_2[r, k^*_{\mathcal{C},r}]$ is ℓ^*-winning in the radius-r game on $\mathcal{C}$, this means Flipper must have reached a winning position at the end of era i. Again by construction of N, each era j lasts at most $2j$ many rounds. It follows that $\mathsf{flip}^\star$ wins the radius-2 Flipper game on G in at most

$$\ell \quad := \quad \sum_{j \in [i]} 2j \quad \leqslant \quad 2i^2 \quad \leqslant \quad \mathrm{const}(\mathcal{C}, r)$$

many rounds. As $G \in \mathcal{C}$ was chosen arbitrary, we have shown that $\mathsf{flip}^\star$ is ℓ-winning in the radius-r game on $\mathcal{C}$, as desired.

For the running time, note that $\mathsf{flip}^\star$ wins in era $i \leqslant \mathrm{const}(\mathcal{C}, r)$. This means in each round only a single query to $\mathsf{flip}_2[r, k]$ for values r, k both of size at most i is made. By Lemma 10.12, each of these calls runs in time $O_i(n^2)$. This is bounded by $O_{\mathcal{C},r}(n^2)$ as desired. ☐

10.3 From Flipper Game To Monadic Stability

We complete the equivalence (1) ⇔ (4) of Theorem 2.1, by proving the remaining implication ¬(1) ⇒ ¬(4). This finishes the proof of Theorem 2.1.

Proposition 10.14. *Let $\mathcal{C}$ be a monadically unstable class. Then there exists an $r \in \mathbb{N}$ such that for every $k \in \mathbb{N}$ and $\ell \in \mathbb{N}$, there exists a graph $G \in \mathcal{C}$ in which Localizer can survive the radius-r budget-k Flipper game for at least ℓ rounds.*

To prove this proposition we will utilize the characterization of monadically stable graph classes by forbidden induced subgraphs, which we recall here.

Proposition 9.1. *A graph class $\mathcal{C}$ is monadically stable if and only if for every $r \geqslant 1$ there exists $k \in \mathbb{N}$ such that $\mathcal{C}$ excludes as induced subgraphs*

- *all flipped star r-crossings of order k, and*
- *all flipped clique r-crossings of order k, and*
- *all flipped half-graphs of order k.*

The first crucial observation is that, for each of the unstable patterns, a large part of the pattern survives in every k-flip.

Lemma 10.15. *Fix $r, k \in \mathbb{N}$ and $\mathbb{G} \in \{$star r-crossing, clique r-crossing, halfgraph$\}$. Let G contain as an induced subgraph a flipped $\mathbb{G}$ of order n. Every k-flip H of G contains as an induced subgraph a flipped $\mathbb{G}$ of order $U_{r,k}(n)$.*

Proof. Let $\mathcal{P}$ be the size k partition of $V(G)$ witnessing that H is a k-flip of G. By Bipartite Ramsey (Lemma 4.15), G contains an induced subgraph G' isomorphic to a flipped $\mathbb{G}$ of order $U_{r,k}(n)$ where the partition $\mathcal{P}$ respects the layering of $\mathbb{G}$: whenever two vertices of G' are in the same layer, then they are in the same part of $\mathcal{P}$. It follows that $H[V(G')]$ is a layer-wise flip of a flipped $\mathbb{G}$. By definition this is again a flipped $\mathbb{G}$. □

The second crucial observation is that, in each of the unstable pattern, there exists a neighborhood which contains a large part of the pattern.

Lemma 10.16. *Fix $r \in \mathbb{N}$ and $\mathbb{G} \in \{$star r-crossing, clique r-crossing, halfgraph$\}$. Let G contain a flipped $\mathbb{G}$ of order n as an induced subgraph. There exists a vertex $v \in V(G)$ such that $G[N_{2r+2}[v]]$ contains a flipped $\mathbb{G}$ of order $U(n)$ as an induced subgraph.*

Proof. In the case of flipped star and clique crossings, any vertex of the flipped crossing can take the role of v. In the case of flipped half-graphs, there might be up to two isolated vertices which are part of the flipped half-graph. Here every non-isolated vertex which is part of the flipped half-graph does the job. □

We can now prove Proposition 10.14.

Proof of Proposition 10.14. By Proposition 9.1, there exists $r' \in \mathbb{N}$ such that $\mathcal{C}$ contains *unstable patterns* (flipped star r'-crossings, flipped clique r'-crossings, or flipped half-graphs) of arbitrary large order as induced subgraphs. We choose $r := 2r' + 2$ and show that for every $\ell, k \in \mathbb{N}$, there exists $G \in \mathcal{C}$ on which Flipper cannot win the radius-r budget-k Flipper game in ℓ rounds. Fix $\ell, k \in \mathbb{N}$. Lemma 10.15 and Lemma 10.16 yields functions $f_1, f_2 : \mathbb{N} \to \mathbb{N}$, such that for every $m \in \mathbb{N}$ in every arena that contains an unstable pattern of order at least $f(m) := \max(f_1(m), f_2(m))$ the following holds.

- For each possible Flipper move, an unstable pattern of size m remains in the arena.
- There exists a Localizer move that keeps an unstable pattern of size m in the arena.

Let $G \in \mathcal{C}$ be a graph which contains an unstable pattern of order $f^{2\ell}(2)$. An easy induction shows that Localizer has a strategy such that, no matter how Flipper plays, after ℓ rounds the arena will still contain at least two vertices. This means Flipper cannot win in ℓ turns as desired. □

This completes proof of the equivalence (1) ⇔ (4) of Theorem 2.1.

10.4 Predictable Flip-Flatness

We use the remainder of the chapter to give the missing proof of predictable flip-flatness (Proposition 10.7). We say a vertex s is *adjacent* to a vertex set B if $N(s) \cap B \neq \emptyset$.

Definition 10.17. A *classifier* in a graph G is a quadruple $\mathfrak{B} = (\mathcal{B}, S, \mathrm{ex}, \mathrm{rep})$, where $\mathcal{B}$ is a family of pairwise disjoint vertex subsets of G, called further *blobs*, S is a non-empty subset of vertices of G, and $\mathrm{ex}\colon V(G) \to \mathcal{B} \cup \{\perp\}$ and $\mathrm{rep}\colon V(G) \to S$ are mappings satisfying the following properties:

(C.1) $S \cap \bigcup \mathcal{B} = \emptyset$; that is, no vertex of S belongs to any blob.

(C.2) Every $s \in S$ is adjacent either to all the blobs in $\mathcal{B}$ or to no blob in $\mathcal{B}$.

(C.3) For all distinct $s, s' \in S$ and each blob $B \in \mathcal{B}$, $N(s) \cap B \neq N(s') \cap B$.

(C.4) For each $v \in \bigcup \mathcal{B}$, we have $\mathrm{ex}(v) \neq \perp$ and $v \in \mathrm{ex}(v)$.

(C.5) For all $v \in V(G)$ and $B \in \mathcal{B} \setminus \{\mathrm{ex}(v)\}$, we have $N(v) \cap B = N(\mathrm{rep}(v)) \cap B$.

The *size* of a classifier $(\mathcal{B}, S, \mathrm{ex}, \mathrm{rep})$ is $|\mathcal{B}|$, and its *order* is $|S|$.

Let us give some intuition. In a classifier we have a family of disjoint blobs $\mathcal{B}$ and a set of representative vertices S. Further, with every vertex v we can associate its *exceptional blob* $\mathrm{ex}(v) \in \mathcal{B}$ and its *representative* $\mathrm{rep}(v) \in S$. The key condition (C.5) says the following: every vertex v behaves in the same way as its representative $\mathrm{rep}(v)$ with respect to all the blobs in $\mathcal{B}$, except for a (single) exceptional blob $\mathrm{ex}(v)$. We allow $\mathrm{ex}(v)$ to be equal to $\perp$, which indicates that v has no exceptional blob (this will be convenient in notation). Condition (C.4) says that if v is contained in some blob $B \in \mathcal{B}$, then in fact B must be the exceptional blob of v. Conditions (C.1), (C.2), and (C.3) are technical assertions that expresses that the representative set S is reasonably chosen.

A classifier naturally partitions the vertex set of the graph, as formalized below.

Definition 10.18. For a classifier $\mathfrak{B} = (\mathcal{B}, S, \mathrm{ex}, \mathrm{rep})$, the *partition raised* by $\mathfrak{B}$ is the partition $\Pi_{\mathfrak{B}}$ of the vertex set of G defined as follows:

$$\Pi_{\mathfrak{B}} := \{\mathrm{rep}^{-1}(s)\colon s \in S\}.$$

For $s \in S$, we write $\Pi_{\mathfrak{B}}(s) := \mathrm{rep}^{-1}(s)$ to indicate the part of $\Pi_{\mathfrak{B}}$ associated with s.

The following observation is easy, but will be the key to our use of classifiers.

Observation 10.19. Let $\mathfrak{B} = (\mathcal{B}, S, \mathrm{ex}, \mathrm{rep})$ be classifier of size at least five in a graph G. Then for every pair of vertices u, v of G, the following conditions are equivalent.

1. u and v are in the same part of $\Pi_{\mathfrak{B}}$.
2. u and v have the same neighborhood in at least three blobs from $\mathcal{B}$.
3. u and v have different neighborhoods in at most two blobs from $\mathcal{B}$.

Proof. Implication (1)→(3) follows by observing that since $\mathrm{rep}(u) = \mathrm{rep}(v)$, u and v must have exactly the same neighborhood in every blob, possibly except for $\mathrm{ex}(u)$ and $\mathrm{ex}(v)$. The implication (3)→(2) is immediate due to $|\mathcal{B}| \geqslant 5$.

Finally, for implication (2)→(1), observe that u and $\mathrm{rep}(u)$ have the same neighborhood in all but at most one blob from $\mathcal{B}$, and similarly for v and $\mathrm{rep}(v)$. Since u and v have the same neighborhood in at least three blobs from $\mathcal{B}$, it follows that $\mathrm{rep}(u)$ and $\mathrm{rep}(v)$ have the same neighborhood in at least one blob from $\mathcal{B}$. By condition (C.3) of Definition 10.17, this means that $\mathrm{rep}(u) = \mathrm{rep}(v)$, so u and v belong to the same part of $\Pi_{\mathfrak{B}}$. □

From Observation 10.19 we can derive a canonicity property for classifiers: whenever two classifiers share at least five blobs in common, the associated partitions are the same. In the next lemma we show an even stronger property: (efficient) predictability for classifiers.

Lemma 10.20. *For every graph G and family $\mathcal{B}^\circ$ consisting of five pairwise disjoint subsets of $V(G)$, there is a unique partition Π° of $V(G)$ with the following property:*

$$\text{for every classifier } \mathfrak{B} = (\mathcal{B}, S, \text{ex}, \text{rep}) \text{ in } G \text{ with } \mathcal{B}^\circ \subseteq \mathcal{B}, \text{ we have } \Pi^\circ = \Pi_{\mathfrak{B}}. \qquad (*)$$

There is an algorithm that given G and $\mathcal{B}^\circ$ and an integer $k \in \mathbb{N}$ in time $O(k \cdot |V(G)|^2)$ either

- *computes a partition Π° of size at most k that satisfies $(*)$, or*
- *concludes that there is no classifier $\mathfrak{B} = (\mathcal{B}, S, \text{ex}, \text{rep})$ in G with $\mathcal{B}^\circ \subseteq \mathcal{B}$ and $|\Pi_{\mathfrak{B}}| \leqslant k$.*

Proof. We first present the construction of Π°. Along the way we also construct a set of representatives S, and at each point vertices $s \in S$ are in one-to-one correspondence with parts $\Pi^\circ(s)$ of Π°. We start with $\Pi^\circ = \emptyset$ and $S = \emptyset$. Then we iterate through the vertices of G in any order, and when considering the next vertex v we include it in the partition as follows:

- If there exists $s \in S$ such that v and s have the same neighborhood in at least three of the blobs of $\mathcal{B}^\circ$, select such s that was added the earliest to S and add v to $\Pi^\circ(s)$.
- Otherwise, if no s as above exists, add v to S and associate with v a new part $\Pi^\circ(v) = \{v\}$.

If at any point the size of S exceeds k, we abort the execution, as we have found a witness for the fact that there is no classifier $\mathfrak{B} = (\mathcal{B}, S, \text{ex}, \text{rep})$ in G with $\mathcal{B}^\circ \subseteq \mathcal{B}$ and $|\Pi_{\mathfrak{B}}| \leqslant k$: the set S contains $k+1$ vertices that pairwise have a different neighborhood on at least three blobs $\mathcal{B}^\circ$, so no two of them can be in the same part of $\Pi_{\mathfrak{B}}$.

It is straightforward to implement the algorithm to work in time $O(k \cdot |V(G)|^2)$. Assume now the execution was not aborted, so Π° has size at most k and is a partition of $V(G)$. We verify that Π° constructed in this manner satisfies $(*)$. Let then $\mathfrak{B} = (\mathcal{B}, S, \text{ex}, \text{rep})$ be any classifier with $\mathcal{B}^\circ \subseteq \mathcal{B}$; we need to argue that $\Pi^\circ = \Pi_{\mathfrak{B}}$. Consider any pair u, v of vertices of G. We need to prove that u, v are in the same part of Π° if and only if they are in the same part of $\Pi_{\mathfrak{B}}$.

For the forward implication, suppose u and v belong to the same part of Π°, say $\Pi^\circ(s)$ for some $s \in S$. By construction, u and s have the same neighborhood in at least three of the blobs of $\mathcal{B}^\circ$. By Observation 10.19, this implies that u and s are in the same part of $\Pi_{\mathfrak{B}}$. Similarly, v and s are in the same part of $\Pi_{\mathfrak{B}}$. By transitivity, u and v are in the same part of $\Pi_{\mathfrak{B}}$.

For the other direction, suppose $u \in \Pi^\circ(s)$ and $v \in \Pi^\circ(s')$ for some $s \neq s'$. By symmetry, we may assume that s' was added to S later than s. Since v was included in $\Pi^\circ(s')$ instead of $\Pi^\circ(s)$, by construction it follows that v and s must have different neighborhoods in at least three different blobs of $\mathcal{B}^\circ$. So by Observation 10.19, v and s belong to different parts of $\Pi_{\mathfrak{B}}$. Since u and s belong to the same part of Π°, by the forward implication they also belong to the same part of $\Pi_{\mathfrak{B}}$. Hence, u and v belong to different parts of $\Pi_{\mathfrak{B}}$. □

We next observe that, at the cost of a moderate loss on the size of a classifier, we may choose the representatives quite freely.

Lemma 10.21. *Let $\mathfrak{B} = (\mathcal{B}, S, \text{ex}, \text{rep})$ be a classifier in a graph G and let S' be any set such that* rep *is a bijection from S' to S. Then there is a classifier $\mathfrak{B}' = (\mathcal{B}', S', \text{ex}', \text{rep}')$ in G such that $\mathcal{B}' \subseteq \mathcal{B}$ and $|\mathcal{B}'| \geqslant |\mathcal{B}| - |S|$.*

Proof. Let $\mathcal{B}'$ be obtained from $\mathcal{B}$ by removing $\text{ex}(s')$ for each $s' \in S'$. Note that by condition (C.4) of Definition 10.17, S' is disjoint from $\bigcup \mathcal{B}'$. Next, for each vertex u set $\text{ex}'(u) := \text{ex}(u)$, except for the case when $\text{ex}(u) \in \mathcal{B} \setminus \mathcal{B}'$; then set $\text{ex}'(u) := \bot$. Since rep is a bijection from S'

to S, for every vertex u there exists exactly one vertex $s' \in S'$ satisfying $\text{rep}(u) = \text{rep}(s')$, and we set $\text{rep}'(u) := s'$.

We claim that $\mathfrak{B}' := (\mathcal{B}', S', \text{ex}', \text{rep}')$ is a classifier. For this, observe that for every $s' \in S'$, since $\text{ex}(s')$ has been removed when constructing $\mathcal{B}'$, we in fact have $N(\text{rep}(s')) \cap B = N(s') \cap B$ for every $B \in \mathcal{B}'$. With this observation in mind, all conditions of Definition 10.17 for $\mathfrak{B}'$ follow directly from those for $\mathfrak{B}$. □

Let G be a well-ordered graph (cf. Definition 10.3). We shall say that a classifier $\mathfrak{B} = (\mathcal{B}, S, \text{ex}, \text{rep})$ in G is *canonical* if the following condition holds: each $s \in S$ is the minimum element of $\Pi_{\mathfrak{B}}(s)$ according to the well-order of G. We note the following.

Corollary 10.22. *Let G be a well-ordered graph, and $\mathfrak{B} = (\mathcal{B}, S, \text{ex}, \text{rep})$ be a classifier in G. Then there is also a canonical classifier $\mathfrak{B}' = (\mathcal{B}', S', \text{ex}', \text{rep}')$ such that $|S'| = |S|$, $\mathcal{B}' \subseteq \mathcal{B}$, $|\mathcal{B}'| \geqslant |\mathcal{B}| - |S|$.*

Proof. It suffices to apply Lemma 10.21 to $S' := \{\min_G \Pi_{\mathfrak{B}}(s) \colon s \in S\}$. □

We use the insulator framework of Part II to extract large canonical classifiers.

Lemma 10.23. *Fix a monadically stable graph class $\mathcal{C}$ and $r \in \mathbb{N}$. For every well-ordered graph $G \in \sigma(\mathcal{C})$ and family $\mathcal{B}_0$ of pairwise disjoint r-balls in G, there exists a canonical classifier $\mathfrak{B} = (\mathcal{B} \subseteq \mathcal{B}_0, S, \text{ex}, \text{rep})$ of size $U_{\mathcal{C},r}(|\mathcal{B}_0|)$ and order $\text{const}(\mathcal{C}, r)$ in G.*

Proof. We first build an orderless insulator (Definition 5.2.) from $\mathcal{B}_0$. Fix an arbitrary enumeration $a_1, \dots, a_n$ of the centers of the balls in $\mathcal{B}_0$. Let A_0 be the orderless grid of height $r + 1$ defined as

$$A_0[i, 1] = \{a_i\} \quad \text{and} \quad A_0[i, \leqslant t] = N^G_{t-1}[a_i] \text{ for all } t \in [r + 1].$$

Hence, for every $i \in [n]$, the column $A_0[i, *]$ is exactly the r-ball centered at a_i. It is easy to verify that $\mathcal{A}_0 := (A_0, \{V(G)\}, \emptyset, \emptyset)$ is an orderless insulator of height $r + 1$ and cost 1. As $\mathcal{C}$ is monadically stable it is in particular monadically dependent and therefore also prepattern-free. Then there exists a bound $t \leqslant \text{const}(\mathcal{C}, r)$ such that G contains no prepattern of order t on any subinsulator of $\mathcal{A}_0$. Applying Lemma 5.53 to the insulator $\mathcal{A}_0$ and the bound t yields

- a subinsulator $\mathcal{A}$ of $\mathcal{A}_0$ that is indexed by a sequence of length $U_{\mathcal{C},r}(|\mathcal{B}_0|)$, and
- a set $S \subseteq V(G) \setminus \mathcal{A}$ of size at most $\text{const}(\mathcal{C}, r)$,

such that either

1. G contains a prepattern of order t on $\mathcal{A}$, or
2. $\mathcal{A}$ is orderable, or
3. S symmetrically samples G on $\mathcal{A}$ with margin 1.

By prepattern-freeness of $\mathcal{C}$ we rule out the first outcome. As $\mathcal{A}$ being orderable (Definition 5.21) requires the existence of a large semi-induced half-graph, monadic stability of $\mathcal{C}$ rules out the second outcome. We are therefore guaranteed the third outcome. We next build a classifier from the sampling set S.

Let $\mathcal{B}$ be the set of columns of $\mathcal{A}$. By definition of an orderless subinsulator (Definition 5.5) and construction of $\mathcal{A}_0$, $\mathcal{B}$ is a subset of $\mathcal{B}_0$ with $|\mathcal{B}| \geqslant U_{\mathcal{C},r}(|\mathcal{B}_0|)$. Since S is a sampling set (Definition 5.37), for every vertex $v \in V(G)$ there exists a representative $\text{rep}(v) \in S$ and an exceptional ball $\text{ex}(v) \in \mathcal{B}$ such that for every ball $B \in \mathcal{B}$ with $B \neq \text{ex}(v)$ we have

$$\text{atp}_G(v/B) = \text{atp}_G(\text{rep}(v)/G).$$

Let us verify the classifier conditions for $\mathfrak{B} := (\mathcal{B}, S, \text{ex}, \text{rep})$.

- Conditions (C.1) and (C.5) are satisfied by construction.
- Condition (C.4) states that any vertex $v \in B \in \mathcal{B}$ has satisfies $\mathrm{ex}(v) = B$. This is the case as we have $(=, v) \in \mathrm{atp}_G(v/B)$ but $(=, v) \notin \mathrm{atp}_G(\mathrm{rep}(v)/B)$.

To also satisfy conditions (C.2) and (C.3) we have to slightly modify $\mathfrak{B}$.

- Condition (C.2) states that each vertex in S is either adjacent to all balls of $\mathcal{B}$ or non-adjacent to all balls of $\mathcal{B}$. By iterating the pigeonhole principle we can pass to a subset of $\mathcal{B}$ of size $|\mathcal{B}|/2^{|S|} \geqslant U_{\mathcal{C},r}(|\mathcal{B}_0|)$ where this condition holds. This requires updating $\mathrm{ex}(v) := \bot$ for every vertex v whose exceptional ball was dropped from $\mathcal{B}$.
- Condition (C.3) states that every two distinct vertices from S have a different neighborhood on every ball of $\mathcal{B}$. By Ramsey's theorem we can pass to a subset of $\mathcal{B}$ of size $U_{|S|}(|\mathcal{B}|) \geqslant U_{\mathcal{C},r}(|\mathcal{B}_0|)$ such that every two distinct vertices from S either
  - have the same neighborhood on every ball of $\mathcal{B}$, or
  - have a different neighborhood on every ball of $\mathcal{B}$.

  By pruning "duplicates" from S, we can assume the latter. We omit the easy details on how to update $\mathrm{ex}(\cdot)$ and $\mathrm{rep}(\cdot)$.

Having constructed the classifier $\mathfrak{B}$ we finish the proof by making it canonical using Corollary 10.22. □

We are finally ready to prove Proposition 10.7 which we restate for convenience.

Proposition 10.7 (Predictable flip-flatness). *There is an algorithm that takes as input $r, k \in \mathbb{N}$, a well-ordered graph G, and a size five set $Z \subseteq V(G)$, and computes in time $O_{r,k}(|V(G)|^2)$ a k-flip* $\mathrm{Predict}(r, k, G, Z)$ *of G with the following properties:*
For every monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$ there is a bound $k_{\mathcal{C},r} \leqslant \mathrm{const}(\mathcal{C}, r)$ and functions $\mathrm{Flip}_{\mathcal{C},r}$ and $\mathrm{Flat}_{\mathcal{C},r}$ such that for all well-ordered graphs $G \in \mathcal{C}$, sets $X, Z \subseteq V(G)$ and integers $k \geqslant k_{\mathcal{C},r}$ we have

- $\mathrm{Flip}_{\mathcal{C},r}(G, X)$ *is a $k_{\mathcal{C},r}$-flip of G,*
- $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *is a size $U_{\mathcal{C},r}(|X|)$ subset of X,*
- $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *is distance-r independent in* $\mathrm{Flip}_{\mathcal{C},r}(G, X)$*, and*
- *if Z is a size 5 subset of* $\mathrm{Flat}_{\mathcal{C},r}(G, X)$ *then* $\mathrm{Predict}(r, k, G, Z) = \mathrm{Flip}_{\mathcal{C},r}(G, X)$.

Proof. The proof proceeds by induction on r.

Case 1: Base Case. For $r = 0$, we may simply set

$$\mathrm{Flip}_{\mathcal{C},0}(G, X) := G \qquad \text{and} \qquad \mathrm{Flat}_{\mathcal{C},0}(G, X) := X.$$

Case 2: Inductive Case. We first define the functions $\mathrm{Flip}_{\mathcal{C},r}$ and $\mathrm{Flat}_{\mathcal{C},r}$. For this, let us consider any well-ordered graph $G \in \sigma(\mathcal{C})$ and set $X \subseteq V(G)$. Let $H := \mathrm{Flip}_{\mathcal{C},r-1}(G, X)$ be the $\mathrm{const}(\mathcal{C}, r)$-flip of G in which the size $U_{\mathcal{C},r}(|X|)$ set $Y_{r-1} := \mathrm{Flat}_{\mathcal{C},r-1}(G, X)$ is $(r-1)$-independent. There is a monadically stable graph class $\mathcal{D}$ that depends only on $\mathcal{C}$ and r such that $H \in \sigma(\mathcal{D})$. For convenience, we denote

$$r' := \lceil r/2 \rceil - 1.$$

Let $\mathcal{B}_0$ be the family of r'-balls in H whose centers are the vertices of Y_{r-1}. Note that the balls of $\mathcal{B}_0$ are pairwise disjoint. We apply Lemma 10.23 to radius r', graph $H \in \sigma(\mathcal{D})$, and family of r'-balls $\mathcal{B}_0$, thus obtaining a canonical classifier $\mathfrak{B} = (\mathcal{B}, S, \mathrm{ex}, \mathrm{rep})$ with

$$|\mathcal{B}| \geqslant U_{\mathcal{D},r'}(|\mathcal{A}|) \geqslant U_{\mathcal{C},r}(|\mathcal{A}|) = U_{\mathcal{C},r}(|Y_{r-1}|) \geqslant U_{\mathcal{C},r}(|X|)$$

and similarly

$$|S| \leqslant \mathrm{const}(\mathcal{D}, r') \leqslant \mathrm{const}(\mathcal{C}, r).$$

Up to applying Ramsey's theorem, we can assume that: either the centers of the balls in $\mathcal{B}$ are pairwise at distance greater than r in H, or they are pairwise at distance exactly r in H. We define Y to be the set of centers of the balls in $\mathcal{B}$ and set $\mathrm{Flat}_{\mathcal{C},r}(G, X) := Y$.

It remains to construct a $\mathrm{const}(\mathcal{C}, r)$-flip $\mathrm{Flip}_{\mathcal{C},r}(G, X)$ of G in which Y is distance-r independent. As H is already a $\mathrm{const}(\mathcal{C}, r)$-flip of G, it suffices to construct a $\mathrm{const}(\mathcal{C}, r)$-flip of H with this property instead (cf. Lemma 4.3).

Case 2.1: The vertices of Y are pairwise at distance greater than r in H. Then Y is already distance-r independent, and we simply set $\mathrm{Flip}_{\mathcal{C},r}(G, X) := H$.

Case 2.2: The vertices of Y are pairwise at distance exactly r in H. We may assume that $|Y| \geqslant 5$, as otherwise there is a trivial 2^4-flip of G in which every vertex of Y is isolated. In what follows, whenever speaking about adjacencies or distances, we mean adjacencies and distances in H.

Recall that $\mathfrak{B}$ is canonical hence

$$s = \min_G \Pi_{\mathfrak{B}}(s) \qquad \text{for every } s \in S.$$

By definition, every vertex of S is adjacent either to all the balls in $\mathcal{B}$, or to none. Further, since vertices of S have pairwise different neighborhoods in every ball $B \in \mathcal{B}$, there is at most one vertex of S that is not adjacent to any ball of $\mathcal{B}$. Let $S_+ \subseteq S$ consist of those vertices of S that are adjacent to every ball in $\mathcal{B}$; thus either $S_+ = S$ or $|S \setminus S_+| = 1$. Let

$$W := \bigcup_{s \in S_+} \Pi_{\mathfrak{B}}(s).$$

We observe that the vertices of W are the ones that keep the vertices of Y at close distance, in the following sense.

Claim 10.24. *For every vertex $v \in V(G)$, the following conditions are equivalent:*

1. *v belongs to W;*
2. *v is at distance exactly $r' + 1$ from all the vertices of Y, possibly except for one;*
3. *v is at distance at most $r' + 1$ from at least two vertices of Y.*

Proof. Implication (2)→(3) is trivial due to $|Y| \geqslant 5$.

For implication (1)→(2) we use that $\mathfrak{B}$ is a classifier. Let $s \in S_+$ be such that $v \in \Pi_{\mathfrak{B}}(s)$. By definition, s is adjacent to all the balls in $\mathcal{B}$. Therefore, v is adjacent to all the balls in $\mathcal{B}$, possibly except for $\mathrm{ex}(v)$. Recalling that $v \in \mathrm{ex}(v)$ in case $v \in \bigcup \mathcal{B}$, we conclude that v is at distance exactly $r' + 1$ from the centers of all the balls in $\mathcal{B}$, that is, vertices of Y, possibly except for one – the center of $\mathrm{ex}(v)$.

We are left with implication (3)→(1). Let $s \in S$ be such that $v \in \Pi_{\mathfrak{B}}(s)$. As v is at distance at most $r' + 1$ from the centers of two balls in $\mathcal{B}$, at least one of them, say B, is different from $\mathrm{ex}(v)$. In particular $v \notin B$, so v being at distance $r' + 1$ from the center of B means that v has to be adjacent to B. As $B \neq \mathrm{ex}(v)$, we infer that s is also adjacent to B. It follows that $s \in S_+$, implying that $v \in W$. ■

Next, we make a case distinction depending on whether r is odd or even. In both cases, we use the following notation. For $s \in S$ and $U \subseteq S$, we write

$$Q_{s,U} := \{v \in \Pi_{\mathfrak{B}}(s) : N_H(v) \cap S = U\}.$$

Further, we let

$$\mathcal{Q} := \{Q_{s,U} : s \in S, U \subseteq S\}. \tag{10.1}$$

Note that $\mathcal{Q}$ is a partition of the vertex set of H into at most $|S| \cdot 2^{|S|} \leqslant \mathrm{const}(\mathcal{C}, r)$ parts, and the definition of $\mathcal{Q}$ only depends on the graph H, partition $\Pi_{\mathfrak{B}}$, and set S. We set $\mathrm{Flip}_{\mathcal{C},r}(G, X) := H \oplus F$ to be the $\mathcal{Q}$-flip of H specified by the symmetric relation $F \subseteq \mathcal{Q}^2$ which we define in the following. In order to later prove the predictability property, it will be crucial that, in both of the following two cases, the definition of F only depends on the partition $\mathcal{Q}$ (and therefore on H, $\Pi_{\mathfrak{B}}$, and S) and the set S_+.

Case 2.2.1: r is odd. We define F as the set of all pairs $(Q_{s_1,U_1}, Q_{s_2,U_2}) \subseteq \mathcal{Q}^2$ satisfying the following conditions:

- $s_1, s_2 \in S_+$;
- $Q_{s_1,U_1} \neq \emptyset$ and $Q_{s_2,U_2} \neq \emptyset$; and
- $s_1 \in U_2$ or $s_2 \in U_1$.

As desired, F is symmetric and depends only on $\mathcal{Q}$ and S_+. The following claim explains the flip set F in more friendly terms.

Claim 10.25. *For any $u_1, u_2 \in V(G)$, the adjacency between the two is flipped in $H \oplus F$ if and only if $u_1, u_2 \in W$ and $\big(u_2 \in N_H(\mathrm{rep}(u_1))$ or $u_1 \in N_H(\mathrm{rep}(u_2))\big)$.*

Proof. Let s_1, U_1, s_2, U_2 be such that $u_1 \in Q_{s_1,U_1}$ and $u_2 \in Q_{s_2,U_2}$; in particular $Q_{s_1,U_1} \neq \emptyset$ and $Q_{s_2,U_2} \neq \emptyset$. The adjacency between u_1 and u_2 was flipped if and only if $(Q_{s_1,U_1}, Q_{s_2,U_2}) \in F$, which in turn is equivalent to the conjunction of conditions $s_1, s_2 \in S_+$ and ($s_1 \in U_2$ or $s_2 \in U_1$). It now remains to note that condition $s_1, s_2 \in S_+$ is equivalent to $u_1, u_2 \in W$, and condition ($s_1 \in U_2$ or $s_2 \in U_1$) is equivalent to ($u_2 \in N_H(\mathrm{rep}(u_1))$ or $u_1 \in N_H(\mathrm{rep}(u_2))$). ■

Further, we note that the vertices of W may only lie outside the balls of $\mathcal{B}$ or on their boundaries.

Claim 10.26. *If $v \in W$, then for every $y \in Y$ we have $\mathrm{dist}_H(v, y) \geqslant r'$.*

Proof. Suppose $\mathrm{dist}_H(v, y) \leqslant r' - 1$ for some $y \in Y$. As $v \in W$, by Claim 10.24 there exists some other $y' \in Y$, $y' \neq y$, such that $\mathrm{dist}_H(v, y') = r' + 1$. Hence, $\mathrm{dist}_H(y, y') \leqslant 2r' = r - 1$. This is a contradiction with the assumption that Y is $(r-1)$-independent in H. ■

We are now ready to argue the following: Y is distance-r independent in $H \oplus F$. See Figure 10.1 for an illustration. For contradiction, suppose in $H \oplus F$ there exists a path P of length at most r connecting some distinct $y_1, y_2 \in Y$. Let $B_1, B_2 \in \mathcal{B}$ be the r'-balls with centers y_1, y_2, respectively. Since the flips of F only affect the adjacency between the vertices of W, and these vertices have to be at distance at least $r' = \frac{r-1}{2}$ from y_1, y_2 due to Claim 10.26, we infer the following: P can be written as

$$P = (y_1, \ldots, v_1, v_2, \ldots, y_2),$$

where $(y_1, \ldots, v_1)$ and $(v_2, \ldots, y_2)$ are paths of length r' in H that are entirely contained in B_1 and in B_2, respectively. In particular, P has length exactly $2r' + 1 = r$ and v_1v_2 is the only edge on P that might have been flipped in $H \oplus F$.

Observe that if the edge v_1v_2 appeared when applying the flip F, then we necessarily have $v_1, v_2 \in W$. Otherwise, if v_1v_2 was present in H, then path P witnesses that already in H, both v_1 and v_2 are at distance at most $r' + 1$ from both y_1 and y_2. By Claim 10.24, this implies that $v_1, v_2 \in W$. So in any case, we have $v_1, v_2 \in W$.

Let $s_1 := \text{rep}(v_1)$ and $s_2 := \text{rep}(v_2)$. Since $v_1 \in B_1$ and $v_2 \in B_2$, we have $\text{ex}(v_1) = B_1$ and $\text{ex}(v_2) = B_2$, hence

$$N_H(s_1) \cap B_2 = N_H(v_1) \cap B_2 \qquad \text{and} \qquad N_H(s_2) \cap B_1 = N_H(v_2) \cap B_1.$$

In particular,

$$v_1, v_2 \text{ are adjacent in } H \quad \Leftrightarrow \quad v_1, s_2 \text{ are adjacent in } H \quad \Leftrightarrow \quad v_1 \in N_H(s_2),$$

and similarly

$$v_1, v_2 \text{ are adjacent in } H \quad \Leftrightarrow \quad s_1, v_2 \text{ are adjacent in } H \quad \Leftrightarrow \quad v_2 \in N_H(s_1).$$

Therefore,

$$v_1, v_2 \text{ are adjacent in } H \quad \Leftrightarrow \quad (v_1 \in N_H(s_2) \text{ or } v_2 \in N_H(s_1)).$$

As $v_1, v_2 \in W$, by Claim 10.25 we conclude that v_1 and v_2 are adjacent in H if and only if their adjacency gets flipped in $H \oplus F$. So v_1 and v_2 are non-adjacent in $H \oplus F$, a contradiction with the existence of the edge v_1v_2 on P.

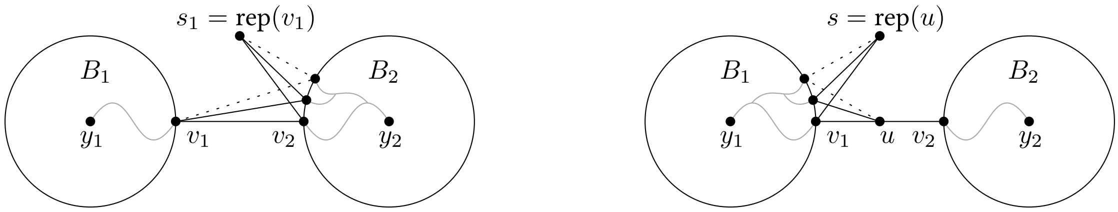

Figure 10.1: The left side depicts Case 2.2.1: v_1 has the same adjacency to B_2 as s_1, hence the edge v_1v_2 is flipped away when applying F if and only if it was present in H.
The right side depicts Case 2.2.2: up to symmetry u has the same adjacency to B_1 as s, hence the edge uv_1 is flipped away when applying F if and only if it was present in H.

Case 2.2.2: r is even. This time, F is defined as the set of all pairs $(Q_{s_1,U_1}, Q_{s_2,U_2}) \in \mathcal{Q}^2$ satisfying the following conditions:

- $Q_{s_1,U_1} \neq \emptyset$ and $Q_{s_2,U_2} \neq \emptyset$; and
- ($s_1 \in S_+$ and $s_1 \in U_2$) or ($s_2 \in S_+$ and $s_2 \in U_1$).

Again, F is symmetric and depends only on $\mathcal{Q}$ and S_+. Also, we may similarly explain flipping according to F as follows.

Claim 10.27. *For any $u_1, u_2 \in V(G)$, applying F flips the adjacency between u_1 and u_2 if and only if $\big(u_1 \in W$ and $u_2 \in N_H(\text{rep}(u_1))\big)$ or $\big(u_2 \in W$ and $u_1 \in N_H(\text{rep}(u_2))\big)$.*

Proof. Analogous to the proof of Claim 10.25, we leave the details to the reader. ∎

Note that Claim 10.25 implies in particular that whenever the adjacency between two vertices is flipped in $H \oplus F$, at least one of them belongs to W. (However, contrary to the odd case, there might be vertices outside W that are affected by the flip.) In this vein, the following observation will be convenient.

Claim 10.28. $W \cap \bigcup \mathcal{B} = \emptyset$.

Proof. For contradiction, suppose there exists $B \in \mathcal{B}$ and $v \in B$ such that $v \in W$. Letting y be the center of B, we have $\mathrm{dist}_H(v, y) \leqslant r'$. By Claim 10.24, there exists another $y' \in Y$, $y' \neq y$, such that $\mathrm{dist}_H(v, y') \leqslant r' + 1$. Hence, $\mathrm{dist}_H(y, y') \leqslant 2r' + 1 = r - 1$, contradicting the distance-$(r-1)$ independence of Y in H. ■

As in the odd case, we are left with arguing that Y is distance-r independent in $H \oplus F$. See Figure 10.1 for an illustration. For contradiction, suppose that there exist distinct $y_1, y_2 \in Y$ and a path P of length at most r that connects y_1 and y_2 in $H \oplus F$. As before, let $B_1, B_2 \in \mathcal{B}$ be the balls with centers y_1, y_2, respectively.

By Claim 10.27, only the vertices of $W \cup \bigcup_{s \in S_+} N_H(s)$ are affected by the flip $H \oplus F$. By Claim 10.28 and as S_+ is disjoint with $\bigcup \mathcal{B}$, all vertices of $W \cup \bigcup_{s \in S_+} N_H(s)$ are at distance (in H) at least r' from all the vertices of Y. Since $r = 2r' + 2$, similarly as in Case 2.2.1 it follows that P has length $2r' + 1 = r - 1$ or $2r' + 2 = r$ and can be written as

$$P = (y_1, \ldots, v_1, v_2, \ldots, y_2) \qquad \text{or} \qquad P = (y_1, \ldots, v_1, u, v_2, \ldots, y_2),$$

where $(y_1, \ldots, v_1)$ and $(v_2, \ldots, y_2)$ are paths of length r' in H entirely contained in B_1 and B_2, respectively.

In the first case, P has length $r - 1$ and is of the form $(y_1, \ldots, v_1, v_2, \ldots, y_2)$. Observe that edge v_1v_2 cannot be present in H, because then P would be entirely contained in H, a contradiction with distance-$(r-1)$ independence of Y in H. On the other hand, note that $v_1, v_2 \notin W$ due to Claim 10.28, so by Claim 10.27 the adjacency between v_1 and v_2 is not flipped in $H \oplus F$. We conclude that v_1 and v_2 remain non-adjacent in $H \oplus F$, a contradiction with the presence of the edge v_1v_2 on P.

In the second case, P has length r and is of the form $(y_1, \ldots, v_1, u, v_2, \ldots, y_2)$. Let us first argue that $u \in W$. If u is adjacent both to v_1 and to v_2 in H, then u is at distance at most $r' + 1$ from both y_1 and y_2 in H, hence that $u \in W$ follows directly from Claim 10.24. On the other hand, if u is non-adjacent in H to one of v_1 or v_2, say to v_1, then the adjacency between u and v_1 must get flipped when applying F. By Claim 10.27 this means that at least one of u and v_1 belongs to W, but it cannot be v_1 due to Claim 10.28. So $u \in W$ in this case as well.

Let $s := \mathrm{rep}(u)$. By symmetry, we may assume that $B_1 \neq \mathrm{ex}(u)$. This means that

$$v_1, u \text{ are adjacent in } H \quad \Leftrightarrow \quad v_1, s \text{ are adjacent in } H \quad \Leftrightarrow \quad v_1 \in N_H(s).$$

Since $u \in W$ and $v_1 \notin W$ (due to Claim 10.28), by Claim 10.27 we conclude that u and v_1 are adjacent in H if and only if their adjacency gets flipped when applying F. So in any case, u and v_1 are non-adjacent in $H \oplus F$. This is a contradiction with the presence of the edge uv_1 on P.

This concludes the construction of the $\mathrm{const}(\mathcal{C}, r)$-flip $\mathrm{Flip}_{\mathcal{C},r}(G, X) = H \oplus F$. Hence, we can choose $k_{\mathcal{C},r} \leqslant \mathrm{const}(\mathcal{C}, r)$ such that $\mathrm{Flip}_{\mathcal{C},r}(G, X)$ is a $k_{\mathcal{C},r}$ of G.

The prediction algorithm. It remains to provide the algorithm $\mathrm{Predict}(r, k, G, Z)$. We follow the definition of $\mathrm{Flip}_{\mathcal{C},r}$ and $\mathrm{Flat}_{\mathcal{C},r}$ and define the algorithm by induction on r. For the base case we can set

$$\mathrm{Predict}(0, k, G, Z) := G$$

for all k, G, Z, which matches the base case of $\mathrm{Flip}_{\mathcal{C},0}$. In the inductive case, for all well-ordered graphs G, sets $Z \subseteq V(G)$ with $|Z| = 5$, and $k \in \mathbb{N}$, we compute the k-flip $\mathrm{Predict}(r, k, G, Z)$ of G as follows.

1. By induction, we compute the flip $H^\circ := \mathrm{Predict}(r-1, k, G, Z)$ of G.
2. If Z is not distance-$(r-1)$ independent in H°, return the original graph G.
3. If Z is distance-r independent in H°, return H°.
4. Otherwise, let $\mathcal{B}^\circ$ consist of the five r'-balls in H° with centers in vertices of Z. Note that the balls of $\mathcal{B}^\circ$ are pairwise disjoint. Apply the algorithm of Lemma 10.20 to the graph G, family $\mathcal{B}^\circ$, and parameter k. If the algorithm concludes that there is no classifier $\mathfrak{B} = (\mathcal{B}, S, \mathrm{ex}, \mathrm{rep})$ in G with $\mathcal{B}^\circ \subseteq \mathcal{B}$ and $|\Pi_{\mathfrak{B}}| \leqslant k$: abort and return the original graph G. Otherwise, we obtain a partition Π°.
5. Let $S^\circ := \{\min_G A : A \in \Pi^\circ\}$ and let S°_+ be the subset of vertices of S° that are adjacent to every ball of $\mathcal{B}^\circ$.
6. Compute the partition $\mathcal{Q}^\circ$ of $V(G)$ from H°, Π°, and S° exactly as $\mathcal{Q}$ was computed from H, $\Pi_{\mathfrak{B}}$, and S in (10.1) above.
7. Compute the relation $F^\circ \subseteq \mathcal{Q}^\circ \times \mathcal{Q}^\circ$ from $\mathcal{Q}^\circ$ and S°_+ exactly as F was computed from $\mathcal{Q}$ and S_+ in Cases 2.2.1 and 2.2.2 above.
8. We check if $H^\circ \oplus F^\circ$ is a k-flip of G using the algorithm from Lemma 4.1. If this is the case, we return $H^\circ \oplus F^\circ$, otherwise we return the original graph G.

By construction the algorithm only returns k-flips of G, as desired. (In particular the original graph G is a k-flip of itself.) We now argue that provided $Y = \mathrm{Flat}_{\mathcal{C},r}(G, X)$, $Z \subseteq Y$ is a set of size 5, and $k \geqslant k_{\mathcal{C},r}$, we have $\mathrm{Flip}_{\mathcal{C},r}(G, X) = \mathrm{Predict}(r, k, G, Z)$. We adopt the notation from the definition of $\mathrm{Flip}_{\mathcal{C},r}$ and $\mathrm{Flat}_{\mathcal{C},r}$ and argue by induction on r. The base case $r = 0$ holds by construction.

For the inductive step we revisit the case distinction above. Without loss of generality, we can assume $k_{\mathcal{C},r} \geqslant k_{\mathcal{C},r-1}$, so also $k \geqslant k_{\mathcal{C},r-1}$. Then, by the induction assumption, we have

$$H^\circ = \mathrm{Predict}(r-1, k, G, Z) = \mathrm{Flip}_{\mathcal{C},r-1}(G, X) = H.$$

In particular, as $Z \subseteq Y$ is distance-$(r-1)$ independent in H, the termination in the second point above cannot happen. Also, if the Y is distance-r independent in H, then the same holds for Z, and we have $\mathrm{Predict}(r, k, G, Z) = H^\circ$ (termination in the third point above). In the definition of $\mathrm{Flip}_{\mathcal{C},r}(G, X)$, Case 2.1 applies, yielding $\mathrm{Flip}_{\mathcal{C},r}(G, X) = H$, as required.

We are left with Case 2.2: the vertices of Y are pairwise at distance exactly r in H. Let $\mathfrak{B} = (\mathcal{B}, S, \mathrm{ex}, \mathrm{rep})$ be the canonical classifier provided by Lemma 10.23 in the construction of Y, whose blobs $\mathcal{B}$ are the r'-balls in H with centers Y. As $Z \subseteq Y$ and $H^\circ = H$, we have that $\mathcal{B}^\circ \subseteq \mathcal{B}$. As witnessed by $|\Pi_{\mathfrak{B}}| \leqslant k_{\mathcal{C},r} \leqslant k$, in the fourth point above, the algorithm of Lemma 10.20 must yield a partition $\Pi^\circ = \Pi_{\mathfrak{B}}$. Since $\mathfrak{B}$ is canonical, we have

$$S = \{\min_G A : A \in \Pi_{\mathfrak{B}}\} = \{\min_G A : A \in \Pi^\circ\} = S^\circ.$$

Similarly, as $\mathfrak{B}$ is a classifier in $H = H^\circ$, we have that a vertex from $S = S^\circ$ is adjacent to every ball of $\mathcal{B}$ if and only if it is adjacent to every ball of $\mathcal{B}^\circ$, and we conclude that $S_+ = S^\circ_+$. As now $H^\circ = H$, $\Pi^\circ = \Pi_{\mathfrak{B}}$, $S^\circ = S$, and $S^\circ_+ = S_+$, both in the definition of $\mathrm{Flip}_{\mathcal{C},r}$ and $\mathrm{Predict}(r, k, G, Z)$, we construct the same partition $\mathcal{Q} = \mathcal{Q}^\circ$. Again $|\mathcal{Q}^\circ| = |\mathcal{Q}| \leqslant k_{\mathcal{C},r} \leqslant k$ so termination in the sixth point above does not happen. Then the construction presented in Cases 2.2.1 and 2.2.2 provides the same relation for $\mathcal{Q}^\circ$ and S°_+, as for $\mathcal{Q}$ and S_+: we have $F^\circ = F$. We conclude that

$$\mathrm{Predict}(r, k, G, Z) = H^\circ \oplus F^\circ = H \oplus F = \mathrm{Flip}_{\mathcal{C},r}(G, X).$$

This graph is a $k_{\mathcal{C},r}$-flip of G and, as $k_{\mathcal{C},r} \leqslant k$, it is not rejected in the last step. Hence, we have $\mathrm{Predict}(r, k, G, Z) = \mathrm{Flip}_{\mathcal{C},r}(G, X)$, as desired.

We finally argue that the graph Predict(r, k, G, Z) can be computed in time $O_{r,k}(|V(G)|^2)$. For this, it is enough to observe that the procedure presented above executes r inductive calls, each of which consists of internal computation that is easy to implement in time $O_{r,k}(|V(G)|^2)$, one call to the algorithm of Lemma 10.20 to compute Π° in time $O(k \cdot |V(G)|^2)$, and one call to the algorithm of Lemma 4.1 which also runs in time $O_k(|V(G)|^2)$. For the runtime bound of the algorithm from Lemma 4.1, we note the following. The algorithm runs in time $O(k \cdot |V(G)|^2)$ if a k-flip of G is supplied as input. While we cannot guarantee that for all values of Z we always give a k-flip to the algorithm, we can guarantee that we only ever supply it with const(k)-flips. This is due to the fact that H° is a k-flip by induction, and we have ensured that Π° also has size at most k, too. Hence, the running time of Lemma 4.1 is still bounded by $O_k(|V(G)|^2)$ and the total algorithm runs in time $O_{r,k}(|V(G)|^2)$. □

The proof of predictable flip-flatness concludes the chapter.

Chapter 11

Neighborhood Covers

In the previous chapters we have shown several combinatorial characterizations for monadic stability (Theorem 2.1). We now work towards the model checking algorithm (Theorem 2.2). Having proved a winning strategy for the Flipper game, we have already established one main ingredient for the algorithm. In this chapter, supply the second ingredient, that is sparse neighborhood covers for monadically stable classes.

For a graph G and vertex subset X, the *weak diameter* of X in G is the maximum distance in G between members of X: $\text{diam}_G(X) := \max_{u,v \in X} \text{dist}_G(u, v)$.

Definition 11.1. Let G be a graph and r be a positive integer. A family $\mathcal{K}$ of subsets of vertices of G is called a *distance-r neighborhood cover* of G if for every vertex u of G there exists $C \in \mathcal{K}$ such that $\text{Ball}_r[u] \subseteq C$. The *diameter* of $\mathcal{K}$ is the maximum weak diameter among the sets of $\mathcal{K}$, while the *overlap* of $\mathcal{K}$ is the maximum number of sets of $\mathcal{K}$ that intersect at a single vertex:

$$\text{diam}(\mathcal{K}) := \max_{C \in \mathcal{K}} \text{diam}_G(C) \qquad \text{and} \qquad \text{overlap}(\mathcal{K}) := \max_{u \in V(G)} |\{C \in \mathcal{K} : u \in C\}|.$$

Elements of a neighborhood cover $\mathcal{K}$ will often be called *clusters*.

The main result of this chapter will be the following.

Theorem 11.2. *There is an algorithm that, given an n-vertex graph G and a radius $r \in \mathbb{N}$, computes a distance-r neighborhood cover of G with diameter at most $4r$ in time $O(n^5)$. For every monadically stable class $\mathcal{C}$ containing G and every $\varepsilon > 0$, the overlap of the cover is bounded by $O_{\mathcal{C},r,\varepsilon}(n^{\varepsilon})$.*

11.1 Neighborhood Complexity

As an important ingredient of Theorem 11.2 and as a result of independent interest, we prove the following theorem.

Theorem 11.3. *Let $\mathcal{C}$ be a monadically stable graph class and $\varepsilon > 0$. Then for every $G \in \mathcal{C}$ and $A \subseteq V(G)$,*

$$|\{N_G[v] \cap A \colon v \in V(G)\}| \leqslant O_{\mathcal{C},\varepsilon}(|A|^{1+\varepsilon}).$$

Given a class $\mathcal{C}$, define the *neighborhood complexity* of $\mathcal{C}$ as the function $\nu_{\mathcal{C}} \colon \mathbb{N} \to \mathbb{N}$ such that

$$\nu_{\mathcal{C}}(n) := \sup_{G \in \mathcal{C}, A \subseteq V(G), |A| = n} |\{N[v] \cap A : v \in V(G)\}|.$$

Note that for every graph class $\mathcal{C}$ we have $\nu_{\mathcal{C}}(n) \leqslant 2^n$ for all $n \in \mathbb{N}$. It is an immediate consequence of the Sauer-Shelah lemma [76, 78, 84] that for every graph class of bounded VC dimension (in particular, every monadically dependent or monadically stable class) there is some constant c such that $\nu_{\mathcal{C}}(n) \leqslant O(n^c)$ for all $n \in \mathbb{N}$. Theorem 11.3 states that every monadically stable graph class $\mathcal{C}$ has *almost linear neighborhood complexity*, that is, $\nu_{\mathcal{C}}(n) \leqslant O_{\mathcal{C},\varepsilon}(n^{1+\varepsilon})$ for all $\varepsilon > 0$. This result is a generalization of an analogous result of Eickmeyer et al. [32] for nowhere dense classes, stated below.

Fact 11.4 ([32]). *Let $\mathcal{C}$ be a nowhere dense graph class and $\varepsilon > 0$. Then for every $G \in \mathcal{C}$ and $A \subseteq V(G)$,*

$$|\{N[v] \cap A : v \in V(G)\}| \leqslant O_{\mathcal{C},\varepsilon}(|A|^{1+\varepsilon}).$$

A similar result holds for all *structurally nowhere dense classes*, that is, classes that can be transduced from a nowhere dense class [70]. However, Theorem 11.3 is incomparable with the statement from [70], as the latter also allows defining neighborhoods using a formula $\varphi(\bar{x}, \bar{y})$ involving tuples of free variables.

In order to prove Theorem 11.3, we will gradually simplify a monadically stable class — while preserving monadic stability, and without decreasing its neighborhood complexity too much — until we arrive at a $K_{t,t}$-free graph class. The next theorem states that monadically stable, $K_{t,t}$-free classes are nowhere dense, so we will be able to conclude using Fact 11.4. The Fact 11.5 below follows from a result of Dvořák [30] (see [66, Corollary 2.3]) and we reprove it in Chapter 13 (Theorem 13.2).

Fact 11.5 (follows from [30]). *Let $\mathcal{C}$ be a monadically stable graph class, and suppose that $\mathcal{C}$ excludes some biclique $K_{t,t}$ as a subgraph. Then $\mathcal{C}$ is nowhere dense.*

Our simplification process will decrease a parameter we call *branching index*, which was introduced by Shelah in model theory and is sometimes referred to as *"Shelah's 2-rank"*. For a bipartite graph $G = (A, B, E)$ and vertex $a \in A$, we denote $\overline{N}_G(a) := B - N_G(a)$.

Definition 11.6. Let $G = (A, B, E)$ be a bipartite graph. The *branching index* of a set $U \subseteq B$, denoted $\mathrm{br}_G(U)$, is defined as

$$\mathrm{br}_G(U) := \begin{cases} -1 & \text{if } U = \emptyset, \\ 1 + \max_{a \in A} \min(\mathrm{br}_G(N(a) \cap U), \mathrm{br}_G(\overline{N}(a) \cap U)) & \text{if } U \neq \emptyset. \end{cases}$$

Note that for $U \subseteq B$ we have that $\mathrm{br}_G(U) = 0$ if and only if U is nonempty and all vertices in U have the same neighborhoods in A. In other words: every vertex in A is either fully adjacent or non-adjacent to U. In particular $|U| = 1$ implies $\mathrm{br}_G(U) = 0$.

For higher values of the branching index, the following perspective might be helpful. Say that $U \subseteq B$ is *split* into sets P, Q by a vertex $a \in A$ if $P = U \cap N_G(a)$ and $Q = U \cap \overline{N}_G(a)$ are both nonempty. Say that U *can be split* into P and Q if P, Q are nonempty and there is some $a \in A$ which splits U into P and Q. Informally, the value $\mathrm{br}_G(U)$ tells us for how many steps we can repeatedly split U into two, four, eight, etc. sets, where in each step, we are required to split each set produced in the previous step into two parts.

A bounded branching index characterizes edge-stable graph classes (Definition 8.1) as made explicit in the following fact proved by Shelah [79, (3) ⇔ (7) in Thm. 2.2 of Chapter II] (see also [49, Lem. 6.7.9] for a presentation that is closer to our use case).

Fact 11.7 ([79], [49, Lem. 6.7.9]). *Let $\mathcal{C}$ be a class of bipartite graphs. $\mathcal{C}$ is edge-stable if and only if there is a number $d \in \mathbb{N}$ such that $\mathrm{br}_G(B) \leqslant d$ for all $G = (A, B, E) \in \mathcal{C}$.*

More precisely, if $\mathcal{C}$ contains no half-graph of order k as a semi-induced half-graphs, then we can choose $d < 2^{k+2} - 2$.

In particular, the branching index is bounded in all monadically stable graph classes. We will use the following statement about definability of the branching index, which can be easily proved by induction on d.

Lemma 11.8. *For every $d \in \mathbb{N}$ there is a first-order sentence β_d over the signature consisting of a binary relation symbol E and unary relation symbols A, B, such that, given a bipartite graph $G = (A, B, E)$, the structure G (where A, B, E are interpreted as the appropriate relations) satisfies β_d if and only if $\mathrm{br}_G(B) \leqslant d$.*

The following proposition, together with Facts 11.4 and 11.5, will later easily yield Theorem 11.3.

Proposition 11.9. *Fix $d \in \mathbb{N}$. There is a transduction T_d with the following properties. Given a bipartite graph $G = (A, B, E)$ with $n := |A| \geqslant 2$ such that no vertices in B have equal neighborhoods and $\mathrm{br}_G(B) \leqslant d$, there is a bipartite graph $G' = (A', B', E') \in \mathsf{T}_d(G)$ with $A' \subseteq A$ and $B' \subseteq B$ such that:*

- *(A.1) $|B'| \geqslant \frac{|B|}{(600 \ln n)^d}$,*
- *(A.2) every vertex $b \in B'$ has at most d neighbors in A' in the graph G',*
- *(A.3) all vertices in B' have distinct neighborhoods in A' in the graph G'.*

The following technical sampling lemma will be used as an ingredient in the proof of Proposition 11.9.

Lemma 11.10. *Let $G = (A, B, E)$ be a bipartite graph such that $|A| \geqslant 2$ and every vertex in B has some neighbor in A. Then there are sets $X \subseteq A$ and $B' \subseteq B$ with $|B'| \geqslant \frac{|B|}{150 \ln |A|}$ such that every vertex $b \in B'$ has exactly one neighbor in X.*

We first give some probabilistic preliminaries. Let X be a random variable which takes values from $a_1, \ldots, a_n$. We denote by $\mathbb{P}[X = a_i]$ the probability that X takes the value a_i. The *expected value* of X is defined as

$$\mathbb{E}[X] = \sum_{i \in [n]} a_i \cdot \mathbb{P}[X = a_i].$$

The *probabilistic method* will be used to show the existence of sets X and B' from Lemma 11.10, as follows. Instead of directly constructing X and a large viable set B', we show that for a randomly chosen X, the expected value of the size of a viable set B' is large. By definition of the expected value, this is sufficient to prove the existence of a set X and a large viable set B' (but does not show how to construct these sets). To impose bounds on the expected values, the following two standard facts from probability theory will come in handy.

Fact 11.11 (Linearity of expectation)**.** *For random variables $X_1, \ldots, X_n$ we have*

$$\mathbb{E}\left[\sum_{i \in [n]} X_i\right] = \sum_{i \in [n]} \mathbb{E}\left[X_i\right].$$

Fact 11.12 (Markov's inequality)**.** *For every non-negative random variable X and real $a > 0$,*

$$\mathbb{P}[X > a] \leqslant \mathbb{E}[X]/a.$$

We are now ready to prove Lemma 11.10.

Proof of Lemma 11.10. Fix a real $\alpha > 1$ to be specified later and denote $n := |A|$. Consider all intervals of the form $[\alpha^i, \alpha^{i+1})$ for some integer $0 \leqslant i \leqslant \log_\alpha n$. For each $b \in B$, the degree of b belongs to exactly one such interval. Therefore, there is some i as above and a set $B_0 \subseteq B$ with

$$B_0 \geqslant |B|/(1 + \log_\alpha n)$$

such that all vertices in B_0 have degree in the interval $[\alpha^i, \alpha^{i+1})$. Set $d = \alpha^i$. Thus, all vertices in B_0 have degree between d and αd.

Pick $X \subseteq A$ by including each vertex of A uniformly at random with probability $1/d$. Consider $b \in B_0$. The expected size of $X \cap N(b)$ is

$$\begin{aligned}
\mathbb{E}[|X \cap N(b)|] &= \mathbb{E}\left[\textstyle\sum_{v \in N(b)} |\{v\} \cap X|\right] \\
&= \textstyle\sum_{v \in N(b)} \mathbb{E}[|\{v\} \cap X|] && \text{(by Fact 11.11)} \\
&= \textstyle\sum_{v \in N(b)} 1 \cdot \mathbb{P}[|v \cap X| = 1] + 0 \cdot \mathbb{P}[|v \cap X| = 0] && \text{(by definition of } \mathbb{E}\text{)} \\
&= \textstyle\sum_{v \in N(b)} 1/d && (\mathbb{P}[v \in X] = 1/d) \\
&= |N(b)| \cdot 1/d.
\end{aligned}$$

Since $d \leqslant |N(b)| \leqslant \alpha d$, we have $1 \leqslant \mathbb{E}[|X \cap N(b)|] \leqslant \alpha$. By Markov's inequality (Fact 11.12),

$$\mathbb{P}\big[|X \cap N(b)| \geqslant 2\big] \leqslant \mathbb{E}\big[|X \cap N(b)|\big]/2 \leqslant \alpha/2.$$

On the other hand,

$$\mathbb{P}\big[|X \cap N(b)| = 0\big] \leqslant (1 - 1/d)^d \leqslant 1/e.$$

This means that

$$\mathbb{P}\big[|X \cap N(b)| = 1\big] = 1 - \mathbb{P}\big[|X \cap N(b)| = 0\big] - \mathbb{P}\big[|X \cap N(b)| \geqslant 2\big] \geqslant 1 - 1/e - \alpha/2 =: \beta.$$

We set $\alpha := 1.1$ and verify that $\beta = 1 - 1/e - \alpha/2 \geqslant \frac{1}{150 \ln \alpha}$. Again by linearity of expectation, the expected number of vertices $b \in B_0$ such that $|X \cap N(b)| = 1$ is at least $\beta|B_0|$. By the probabilistic method, there exists an assignment to X reaching the expected value, so let us fix X according to this assignment. We then set B' to be those elements of B_0 with exactly one neighbor in X. Since $n \geqslant 2$, we have $1 + \log_\alpha n < 2\log_\alpha n$. Therefore,

$$|B'| \geqslant \beta \cdot |B_0| \geqslant \beta \cdot \frac{|B|}{1 + \log_\alpha n} > \frac{1}{150 \ln \alpha} \cdot \frac{|B|}{2\log_\alpha n} = \frac{|B|}{150 \ln \alpha \cdot 2 \cdot \frac{\ln n}{\ln \alpha}} \geqslant \frac{|B|}{300 \ln n}.$$

This concludes the proof. □

The next lemma is the main engine of the proof of Proposition 11.9, and hence of Theorem 11.3. The central definitions of this lemma are also depicted in Figure 11.1.

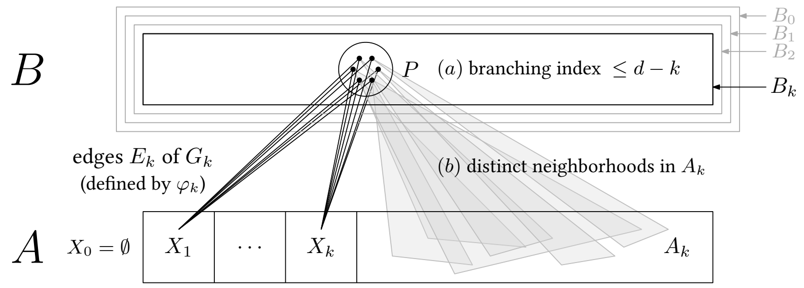

Figure 11.1: Illustration of central definitions in Lemma 11.13

Lemma 11.13. *Fix $d \in \mathbb{N}$. For every $0 \leqslant k \leqslant d$ there is a formula $\varphi_k(x, y)$ in the signature consisting of a binary relation symbol E and unary relation symbols A, B, $X_0, \ldots, X_{k-1}$, $B_0, \ldots, B_{k-1}$, such that the following holds. Given a bipartite graph $G = (A, B, E)$ with $n := |A| \geqslant 2$ such that no two vertices of B have equal neighborhoods and $\mathrm{br}_G(B) \leqslant d$, there exist sets $B = B_0 \supseteq \ldots \supseteq B_k$, pairwise disjoint sets $X_0, \ldots, X_k \subseteq A$, and a bipartite graph $G_k = \big(A, B, E(G_k)\big)$ such that:*

(B.1) $|B_k| \geqslant \frac{|B|}{(600 \ln n)^k}$;

(B.2) $E(G_k) = \{ab \in A \times B : (G, A, B, X_0, \ldots, X_{k-1}, B_0, \ldots, B_{k-1}) \models \varphi_k(a, b)\}$;

(B.3) every $b \in B_k$ has at most k neighbors in $X_0 \cup \cdots \cup X_k$ in the graph G_k; and

(B.4) for every $P \subseteq B_k$ such that all $b \in P$ have the same neighborhood in $X_0 \cup \cdots \cup X_k$ in the graph G_k,

- *(a)* $\mathrm{br}_{G_k}(P) \leqslant d - k$, *and*
- *(b) no two distinct vertices in P have equal neighborhoods in $A - (X_0 \cup \cdots \cup X_k)$ in the graph G_k.*

Proof. We proceed by induction on k. For $k = 0$ the formula $\varphi_0(x, y) := E(x, y)$ satisfies the required conditions. Namely, for a given graph $G = (A, B, E)$ we define $B_0 := B$, $G_0 := G$, $X_0 := \emptyset$, and the required conditions trivially hold.

Assuming the statement holds for some value $k \geqslant 0$, we prove it for $k + 1 \leqslant d$. Let $\varphi_k(x, y)$ be as in the statement. Consider a bipartite graph $G = (A, B, E)$ and let $B = B_0 \supseteq \ldots \supseteq B_k$, $X_0, \ldots, X_k \subseteq A$, and G_k be given by the induction assumption. Denote $A_k := A-(X_0 \cup \cdots \cup X_k)$.

By a *k-class* we mean an inclusion-wise maximal set $P \subseteq B_k$ of vertices in $b \in B_k$ which have equal neighborhoods in $X_0 \cup \cdots \cup X_k$ in the graph G_k. By assumption, $\mathrm{br}_{G_k}(P) \leqslant d - k$ for every k-class P. Therefore, for every vertex $a \in A$ and k-class P, we have

$$\min\big(\mathrm{br}_{G_k}(N_{G_k}(a) \cap P), \mathrm{br}_{G_k}(\overline{N}_{G_k}(a) \cap P)\big) < d - k. \tag{11.1}$$

Define a relation $E_0 \subseteq A_k \times B_k \subseteq A \times B$ as

$$\begin{aligned} E_0 \;&:=\; \bigcup \big\{\{a\} \times P \;\big|\; a \in A_k, P \text{ is a } k\text{-class with } \mathrm{br}_{G_k}(N_{G_k}(a) \cap P) = d - k\big\} \\ &=\; \big\{ab \;\big|\; a \in A_k \text{ and } b \in P \text{ for some } k\text{-class } P \text{ with } \mathrm{br}_{G_k}(N_{G_k}(a) \cap P) = d - k\big\}. \end{aligned}$$

Let $G_{k+1} := (A, B, E(G_k) \triangle E_0)$, where $\triangle$ denotes the symmetric difference. So for a pair ab with $a \in A$ and $b \in B$, we have that $ab \in E(G_{k+1})$ if and only if ab belongs to exactly one of the sets $E(G_k)$ and E_0. The relation $E(G_{k+1})$ is definable by a first-order formula, as stated in the next claim.

Claim 11.14. *There is a first-order formula $\varphi_{k+1}(x, y)$, which is independent of G, such that for all $a \in A$ and $b \in B_k$ we have*

$$(G, A, B, X_0, \ldots, X_k, B_0, \ldots, B_k) \models \varphi_{k+1}(a, b) \quad \Leftrightarrow \quad ab \in E(G_{k+1}).$$

Here $\varphi_{k+1}(x,y)$ is over the signature $\{E, A, B, X_0, \ldots, X_k, B_0, \ldots, B_k\}$.

Proof. By induction, the edge relation of G_k is definable by the formula $\varphi_k(x,y)$ over the signature

$$\{E, A, B, X_0, \ldots, X_{k-1}, B_0, \ldots, B_{k-1}\}.$$

Adding the predicates A_k and B_k lets us define the equivalence relation stating whether two vertices from B_k belong to the same k-class. Combining this with the formula β_{d-k} obtained in Lemma 11.8, we construct a formula $\psi(x,y)$ defining the relation E_0 as defined above. Then $\varphi_{k+1}(x,y)$ is defined as the XOR of the formulas ψ and φ_k. We omit the easy details. ■

Next, we note that the construction of G_{k+1} from G_k did not affect the neighborhoods in $X_0 \cup \cdots \cup X_k$, nor the branching indices of subsets of k-classes.

Claim 11.15. *If $b \in B$, then the neighborhood of b in $X_0 \cup \cdots \cup X_k$ is the same when considered in G_k, and when considered in G_{k+1}.*

Proof. We have that $E(G_{k+1}) = E(G_k) \triangle E_0$, with $E_0 \subseteq A_k \times B_k$ and A_k disjoint from $X_0 \cup \cdots \cup X_k$. ■

Claim 11.16. *Let $P \subseteq B_k$ be a k-class and $Q \subseteq P$. Then $\mathrm{br}_{G_{k+1}}(Q) = \mathrm{br}_{G_k}(Q)$.*

Proof. We show that

$$\{N_{G_k}(a) \cap Q, \overline{N}_{G_k}(a) \cap Q\} = \{N_{G_{k+1}}(a) \cap Q, \overline{N}_{G_{k+1}}(a) \cap Q\} \qquad \text{for every } a \in A. \quad (11.2)$$

So, the two parts into which a splits Q are the same in G_k as in G_{k+1} (where the neighborhood is possibly swapped with the non-neighborhood). Claim 11.16 then follows from (11.2) by a straightforward induction.

Towards (11.2), fix $a \in A$. If $a \in X_0 \cup \cdots \cup X_k$, the statement follows from Claim 11.15. Now, suppose $a \in A - (X_0 \cup \cdots \cup X_k) = A_k$. Then the set $\{a\} \times P$ is either disjoint from E_0, or is contained in E_0. In the first case, we have that $N_{G_k}(a) \cap P = N_{G_{k+1}}(a) \cap P$, while in the latter case, we have that $N_{G_k}(a) \cap P = \overline{N}_{G_{k+1}}(a) \cap P$. Either way, (11.2) follows. ■

Let $B_- \subseteq B_k$ be the set of those vertices $b \in B_k$ such that b has no neighbor in A_k in the graph G_{k+1}, and let $B_+ := B_k - B_-$. We consider two cases, depending on which of the sets B_+, B_- is larger.

Case 1: $|\mathbf{B}_+| \geqslant |\mathbf{B_k}|/2$. Apply Lemma 11.10 to $G_{k+1}[A_k, B_+]$, obtaining sets $B_{k+1} \subseteq B_+ \subseteq B_k$ and $X_{k+1} \subseteq A_k$ such that $|B_{k+1}| \geqslant \frac{|B_+|}{300 \ln n} \geqslant \frac{|B_k|}{600 \ln n}$ and every vertex in B_k has exactly one neighbor in X_{k+1} in G_{k+1}. We check the required properties of A_{k+1} and B_{k+1}. We have $|B_{k+1}| \geqslant \frac{|B_k|}{600 \ln n} \geqslant \frac{|B|}{(600 \ln n)^{k+1}}$ by the induction assumption, so condition (B.1) holds. Condition (B.2) holds by Claim 11.14.

To verify condition (B.3), let $b \in B_{k+1}$. We show that b has at most $k+1$ neighbors in $X_0 \cup \ldots \cup X_k \cup X_{k+1}$ in the graph G_{k+1}. By Claim 11.15, the adjacency between b and $X_0, \ldots, X_k$ in the graph G_{k+1} is the same as in the graph G_k. Therefore, b has at most k neighbors in $X_0 \cup \ldots \cup X_k$ in the graph G_{k+1}, as it does so in the graph G_k by assumption. Furthermore, b has exactly one neighbor in X_{k+1} by construction. This verifies condition (B.3).

Finally, we verify condition (B.4). Let $P' \subseteq B_{k+1}$ be a $(k+1)$-class. By Claim 11.15 and since every vertex in B_{k+1} has exactly one neighbor in X_{k+1}, we can write $P' = N_{G_{k+1}}(a_0) \cap P \cap B_{k+1}$ for some k-class $P \subseteq B_k$ and some $a_0 \in X_{k+1}$. We need to show that (a) $\mathrm{br}_{G_{k+1}}(P') \leqslant d-k-1$, and that (b) P' does not contain distinct vertices with equal neighborhoods in A_{k+1} in the graph G_{k+1}.

We first verify (a), that is, $\mathrm{br}_{G_{k+1}}(P') < d-k$. For all $b \in P$ we have that $(a_0, b) \in E_0$ if and only if $\mathrm{br}_{G_k}(N_{G_k}(a_0) \cap P) = d-k$. Suppose first that

$$\mathrm{br}_{G_k}(N_{G_k}(a_0) \cap P) < d-k. \tag{11.3}$$

Then $(a_0, b) \notin E_0$ for all $b \in P$. As $E_0 = E(G_k) \triangle E(G_{k+1})$, it follows that the neighborhood of a_0 in P is the same when evaluated in G_k and when evaluated in G_{k+1}. Therefore,

$$N_{G_k}(a_0) \cap P \cap B_{k+1} = N_{G_{k+1}}(a_0) \cap P \cap B_{k+1} = P'.$$

In particular $P' \subseteq N_{G_k}(a_0) \cap P$, so $\mathrm{br}_{G_k}(P') < d-k$ by (11.3) and the monotonicity of the branching index. Then also $\mathrm{br}_{G_{k+1}}(P') < d-k$ by Claim 11.16. This confirms (a) in the considered case. Now suppose that

$$\mathrm{br}_{G_k}(N_{G_k}(a_0) \cap P) \geqslant d-k. \tag{11.4}$$

Then $\mathrm{br}_{G_k}(\overline{N}_{G_k}(a_0) \cap P) < d-k$ holds by (11.1). Dually to the previous case, we have that $(a_0, b) \in E_0$ for all $b \in P$. By a reasoning dual to the one above,

$$\overline{N}_{G_k}(a_0) \cap P \cap B_{k+1} = N_{G_{k+1}}(a_0) \cap P \cap B_{k+1} = P'.$$

Again, we conclude that $\mathrm{br}_{G_{k+1}}(P') < d-k$, confirming (a).

We now verify (b), that is, we show that P' does not contain any pair of distinct vertices with equal neighborhoods in A_{k+1} in the graph G_{k+1}. Let $b, b' \in P'$ be distinct. Then $b, b' \in P$, so by assumption, b and b' have distinct neighborhoods in A_k in the graph G_k. Let $a \in A_k$ be such that

$$(a, b) \in E(G_k) \Leftrightarrow (a, b') \notin E(G_k).$$

Since $b, b' \in P$ it follows that

$$(a, b) \in E_0 \Leftrightarrow (a, b') \in E_0.$$

As $E(G_{k+1}) = E(G_k) \triangle E_0$, it follows that $(a, b) \in E(G_{k+1}) \Leftrightarrow (a, b') \notin E(G_{k+1})$. Since a_0 is the unique neighbor of both b and b' in X_{k+1} in the graph G_{k+1}, and a is adjacent in G_{k+1} either to b or b', we conclude $a \notin X_{k+1}$. Therefore, a witnesses that b and b' have distinct neighborhoods in $A_{k+1} = A_k \setminus X_{k+1}$ in the graph G_{k+1}.

Thus, we verified condition (B.4), and completed Case 1.

Case 2: $|\mathbf{B}_-| > |\mathbf{B_k}|/\mathbf{2}$. Let $B_{k+1} := B_- \subseteq B_k$ consist of the vertices with no neighbors in A_k in the graph G_{k+1}. Set $X_{k+1} := \emptyset$. We verify the required conditions for this choice of B_{k+1} and X_{k+1}.

Condition (B.1) holds as $|B_{k+1}| \geqslant |B_k|/2 \geqslant \frac{|B|}{2 \cdot (600 \ln n)^k} \geqslant \frac{|B|}{(600 \ln n)^{k+1}}$, condition (B.2) holds by Claim 11.14, and condition (B.3) holds by Claim 11.15. To prove condition (B.4), we show the following

Claim 11.17. *No two distinct vertices of B_{k+1} have equal neighborhoods in $X_0 \cup \cdots \cup X_k$ in G_{k+1}.*

Proof. Suppose that distinct $b, b' \in B_{k+1}$ have equal neighborhoods in $X_0 \cup \cdots \cup X_k$ in G_{k+1}. Using Claim 11.15, let P be the k-class such that $b, b' \in P$. By condition (B.4).(b), we have that b and b' have different neighborhoods in A_k in G_k. Let $a \in A_k$ be adjacent to exactly one of b, b' in G_k. As b and b' belong to the same k-class P, it follows by definition of E_0 that $ab \in E_0 \Leftrightarrow ab' \in E_0$. As $E(G_k) = E(G_{k+1}) \triangle E_0$, we observe that a is adjacent to exactly one of b, b' in G_{k+1}. Moreover, as $b, b' \in B_-$, and vertices in B_- have no neighbors in A_k in the graph G_k, it follows that $a \notin A_k$, so $a \in X_0 \cup \cdots \cup X_k$. Therefore, b and b' have distinct neighborhoods in $X_0 \cup \cdots \cup X_k$ in G_{k+1}, a contradiction which completes the proof. ■

As $X_0 \cup \cdots \cup X_k = X_0 \cup \cdots \cup X_{k+1}$ and by Claim 11.17, we have that each $(k+1)$-class P in G_{k+1} contains only a single vertex. This means we have $\mathrm{br}_{G_{k+1}}(P) = 0$ which satisfies condition (B.4).(a). Moreover, (B.4).(b) is vacuously true. This completes Case 2, and the proof of the lemma. □

Next, we proceed to Proposition 11.9, which is obtained by setting $k = d$ in Lemma 11.13 and studying the consequences of condition (B.4) in this case. We repeat the statement.

Proposition 11.9. *Fix $d \in \mathbb{N}$. There is a transduction T_d with the following properties. Given a bipartite graph $G = (A, B, E)$ with $n := |A| \geqslant 2$ such that no vertices in B have equal neighborhoods and $\mathrm{br}_G(B) \leqslant d$, there is a bipartite graph $G' = (A', B', E') \in \mathsf{T}_d(G)$ with $A' \subseteq A$ and $B' \subseteq B$ such that:*

(A.1) $|B'| \geqslant \frac{|B|}{(600 \ln n)^d}$,

(A.2) every vertex $b \in B'$ has at most d neighbors in A' in the graph G',

(A.3) all vertices in B' have distinct neighborhoods in A' in the graph G'.

Proof. Apply Lemma 11.13 to $k = d$, obtaining a formula $\varphi_d(x, y)$ involving the edge relation E and some unary predicates. Let T_d be the transduction that first assigns these unary predicates, then applies $\varphi_d(x, y)$, and finally takes an arbitrary subgraph. Given a bipartite graph $G = (A, B, E)$, let $B_0, \ldots, B_d, X_0, \ldots, X_d$ and $G_d = (A, B, E')$ be as in the statement of Lemma 11.13. Set $A' := X_0 \cup \cdots \cup X_d$, $B' := B_d$. Let $G' := G_d[A', B']$ be the bipartite graph induced on A' and B'. Then $G' \in \mathsf{T}_d(G)$.

The conditions (A.1) and (A.2) follow immediately from Lemma 11.13. We verify condition (A.3). Suppose $b, b' \in B'$ have equal neighborhoods in A' in the graph G'. By condition (B.4).(a) in Lemma 11.13 applied to $P = \{b, b'\}$, we have that $\mathrm{br}_{G_d}(P) = 0$. Hence, b and b' have equal neighborhoods in G_d. By condition (B.4).(b), we get $b = b'$. This completes the proof of Proposition 11.9. □

The next lemma combines Proposition 11.9 with Fact 11.4 and Fact 11.5.

Lemma 11.18. *Fix $\varepsilon > 0$ and let $\mathcal{C}$ be a monadically stable class of bipartite graphs. Then for every $G = (A, B, E) \in \mathcal{C}$ such that no two vertices in B have equal neighborhoods in A, we have that*

$$|B| \leqslant O_{\mathcal{C},\varepsilon}(|A|^{1+\varepsilon}).$$

Proof. Let d be as in Fact 11.7, so that $\mathrm{br}_G(B) \leqslant d$ for all $G = (A, B, E) \in \mathcal{C}$, Let T_d be as in Proposition 11.9. Without loss of generality, we may assume $|A| \geqslant 2$ and that B has no two vertices with equal neighborhoods in A, for all $(A, B, E) \in \mathcal{C}$. We associate with every $G \in \mathcal{C}$ a bipartite graph $F(G) \in \mathsf{T}_d(G)$ satisfying the conditions listed in Proposition 11.9. Let $\mathcal{D} = \{F(G) : G \in \mathcal{C}\}$. Then $\mathcal{D}$ is monadically stable, as $\mathcal{D} \subseteq \mathsf{T}_d(\mathcal{C})$ and $\mathcal{C}$ is monadically stable. Moreover, the class $\mathcal{D}$ avoids $K_{d+1,d+1}$ as a subgraph, by condition (A.2). Therefore, Fact 11.5 implies that $\mathcal{D}$ is nowhere dense. By (A.3), for every graph $(A', B', E') \in \mathcal{D}$, there is no pair of vertices in B' with equal neighborhoods in A'. Consider $G = (A, B, E) \in \mathcal{C}$ and $F(G) = (A', B', E') \in \mathcal{D}$. By Fact 11.4 we have for $\delta := \varepsilon/2$ that

$$|B'| \leqslant O_{\mathcal{D},\delta}\left(|A'|^{1+\delta}\right).$$

On the other hand, by condition (A.1) we infer that

$$|B'| \geqslant |B|/(600 \ln |A|)^d.$$

As $|A'| \leqslant |A|$, we obtain

$$|B| \leqslant |B'| \cdot (600 \ln |A|)^d \leqslant (600 \ln |A|)^d \cdot O_{\mathcal{D},\delta}\left(|A|^{1+\delta}\right) \leqslant O_{\mathcal{C},d,\varepsilon}\left(|A|^{1+\varepsilon}\right).$$ □

Theorem 11.3, restated for convenience, now follows easily.

Theorem 11.3. *Let $\mathcal{C}$ be a monadically stable graph class and $\varepsilon > 0$. Then for every $G \in \mathcal{C}$ and $A \subseteq V(G)$,*

$$|\{N_G[v] \cap A \colon v \in V(G)\}| \leqslant O_{\mathcal{C},\varepsilon}(|A|^{1+\varepsilon}).$$

Proof. Define the following class of bipartite graphs $\mathcal{B}$:

$$\mathcal{B} := \{G[A, B] : G \in \mathcal{C}, A, B \subseteq V(G), A \cap B = \emptyset\}.$$

As $\mathcal{C}$ is monadically stable, it follows that $\mathcal{B}$ is as well. By Lemma 11.18, for all $(A, B, E) \in \mathcal{B}$ such that no two vertices in B have equal neighborhoods we have:

$$|B| \leqslant O_{\mathcal{B},\varepsilon}\left(|A|^{1+\varepsilon}\right).$$

Let $G \in \mathcal{C}$ and $A \subseteq V(G)$. Choose an inclusion-maximal set $B \subseteq V(G) - A$ such that no two vertices of B have equal neighborhoods in A, in the graph G. Then no two vertices of B have equal neighborhoods in A in the bipartite graph $G[A, B] \in \mathcal{B}$. Therefore:

$$\begin{aligned} |\{N[v] \cap A : v \in V(G)\}| &\leqslant |\{N[v] \cap A : v \in A\}| + |\{N[v] \cap A : v \in B\}| \\ &\leqslant |A| + |B| \leqslant |A| + O_{\mathcal{B},\varepsilon}\left(|A|^{1+\varepsilon}\right) \leqslant O_{\mathcal{C},\varepsilon}\left(|A|^{1+\varepsilon}\right). \end{aligned}$$

□

11.2 Neighborhood Covers via Welzl Orders

We first reduce the problem of finding a distance-r cover in a graph G, to finding a distance-1 cover in an interpretation of G. Let the *rth power* of G, denoted G^r, be the graph on the same vertex set as G where vertices u, v are adjacent if and only if the distance between u and v in G is at most r. Note that for every fixed r, G^r can be easily interpreted in G using a formula that checks whether the distance between u and v is at most r. The next lemma shows that finding a distance-1 neighborhood cover in G^r immediately yields a distance-r neighborhood cover in G with the same overlap.

Lemma 11.19. *Let G be a graph, r be a positive integer, and $\mathcal{K}$ be a distance-1 neighborhood cover of G^r of diameter d. Then $\mathcal{K}$ is also a distance-r neighborhood cover of G of diameter at most $d \cdot r$.*

Proof. That $\mathcal{K}$ is a distance-r neighborhood cover of G follows immediately from the observation that for every vertex u, $N_1^{G^r}[u] = N_r^G[u]$. That the weak diameter of $\mathcal{K}$ is at most $d \cdot r$ follows immediately from the triangle inequality and the definition of the graph G^r. □

In this section we provide a construction of neighborhood covers with small overlap for monadically stable graph classes. Formally, we prove the following result.

Theorem 11.2. *There is an algorithm that, given an n-vertex graph G and a radius $r \in \mathbb{N}$, computes a distance-r neighborhood cover of G with diameter at most $4r$ in time $O(n^5)$. For every monadically stable class $\mathcal{C}$ containing G and every $\varepsilon > 0$, the overlap of the cover is bounded by $O_{\mathcal{C},r,\varepsilon}(n^\varepsilon)$.*

We note that the algorithm of Theorem 11.2 does not depend on the class $\mathcal{C}$ or the value of ε: it is a single algorithm that, when supplied with a graph $G \in \mathcal{C}$ and radius $r \in \mathbb{N}$, will always output a neighborhood cover of G with diameter and overlap bounded as asserted.

The main ingredient towards proving Theorem 11.2 will be a tool introduced by Welzl [85] in the context of geometric range queries, called *spanning paths with low crossing number*, which we will call *Welzl orders*. To introduce them, we need some definitions. We remark that for convenience, our terminology slightly differs from that of Welzl.

Consider a set system $\mathcal{S} = (U, \mathcal{F})$, where U is a finite universe and $\mathcal{F}$ is a family of subsets of U. We call the elements of U and $\mathcal{F}$ the *points* and *ranges* in $\mathcal{S}$. The *(primal) shatter function* of $\mathcal{S}$ is the function $\pi_{\mathcal{S}}(\cdot)$ that assigns each positive integer n the value

$$\pi_{\mathcal{S}}(n) := \max_{A \subseteq U, |A| \leqslant n} |\{X \cap A \colon X \in \mathcal{F}\}|.$$

In other words, $\pi_{\mathcal{S}}(n)$ is the largest number of *traces* that the sets from $\mathcal{F}$ leave on a subset $A \subseteq U$ of size n, where the trace left by $X \in \mathcal{F}$ on A is $X \cap A$. For example, the Sauer-Shelah Lemma states that if the VC dimension of $\mathcal{S}$ is d, then $\pi_{\mathcal{S}}(n) \leqslant O(n^d)$. On the other hand, from Theorem 11.3 we immediately obtain the following.

Corollary 11.20. *Let $\mathcal{C}$ be a monadically stable graph class and $\mathcal{D}$ be the class of set systems of closed neighborhoods of graphs in $\mathcal{C}$, that is,*

$$\mathcal{D} := \big\{ \big(V(G), \{N_G[u] \colon u \in V(G)\}\big) \colon G \in \mathcal{C} \big\}.$$

Then for every $\mathcal{S} \in \mathcal{D}$, we have $\pi_{\mathcal{S}}(n) \leqslant O_{\mathcal{C},\varepsilon}(n^{1+\varepsilon})$ for all $\varepsilon > 0$.

Given a set system $\mathcal{S} = \{U, \mathcal{F}\}$, we define its *dual* to be the set system $\mathcal{S}^* := \{\mathcal{F}, U^*\}$ where

$$U^* := \bigcup_{u \in U} R_u \quad \text{with} \quad R_u := \{F \in \mathcal{F} : u \in F\}.$$

The *dual shatter function* $\pi^*_{\mathcal{S}}(\cdot) := \pi_{\mathcal{S}^*}(\cdot)$ of $\mathcal{S}$ is the shatter function of its dual $\mathcal{S}^*$. In general set systems the primal and dual shatter function bound each other by polynomial factors. For set systems arising from neighborhoods in undirected graphs, this link is much stronger, due to the symmetry of the edge relation.

Lemma 11.21. *For every graph G, set system $\mathcal{S} = \big(V(G), \{N_G[u] \colon u \in V(G)\}\big)$, and $n \in \mathbb{N}$, we have*

$$\pi^*_{\mathcal{S}}(n) \leqslant \pi_{\mathcal{S}}(n).$$

We remark that, again for graphs, the reverse direction also holds: $\pi^*_{\mathcal{S}}(\cdot) = \pi_{\mathcal{S}}(\cdot)$. However, we will only use the bound stated above, which we prove below.

Proof. Using the symmetry of the edge relation, for every $u \in V(G)$ and range R_u in $\mathcal{S}^*$ we have

$$R_u = \{N_G[v] : u \in N_G[v], v \in V(G)\} = \{N_G[v] : v \in N_G[u]\}. \tag{11.5}$$

Consider vertices $v_1, \ldots, v_n \in V(G)$ and $u_1, \ldots, u_m \in V(G)$ witnessing that $\pi^*_{\mathcal{S}}(n) \geqslant m$. This means $A^* = \{N_G[v_1], \ldots, N_G[v_n]\}$ is a set of n points in $\mathcal{S}^*$ and $B^* := \{R_{u_1}, \ldots, R_{u_m}\}$ is a set of m ranges in $\mathcal{S}^*$ with pairwise different traces on A^*. By (11.5), we have $N_G[u_i] \neq N_G[u_j]$ for all $i \neq j \in [m]$. Then $A := \{v_1, \ldots, v_n\}$ is a set of n points in $\mathcal{S}$ and $B := \{N_G[u_1], \ldots, N_G[u_m]\}$ is a set of m ranges in $\mathcal{S}$.

To prove $\pi_{\mathcal{S}}(n) \geqslant m$, it remains to show that the ranges in B have pairwise different traces in A. Consider two distinct ranges $N_G[u]$ and $N_G[v]$ from B. Then R_u and R_v have pairwise different traces on A^*. This is witnessed by a vertex $d \in A$ such that $N_G[d] \in R_u \triangle R_v$, where $\triangle$ denotes the symmetric difference. By (11.5), we have $d \in N_G[u] \triangle N_G[v]$, which proves that $N_G[u]$ and $N_G[v]$ have different traces on A. □

Next, for a set system $\mathcal{S} = (U, \mathcal{F})$ and a total order $\preccurlyeq$ on U, we define the *crossing number* of $\preccurlyeq$ as follows. For $X \in \mathcal{F}$, the *crossing number* of X with respect to $\preccurlyeq$ is the number of pairs (u, u') of elements of U such that

- u' is the immediate successor of u in $\preccurlyeq$, and
- exactly one of u and u' belongs to X.

Note that this is equivalent to the following: the crossing number of X is the least k such that $\preccurlyeq$ can be partitioned into $k+1$ intervals so that every interval is either contained in or disjoint from X. Then the crossing number of $\preccurlyeq$ is the maximum crossing number of any $X \in \mathcal{F}$ with respect to $\preccurlyeq$. The following statement was proved in [85].

Fact 11.22 (Thm. 4.2 and Lem. 3.3 of [85], see also Thm. 4.3 of [10])**.** *Suppose $\mathcal{S} = (U, \mathcal{F})$ is a set system with dual shatter function $\pi^*_{\mathcal{S}}(n) \leqslant O(n^d)$, where $d > 1$ is a real. Then there exists a total order $\preccurlyeq$ on U with crossing number bounded by $O(|U|^{1-1/d} \cdot \log |U|)$.*

The proofs of Fact 11.22 given in [85] and [10] are constructive, but no precise runtime analysis is given. In Section 11.3, we analyze the construction and give a runtime bound:

Lemma 11.23. *There is an algorithm that given a set system $\mathcal{S} = (V(G), \{N[v] : v \in V(G)\})$ arising from an n-vertex graph G, computes an order as in Fact 11.22 in time $O(n^5)$.*

We note that the construction of Fact 11.22 does not need to be supplied with the value of d: it is a single algorithm that, given $\mathcal{S}$, computes a total order $\preccurlyeq$, and the guarantee on the crossing number of $\preccurlyeq$ follows from the assumption on the growth function of $\mathcal{S}$.

Next, we show that, given a total order with a low crossing number, we can construct a neighborhood cover with a small overlap and constant diameter using a relatively easy greedy construction.

Lemma 11.24. *Suppose $G = (V, E)$ is a graph, $\mathcal{S} := (V, \{N[u] \colon u \in V\})$ is the set system of closed neighborhoods in G, and $\preccurlyeq$ is a total order on V with crossing number k (with respect to $\mathcal{S}$). Then G admits a distance-1 neighborhood cover with diameter at most 4 and overlap at most $k+1$, and such a neighborhood cover can be computed, given G and $\preccurlyeq$, in time $O(|V|^3)$.*

Proof. We need some auxiliary definitions about the order $\preccurlyeq$. An *interval* is a set $I \subseteq V$ that is convex in $\preccurlyeq$: $u \preccurlyeq v \preccurlyeq w$ and $u, w \in I$ entails $v \in I$. A *prefix* of an interval I is an interval J such that $u, v \in I$, $u \preccurlyeq v$ and $v \in J$ entails $u \in J$. An interval I is *compact* if $I \subseteq N[u]$ for some $u \in V$. We perform the following greedy construction of a partition $\mathcal{I}$ of V into intervals:

- Start with $\mathcal{I} := \emptyset$.
- As long as $V \setminus \bigcup \mathcal{I} \neq \emptyset$, let I be the largest prefix of $V \setminus \bigcup \mathcal{I}$ that is compact. Then add I to $\mathcal{I}$.

Thus, $\mathcal{I}$ consists of compact nonempty intervals. A straightforward implementation of the procedure presented above computes $\mathcal{I}$ in time $O(|V|^3)$.

We claim that $\mathcal{K} := \{N[I] \colon I \in \mathcal{I}\}$ is a neighborhood cover of G of diameter at most 4 and overlap at most $k+1$. That $\mathcal{K}$ is a neighborhood cover is clear: if u is a vertex and $I \in \mathcal{I}$ is such that $u \in I$, then $N[u] \subseteq N[I]$. Also observe that the compactness of every $I \in \mathcal{I}$ implies that $N[I]$ has weak diameter at most 4. We are left with proving the claimed bound on the overlap.

Fix any vertex $v \in V$. Call a vertex $u \in V$ a *crossing* for v if u has a successor u' in $\preccurlyeq$ and $N[v]$ contains exactly one of the vertices u and u'. Since $\mathcal{S}$ has crossing number k, there are at most k distinct crossings for v. We claim the following:

$$\text{for every interval } I \in \mathcal{I} \text{ such that } v \in N[I],$$
$$I \text{ contains a crossing for } v \text{ or } I \text{ contains the } \preccurlyeq\text{-largest element of } V.$$

Since there can be at most k crossings for v and at most one interval can contain the $\preccurlyeq$-largest element, this claim will conclude the proof: it implies that v belongs to at most $k+1$ clusters of $\mathcal{K}$.

To show the claim, first note that $v \in N[I]$ implies that $I \cap N[v] \neq \emptyset$. If also $I \setminus N[v] \neq \emptyset$ then clearly I contains a crossing for v, so assume otherwise: $I \subseteq N[v]$. Let u be the largest element of I in the $\preccurlyeq$ order. Unless u is actually the $\preccurlyeq$-largest element of V, u has a successor u' and $u' \in I' \in \mathcal{I}$ for some $I' \neq I$. Also, unless u itself is a crossing for v, we have $u' \in N[v]$. We now observe that $I \cup \{u'\}$ is an interval that is compact, as witnessed by the vertex v, and is strictly larger than I. This contradicts the construction of $\mathcal{I}$: in the round when I was added to $\mathcal{I}$, we could have added the larger interval $I \cup \{u'\}$ instead. This concludes the proof of the claim and of the lemma. □

We may now combine all the gathered tools and prove Theorem 11.2.

Proof of Theorem 11.2. Let $G \in \mathcal{C}$ be the input graph. By Lemma 11.19, it suffices to compute a distance-1 neighborhood cover with diameter 4 for the rth power $G^r = (V, E) \in \mathcal{C}^r$ of G, where $\mathcal{C}^r$ is the class that contains the rth power of every graph in $\mathcal{C}$. As $\mathcal{C}$ interprets $\mathcal{C}^r$, the latter class is still monadically stable. The graph G^r can be computed in time $O(r \cdot |V|^2)$ from G, and we can assume $r \leqslant |V|$ without loss of generality. Let $\mathcal{S} := (V, \{N[u] \colon u \in V\})$ be the set system of closed neighborhoods in G^r. By Lemma 11.21 and Corollary 11.20, we have $\pi^*_{\mathcal{S}}(n) \leqslant \pi_{\mathcal{S}}(n) \leqslant O_{\mathcal{C},r,\varepsilon}(n^{1+\varepsilon})$ for every $\varepsilon > 0$. Apply the algorithm of Fact 11.22 to $\mathcal{S}$, to obtain a total order $\preccurlyeq$ on V such that the crossing number of $\preccurlyeq$ is bounded by $O_{\mathcal{C},r,\varepsilon}(|V|^{1-\frac{1}{1+\varepsilon}} \cdot \log |V|) \leqslant O_{\mathcal{C},r,\varepsilon}(|V|^{\varepsilon})$ for every $\varepsilon > 0$. By Lemma 11.23, this application takes time $O(|V|^5)$. It now suffices to apply the algorithm of Lemma 11.24 to G^r and $\preccurlyeq$. □

11.3 Computing Welzl Orders

Fact 11.22 (Thm. 4.2 and Lem. 3.3 of [85], see also Thm. 4.3 of [10])**.** *Suppose $\mathcal{S} = (U, \mathcal{F})$ is a set system with dual shatter function $\pi^*_{\mathcal{S}}(n) \leqslant O(n^d)$, where $d > 1$ is a real. Then there exists a total order $\preccurlyeq$ on U with crossing number bounded by $O(|U|^{1-1/d} \cdot \log |U|)$.*

Lemma 11.23. *There is an algorithm that given a set system $\mathcal{S} = (V(G), \{N[v] : v \in V(G)\})$ arising from an n-vertex graph G, computes an order as in Fact 11.22 in time $O(n^5)$.*

Proof. Fix a set system $\mathcal{S} = (U = V(G), \mathcal{F} = \{N[v] : v \in V(G)\})$ arising from an n-vertex graph G. We summarize the construction given in [10, Thm. 4.3] and analyze its running time. The source uses slightly different notation. The goal is to produce a *spanning path* with low *crossing number*. A spanning path for a subset $X \subseteq U$ is a symmetric, irreflexiv relation $P \subseteq X^2$ such that the graph with vertices X and edges P is a path. A range $R \in \mathcal{F}$ *crosses* an edge $(u, v) \in P$ if $|\{u, v\} \cap R| = 1$. Given a spanning path P and a range $R \in \mathcal{F}$, we denote by $\mathrm{cn}(R)$ the number of edges crossed by R in P. The *crossing number* of P is defined as $\max_{R \in \mathcal{F}} \mathrm{cn}(R)$. It is easy to see that a spanning path for X with low crossing number naturally corresponds to a total order of X with low crossing number, as defined in Chapter 11. We analogously say that P is a *spanning tree (forest)* for X if the graph (X, P) forms a tree (forest).

The first observation is that any spanning tree T for a set X can be converted into a spanning path P for the same set, whose crossing number is at most twice that of T ([10, Lem. 3.1]). P is obtained from T by a simple depth-first traversal that can be executed in time $O(n)$. This reduces the problem to computing a spanning tree with low crossing number.

The spanning tree is constructed in [10, Thm. 4.3]. The edges of the spanning tree are inserted iteratively. In each iteration a subalgorithm computes $\frac{1}{2}|U|$ edges that form a forest. The edges are inserted between the elements of U. Then U is pruned by keeping only one vertex of each tree of the forest and the next iteration starts. After $O(\log(n))$ rounds a complete spanning tree has been constructed.

We next describe the subalgorithm [10, Lem. 4.2] that constructs a forest. This algorithm runs on a system $(U_1, \mathcal{F}_1)$ where $U_1 \subseteq U$ and $\mathcal{F}_1$ contains the ranges from $\mathcal{F}$, each augmented with a weight that is initially set to 1. The algorithm starts with an edgeless forest. In each step of the algorithm, we pick an edge $(u, v) \in U_i^2$ where the sum of the weights of the ranges from $\mathcal{F}_i$ crossing it is minimal. We add (u, v) to the forest, obtain U_{i+1} from U_i by removing either u or v, and obtain $\mathcal{F}_{i+1}$ from $\mathcal{F}_i$ by doubling the weights of all the ranges that cross (u, v). The procedure is repeated until the forest contains at least $\frac{1}{2}|U_1|$ edges. This finishes the description of the algorithm to build a spanning path with low crossing number.

Let us now bound the running time of the algorithm. During the ith iteration of the subalgorithm [10, Lem. 4.2], we maintain a data structure in which

- each range in $\mathcal{F}_i$ is linked to all edges from U_i^2 it crosses, and
- each edge from U_i^2 stores the summed weight of all the ranges that cross it.

Recall that both U_i and $\mathcal{F}_i$ have size at most n. Assuming we can test adjacency between two vertices of G in constant time, an initial data structure for U_1 and F_1 can be constructed in time $O(n^3)$. Now in a single iteration:

1. We sort the edges from U_i^2 by their summed weights to find an edge e_i with minimum summed weight. This requires $O(n^2 \cdot \log(n))$ weight comparisons.
2. We add e_i to the forest, remove it from U_{i+1}, double the weight of every range R that crosses e_i and update the summed weights of each edge that R crosses. This requires a total of $O(n^3)$ weight additions.

Due to the doubling, the weights can get quite large. However since the number of iterations is bounded by n, the weights can be represented by bit arrays of length $O(n)$, on which comparisons and summations can be carried out in time $O(n)$. The running time to insert a single edge can therefore be bounded by $O(n^4)$. Each inserted edge becomes an edge of the spanning tree, which has a total of $n - 1$ edges. We can therefore bound the running time of all the calls to the subalgorithm by $O(n^5)$. This dominates the total running time. □

We remark that the running time of the algorithm could be improved, for example by introducing a rounding mechanism to avoid large weights (see e.g. [59, Sec. 5]).

Chapter 12

Model Checking

In this chapter we present the model checking algorithm for monadically stable graph classes.

12.1 Preliminaries

We introduce some additional definitions.

Unary expansions. In this chapter, we allow graphs to be expanded by unary predicates. More precisely, a graph G is as relational structure with a finite universe $V(G)$ over a finite signature Σ consisting of the binary, irreflexive edge relation E and a finite number of unary predicates. Most of the time, the signature Σ will be clear from the context and we will not mention it explicitly. For a graph G and a unary predicate P in the signature of G, we say that P is *interpreted* by $\{v \in V(G) : G \models P(v)\}$. This definition of a graph is similar to the definition of a coloring of a graph, that was used in previous chapters. However, we do not require that the unary predicates from Σ partition the vertex set: Instead of being assigned exactly one color, a vertex may satisfy multiple unary predicates (or also none at all). In order to not confuse the two notions, this chapter will only use the above definition of a graph and never use colorings.

We will commonly construct *unary expansions*, that is, we expand graphs with additional unary predicates. We introduce some convenient notation for this task. For a graph G over the signature Σ and a subset of its vertices $W \subseteq V(G)$, we write $G\langle X \mapsto W\rangle$ for the graph G over the signature $\Sigma \cup \{X\}$ where the predicate X is interpreted as W. We write $G\langle W\rangle$ as a shorthand for $G\langle W \mapsto W\rangle$, where (by slight abuse of notation) we identify a relation symbol with its interpretation. For a family $\mathcal{U} = \{U_1, \dots, U_t\}$ of subsets of $V(G)$, we write $G\langle \mathcal{U}\rangle$ or $G\langle U_1, \dots, U_t\rangle$ as a shorthand for $G\langle U_1\rangle \dots \langle U_t\rangle$.

Types. Let G be a graph and $\bar{a} \in V(G)^{|\bar{a}|}$ be a tuple in G. We denote by $\mathrm{tp}_q(G, \bar{a})$ the finite set of all normalized formulas $\varphi(\bar{x})$ with $|\bar{x}| = |\bar{a}|$ and quantifier rank at most q over the signature of G such that $G \models \varphi(\bar{a})$. We write $\mathrm{tp}_q(G) := \mathrm{tp}_q(G, \emptyset)$ for the set of all normalized sentences of quantifier rank at most q that hold in G.

12.2 Guarded Formulas and Local Types

12.2.1 Guarded Formulas

Given a set of unary predicates $\mathcal{U}$, we say a formula is *$\mathcal{U}$-guarded* if every quantifier is of the form $\exists x \in U$ or $\forall x \in U$ for some $U \in \mathcal{U}$. Our model checking algorithm crucially builds on the simple observation that when evaluating guarded sentences, we can ignore all vertices outside the guarding sets.

Observation 12.1. Given a graph G and a family $\mathcal{U} = \{U_1, \ldots, U_t\}$ of subsets of $V(G)$. Interpreting each set from $\mathcal{U}$ as a unary predicate, we have for every $\mathcal{U}$-guarded sentence φ that

$$G\langle U_1, \ldots, U_t\rangle \models \varphi \quad \Leftrightarrow \quad G\langle U_1, \ldots, U_t\rangle[U_1 \cup \ldots \cup U_t] \models \varphi.$$

Our goal is to compute a representative set of guards $\mathcal{U} = \{U_1, \ldots, U_t\}$ such that we can translate our input formula φ into an equivalent $\mathcal{U}$-guarded formula. Here, crucially, the size t of $\mathcal{U}$ shall depend only on φ (and eventually the depth of the recursion on $\mathcal{C}$). Assume for now that we have recursively computed a large set of candidate guards $\{V_1, \ldots, V_m\}$. Then the selection of the set $\mathcal{U}$ is based on the following key theorem that we prove in the remainder of this section.

Theorem 12.14. *Let G be a graph and let $A, B \subseteq V(G)$ be vertex sets such that* $\mathrm{dist}(A, B) > 2^k$ *and* $\mathrm{tp}_k(G\langle X \mapsto A\rangle\, [N_{2^{k-1}-1}[A]]) = \mathrm{tp}_k(G\langle X \mapsto B\rangle\, [N_{2^{k-1}-1}[B]])$*. Let $\bar{w} \in V(G)^{|\bar{y}|}$ be vertices with* $\mathrm{dist}(\bar{w}, A \cup B) \geqslant 2^k$*. Then for every formula $\varphi(\bar{y}, x)$ of quantifier rank at most $k-1$ in the signature of G we have $G\langle A\rangle \models \exists x \in A\ \varphi(\bar{w}, x) \Leftrightarrow G\langle B\rangle \models \exists x \in B\ \varphi(\bar{w}, x)$.*

Intuitively, the theorem states the following. Given a graph G and two sets A and B whose neighborhoods look alike and which are far from each other. If we find an element $a \in A$ satisfying a first-order property $\varphi(x)$, then we also find an element $b \in B$ satisfying the same property, which will allow us to restrict quantification to appropriately chosen guard sets.

Note that the locality radius in the theorem naturally corresponds to distances that can be expressed with $k-1$ quantifiers. The proof of the theorem is based on the notion of *local types*, which were introduced in [40]. Local types over graphs capture the locality properties of first-order logic by identifying the semantic restriction to 2^{k-1}-neighborhoods with the ability of first-order logic to syntactically make these restrictions. We remark that the results of this section do *not* directly translate to structures with relations of arity greater than 2, since defining distances in (the Gaifman graph of) such structures may require the use of additional quantifiers. Some of the results we prove here were proved in a different notation already in [40] and in the lecture notes of Szymon Toruńczyk [82], while some lemmas and in particular the main theorem of this section, Theorem 12.14, is new. We provide all proofs for consistency and completeness.

Our proof of Theorem 12.14 proceeds as follows. In Section 12.2.2 we recall the notion of Ehrenfeucht-Fraïssé games (short EF-games), which are a classical tool of finite model theory to understand the expressiveness of first-order logic. We introduce a local variant of the games in Section 12.2.3, where all moves are restricted to the local neighborhoods of elements that were played before. Classically, EF-games can be played on two different structures. In Section 12.2.4 we show that, when playing on the same graph, local games determine global games. We relate local games with local types in Section 12.2.5. Up to this point, most of the results were provided in a similar form already in [40]. Towards the proof of Theorem 12.14 we now extend the framework and incorporate guards into local games in Section 12.2.6.

12.2.2 Games

The EF-game is played by two players called *Spoiler* and *Duplicator* on two structures. It is Spoiler's goal to distinguish the two structures, while Duplicator wants to show that the structures cannot be distinguished. The connection with first-order logic is as follows: Duplicator has a winning strategy in the q-round EF-game on two structures if and only if the two structures satisfy the same sentences of quantifier rank at most q. In this work, we consider only games that are played on a single graph with different distinguished vertices $\bar{a}$ and $\bar{b}$. We refer to the literature for extensive background on EF-games, for example, to the textbook [57].

Each position of the game is a tuple $(G, \bar{a}, \bar{b}, k)$ consisting of a graph G that is fixed throughout the game, two non-empty tuples of vertices $\bar{a}, \bar{b}$ of equal length, and a counter $k \in \mathbb{N}$ that keeps track of the number of rounds that are still to play. The game starts in some position $(G, \bar{a}_0, \bar{b}_0, q)$. If we are currently at a position $(G, \bar{a}, \bar{b}, k)$, one round of the game proceeds as follows.

- Spoiler selects a vertex of G as a_k (he makes an *a-move*) or as b_k (he makes a *b-move*).
- If Spoiler made an a-move, then Duplicator has to reply with a b-move, that is, select a vertex of G as b_k, or if he made a b-move, then she has to reply with an a-move, that is, select a vertex of G as a_k.
- The game continues at position $(G, \bar{a}a_k, \bar{b}b_k, k-1)$.

The game terminates when $k = 0$. Assume that a final position $(G, a_\ell, \ldots, a_1, b_\ell, \ldots, b_1, 0)$ is reached ($\ell = k + |\bar{a}|$). We say this is a *winning position for Duplicator* if $(a_\ell, \ldots, a_1, b_\ell, \ldots, b_1)$ defines a *partial automorphism* on G if for all $1 \leqslant i, j \leqslant \ell$,

- $a_i = a_j \Leftrightarrow b_i = b_j$,
- a_i and b_i satisfy the same unary predicates in G, and
- $(a_i, a_j) \in E(G) \Leftrightarrow (b_i, b_j) \in E(G)$.

In the language of types: $\mathrm{atp}_G(a_\ell \ldots a_1) = \mathrm{atp}_G(b_\ell \ldots b_1)$.

We say that Duplicator has a *winning strategy* from a position if she can play such that she reaches – no matter how Spoiler plays – a winning position for Duplicator. Otherwise, we say Spoiler has a *winning strategy*. We write $(G, \bar{a}) \cong_k (G, \bar{b})$ if Duplicator has a winning strategy from position $(G, \bar{a}, \bar{b}, k)$. The proof of the following classical result can be found, for example, in [57, Theorem 3.9].

Lemma 12.2. *$(G, \bar{a}) \cong_k (G, \bar{b})$ if and only if* $\mathrm{tp}_k(G, \bar{a}) = \mathrm{tp}_k(G, \bar{b})$.

12.2.3 Local Games

It is well known that first-order logic can express only local properties of graphs. In particular, for every k and $d \leqslant 2^k$ there exists a formula of quantifier rank k that can determine if the distance between two elements is exactly d, while there is no formula with k quantifiers that can distinguish between distances strictly greater than 2^k. This fact motivates our next definition of *local games*. The key observation is that in a position $(G, \bar{a}, \bar{b}, k)$ when an element a_k at distance at most 2^{k-1} from $\bar{a}$ is chosen by Spoiler, then Duplicator must respond with an element b_k at the exactly same distance to $\bar{b}$, (and vice versa) as otherwise Spoiler can change his strategy to simply point out the difference in distances. On the other hand, these locality properties imply that Spoiler will never select an element at distance greater than 2^{k-1} from both $\bar{a}$ and $\bar{b}$, as this element could simply be copied by Duplicator.

At a position $(G, \bar{a}, \bar{b}, k)$, we say that a move (by Spoiler or Duplicator) is *local* if it is an a-move and contained in $N_{2^{k-1}}[\bar{a}]$ or if it is a b-move and contained in $N_{2^{k-1}}[\bar{b}]$. We define the *local EF-game* as the EF-game where we require that both players are only allowed to play local moves. We call the regular EF-game *global* to distinguish it from the local game. We write $(G, \bar{a}) \cong_k^{\mathrm{local}} (G, \bar{b})$ if Duplicator has a winning strategy for the local game from position $(G, \bar{a}, \bar{b}, k)$.

12.2.4 Local Games Determine Global Games

We will argue that the local and global EF-games are equivalent when we start from positions $(G, \bar{a}, \bar{b}, k)$ such that $\bar{a}$ and $\bar{b}$ are at distance greater than 2^{k+1}. Towards this goal, we formally prove the above observations. First, we observe that Duplicator has to respond to a local move of Spoiler with her own local move.

Lemma 12.3 (see Lemma 9.2 of [82]). *Consider the global game at a position $(G, \bar{a}, \bar{b}, k)$. Assume Spoiler made a local a-move $a_k \in N_{2^{k-1}}[\bar{a}]$, say $a_k \in N_{2^{k-1}}[a_j]$ for $a_j \in \bar{a}$ and Duplicator answers with a b-move $b_k \notin N_{2^{k-1}}[\bar{b}_j]$, or symmetrically, Spoiler made a local b-move $b_k \in N_{2^{k-1}}[\bar{b}]$, say $b_k \in N_{2^{k-1}}[b_j]$ for some $b_j \in \bar{b}$ and Duplicator answers with an a-move $a_k \notin N_{2^{k-1}}[\bar{a}_j]$. Then Spoiler has a winning strategy for the remaining global game from position $(G, \bar{a}a_k, \bar{b}b_k, k-1)$. In particular, if Duplicator answers with a non-local move to a local move, she loses the game.*

Proof. We prove the statement by induction on k. By symmetry, we may assume that Spoiler makes an a-move. For $k = 1$ the claim is true, as $a_1 \in N_1[a_j] \Leftrightarrow b_1 \in N_1[b_j]$ is necessary for $(\bar{a}a_1, \bar{b}b_1)$ to be a partial automorphism.

Now assume $k > 1$. As $a_k \in N_{2^{k-1}}[a_j]$, Spoiler can play $a_{k-1} \in N_{2^{k-2}}[a_k] \cap N_{2^{k-2}}[a_j]$ as his next move. As $b_k \notin N_{2^{k-1}}[b_j]$, no matter which b_{k-1} Duplicator plays as a response, either $b_{k-1} \notin N_{2^{k-2}}[b_k]$ or $b_{k-1} \notin N_{2^{k-2}}[b_j]$. If $b_{k-1} \notin N_{2^{k-2}}[b_k]$, then by induction hypothesis applied to position $(G, a_k, b_k, k-1)$, Spoiler wins from position $(G, a_k a_{k-1}, b_k b_{k-1}, k-2)$. If $b_{k-1} \notin N_{2^{k-2}}[b_j]$, then by induction hypothesis applied to position $(G, a_j, b_j, k-1)$, Spoiler wins from position $(G, a_j a_{k-1}, b_j b_{k-1}, k-2)$. Since adding more preselected vertices only helps Spoiler, he would win in particular the remaining game from position $(G, \bar{a}a_k a_{k-1}, \bar{b}b_k, \bar{b}_{k-1}, k-2)$. □

Lemma 12.4 (see Lemma 3.5 of [40]). *Consider tuples of vertices $\bar{a}, \bar{a}'$, $\bar{b}, \bar{b}'$ in a graph G such that $\operatorname{dist}(\bar{a}, \bar{a}') > 2^k$ and $\operatorname{dist}(\bar{b}, \bar{b}') > 2^k$. Then $(G, \bar{a}\bar{a}') \cong_k^{\text{local}} (G, \bar{b}\bar{b}')$ if and only if both $(G, \bar{a}) \cong_k^{\text{local}} (G, \bar{b})$ and $(G, \bar{a}') \cong_k^{\text{local}} (G, \bar{b}')$.*

Proof. We prove the statement by induction on k. For $k = 0$, observe that

$$\operatorname{dist}(\bar{a}, \bar{a}'), \operatorname{dist}(\bar{b}, \bar{b}') > 2^0 = 1,$$

and thus there are no edges between $\bar{a}$ and $\bar{a}'$ or between $\bar{b}$ and $\bar{b}'$ in G. This means $(\bar{a}\bar{a}', \bar{b}\bar{b}')$ is a partial automorphism if and only if both $(\bar{a}, \bar{b})$ and $(\bar{a}', \bar{b}')$ are partial automorphisms. This proves the statement for $k = 0$. Next, assume the statement holds for $k - 1$ and we will prove it for k.

Assume $(G, \bar{a}) \cong_k^{\text{local}} (G, \bar{b})$ and $(G, \bar{a}') \cong_k^{\text{local}} (G, \bar{b}')$. Note that in particular $(G, \bar{a}') \cong_{k-1}^{\text{local}} (G, \bar{b}')$. We consider the local game at position $(G, \bar{a}\bar{a}', \bar{b}\bar{b}', k)$ and show that Duplicator has a winning strategy. By symmetry, without loss of generality, Spoiler starts with an a-move $a_k \in N_{2^{k-1}}[\bar{a}]$. Duplicator responds according to the winning strategy for the local game at position $(G, \bar{a}, \bar{b}, k)$ yielding $b_k \in N_{2^{k-1}}[\bar{b}]$ such that $(G, \bar{a}a_k) \cong_{k-1}^{\text{local}} (G, \bar{b}b_k)$. By assumption, $\operatorname{dist}(\bar{a}, \bar{a}') > 2^k$ and $\operatorname{dist}(\bar{b}, \bar{b}') > 2^k$, and thus $\operatorname{dist}(\bar{a}a_k, \bar{a}') > 2^{k-1}$ and $\operatorname{dist}(\bar{b}b_k, \bar{b}') > 2^{k-1}$. By induction, since $(G, \bar{a}a_k) \cong_{k-1}^{\text{local}} (G, \bar{b}b_k)$ and $(G, \bar{a}') \cong_{k-1}^{\text{local}} (G, \bar{b}')$, we have $(G, \bar{a}a_k\bar{a}') \cong_{k-1}^{\text{local}} (G, \bar{b}b_k\bar{b}')$. Since we made no assumptions on Spoiler's local move, this implies $(G, \bar{a}\bar{a}') \cong_k^{\text{local}} (G, \bar{b}\bar{b}')$.

Conversely, assume $(G, \bar{a}) \not\cong_k^{\text{local}} (G, \bar{b})$ or $(G, \bar{a}') \not\cong_k^{\text{local}} (G, \bar{b}')$. Without loss of generality, $(G, \bar{a}) \not\cong_k^{\text{local}} (G, \bar{b})$. We consider the local game at position $(G, \bar{a}\bar{a}', \bar{b}\bar{b}', k)$. Spoiler chooses $a_k \in N_{2^{k-1}}[\bar{a}]$ according to his winning strategy at position $(G, \bar{a}, \bar{b}, k)$. By Lemma 12.3, Duplicator responds with $b_k \in N_{2^{k-1}}[\bar{b}]$, which is a valid turn in the local game on position $(G, \bar{a}, \bar{b}, k)$. As Spoiler played according to his winning strategy on that position, we have $(G, \bar{a}a_k) \not\cong_{k-1}^{\text{local}} (G, \bar{b}b_k)$. Again we have $\operatorname{dist}(\bar{a}a_k, \bar{a}') > 2^{k-1}$ and $\operatorname{dist}(\bar{b}b_k, \bar{b}') > 2^{k-1}$ and, by induction, we have $(G, \bar{a}a_k\bar{a}') \not\cong_{k-1}^{\text{local}} (G, \bar{b}b_k\bar{b}')$. Since we made no assumptions on Duplicator's local move this implies $(G, \bar{a}\bar{a}') \not\cong_k^{\text{local}} (G, \bar{b}\bar{b}')$. □

Theorem 12.5 (see Lemma 9.4 of [82]). *Consider a graph G with tuples $\bar{a}$, $\bar{b}$ such that $\operatorname{dist}(\bar{a}, \bar{b}) > 2^{k+1}$. Then*

$$(G, \bar{a}) \cong_k (G, \bar{b}) \quad \Leftrightarrow \quad (G, \bar{a}) \cong_k^{\text{local}} (G, \bar{b}).$$

Proof. The forward direction is easy. Duplicator's winning strategy for the global game when Spoiler makes only local moves is also a winning strategy for the local game, since by Lemma 12.3 her winning strategy responds locally to local moves.

We prove the backward direction by induction on k. For $k = 0$ the global and local game are the same, hence the statement is true. Assume it holds for $k-1$ and we will prove that it also holds for k. We first prove the following claim.

Claim 12.6. *If Duplicator a winning strategy for the game from position $(G, \bar{a}, \bar{b}, k)$ where the first round is local and the remaining rounds are global, then she also has a winning strategy for the global game from position $(G, \bar{a}, \bar{b}, k)$.*

Proof. We will give a winning strategy for Duplicator for the global game at position $(G, \bar{a}, \bar{b}, k)$. Without loss of generality, we can assume Spoiler starts the game with an a-move. If Spoiler opens with a local move, then Duplicator can respond according to her given first-local-then-global winning strategy for the position and win the game.

Thus, we may assume that Spoiler opens with a non-local move $a_k \notin N_{2^{k-1}}[\bar{a}]$. We start by arguing that there exists an element $b_k \notin N_{2^{k-1}}[\bar{b}]$ with $(G, a_k) \cong^{\text{local}}_{k-1} (G, b_k)$: if $a_k \notin N_{2^{k-1}}[\bar{b}]$, then we can choose $b_k = a_k$ and are done. Thus assume $a_k \in N_{2^{k-1}}[\bar{b}]$. We will use a role-swapping argument: if Spoiler had played the local element $a_k \in N_{2^{k-1}}[\bar{b}]$ as a b-move (under the name b_k), then Duplicator's winning strategy of the first-local-then-global game would have replied by Lemma 12.3 with an element $b_k \in N_{2^{k-1}}[\bar{a}]$ as an a-move (under the name a_k). Duplicator would win the remaining $(k-1)$-round global game, or in other words $(G, \bar{b}a_k) \cong_{k-1} (G, \bar{a}b_k)$. By the forward direction of the theorem, which was already proved above, also $(G, \bar{b}a_k) \cong^{\text{local}}_{k-1} (G, \bar{a}b_k)$. In particular, we also have $(G, a_k) \cong^{\text{local}}_{k-1} (G, b_k)$. Since $b_k \in N_{2^{k-1}}[\bar{a}]$ and $\text{dist}(\bar{a}, \bar{b}) > 2^{k+1} > 2^k$ we have $b_k \notin N_{2^{k-1}}[\bar{b}]$. This yields the desired $b_k \notin N_{2^{k-1}}[\bar{b}]$ with $(G, a_k) \cong^{\text{local}}_{k-1} (G, b_k)$.

Duplicator now responds to Spoiler's move a_k with b_k. Since Duplicator has a k-round winning strategy for the game with preselected tuples $\bar{a}, \bar{b}$ where the first round is local and the remaining rounds are global, she in particular has a $(k-1)$-round winning strategy for these tuples, that is, $(G, \bar{a}) \cong_{k-1} (G, \bar{b})$, and hence $(G, \bar{a}) \cong^{\text{local}}_{k-1} (G, \bar{b})$ by the forward direction of the theorem. Since $a_k \notin N_{2^{k-1}}[\bar{a}]$, $b_k \notin N_{2^{k-1}}[\bar{b}]$, we have $\text{dist}(\bar{a}, a_k) > 2^{k-1}$ and $\text{dist}(\bar{b}, b_k) > 2^{k-1}$. We can thus apply Lemma 12.4 for $k-1$ to the tuples $\bar{a}, a_k, \bar{b}, b_k$: Since $(G, a_k) \cong^{\text{local}}_{k-1} (G, b_k)$ and $(G, \bar{a}) \cong^{\text{local}}_{k-1} (G, \bar{b})$, it follows that $(G, \bar{a}a_k) \cong^{\text{local}}_{k-1} (G, \bar{b}b_k)$. By induction hypothesis we have $(G, \bar{a}a_k) \cong_{k-1} (G, \bar{b}b_k)$. Since we made no assumptions on Spoiler's first move, Duplicator has a winning strategy in the global game from position $(G, \bar{a}, \bar{b}, k)$. ■

We are ready to prove the backwards direction of the statement. Assume $(G, \bar{a}) \cong^{\text{local}}_{k} (G, \bar{b})$. By the previous claim it suffices to show that Duplicator has a winning strategy for the game from position $(G, \bar{a}, \bar{b}, k)$ where the first round is local and the remaining rounds are global. Hence, let $a_k \in N_{2^{k-1}}[\bar{a}]$ be a local a-move of Spoiler. We let $b_k \in N_{2^{k-1}}[\bar{b}]$ be Duplicator's response that she would play as a winning move in the local game, that is, we have $(G, \bar{a}a_k) \cong^{\text{local}}_{k-1} (G, \bar{b}b_k)$.

Since $\text{dist}(\bar{a}, \bar{b}) > 2^{k+1}$, and $a_k \in N_{2^{k-1}}[\bar{a}]$, $b_k \in N_{2^{k-1}}[\bar{b}]$, it follows that $\text{dist}(\bar{a}a_k, \bar{b}b_k) > 2^k$ (see Figure 12.1).

By induction hypothesis we have $(G, \bar{a}a_k) \cong_{k-1} (G, \bar{b}b_k)$. As the first move was an arbitrary local a-move, this yields a winning strategy of Duplicator for the game at position $(G, \bar{a}, \bar{b}, k)$ where the first round is local and the remaining rounds are global. By the previous claim, $(G, \bar{a}) \cong_k (G, \bar{b})$. □

While we do not prove, whether the distance requirement $\text{dist}(\bar{a}, \bar{b}) > 2^{k+1}$ in Theorem 12.5 is tight, the example in Figure 12.2 illustrates that some form of distance requirement is necessary for the theorem to hold. Let $k = 1$. Then $(G, a) \cong^{\text{local}}_{k} (G, b)$. However, $(G, a) \not\cong_k (G, b)$ as in

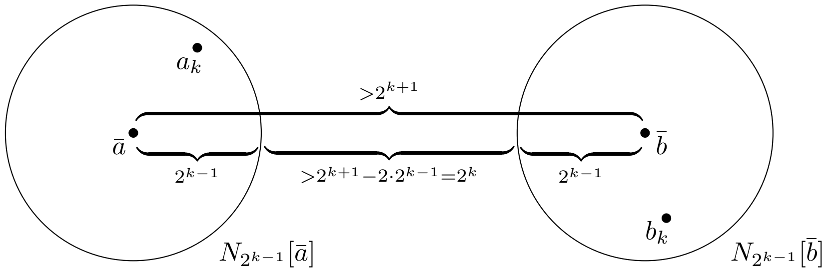

Figure 12.1: Bounding $\mathrm{dist}(\bar{a}a_k, \bar{b}b_k)$ from below.

the global game, Spoiler can choose the uppermost red element as a global b-move. Duplicator cannot reply with a red element that is not adjacent to a, and hence loses the game.

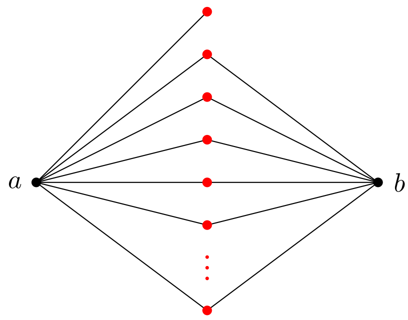

Figure 12.2: An example illustrating the need for a distance constraint in Theorem 12.5.

12.2.5 Local Games and Local Types

We now establish the connection between local games and local first-order logic. Unlike in Gaifman's Locality Theorem, we do not increase the quantifier rank when localizing formulas.

Let G be a graph and $\bar{a}$ be a tuple of vertices of G. It is well known that $\mathrm{tp}_q(G, \bar{a}) = \mathrm{tp}_q(G, \bar{b})$ if and only if $(G, \bar{a}) \cong_q (G, \bar{b})$. The *localization* of a formula φ with free variables is the formula with the same free variables as φ that replaces every subformula $\exists x\ \psi(x, \bar{y})$ with $\exists x \in N_{2^k}[\bar{y}]\ \psi(x, \bar{y})$ (or more precisely $\exists x\ x \in N_{2^k}[\bar{y}] \land \psi(x, \bar{y})$), where k is the quantifier rank of ψ. Likewise, every subformula $\forall x\ \psi(x, \bar{y})$ is replaced with $\forall x \in N_{2^k}[\bar{y}]\ \psi(x, \bar{y})$ (or more precisely $\forall x\ x \in N_{2^k}[\bar{y}] \rightarrow \psi(x, \bar{y})$).

We call a formula *local* if it is the localization of some formula. As the following lemma shows, localizing a formula does not change its quantifier rank.

Lemma 12.7. *There exists a formula with quantifier rank k and free variables $x\bar{y}$ expressing that $x \in N_{2^k}[\bar{y}]$.*

Proof. We can check whether $x \in N_{2^0}[\bar{y}]$ using the quantifier-free formula

$$\bigvee_{y \in \bar{y}} E(x, y) \lor x = y.$$

For $k > 0$, we note that $x \in N_{2^k}[\bar{y}]$ if and only if $\exists z (z \in N_{2^{k-1}}[x] \land z \in N_{2^{k-1}}[\bar{y}])$. □

Let G be a graph and $\bar{a} \in V(G)^{|\bar{a}|}$ be a tuple in G. We partition the finite set of all (normalized) local formulas $\varphi(\bar{x})$ with $|\bar{x}| = |\bar{a}|$ and quantifier rank at most q over the signature of G into the sets $\mathrm{tp}_q^{\mathrm{local}}(G,\bar{a})$ and $\overline{\mathrm{tp}}_q^{\mathrm{local}}(G,\bar{a})$ such that $\varphi(\bar{x}) \in \mathrm{tp}_q^{\mathrm{local}}(G,\bar{a})$ if and only if $G \models \varphi(\bar{a})$ and conversely $\varphi(\bar{x}) \in \overline{\mathrm{tp}}_q^{\mathrm{local}}(G,\bar{a})$ if and only if $G \not\models \varphi(\bar{a})$. We call $\mathrm{tp}_q^{\mathrm{local}}(G,\bar{a})$ the *local q-type of $\bar{a}$ in G*.

We next relate local types and local games.

Lemma 12.8. *If $\mathrm{tp}_k^{\mathrm{local}}(G,\bar{a}) = \mathrm{tp}_k^{\mathrm{local}}(G,\bar{b})$, then $(G,\bar{a}) \cong_k^{\mathrm{local}} (G,\bar{b})$.*

Proof. We prove the claim by induction on k. For $k=0$, $\mathrm{tp}_k^{\mathrm{local}}(G,\bar{a})$ and $\mathrm{tp}_k^{\mathrm{local}}(G,\bar{b})$ are known as the atomic types of $\bar{a}$ and $\bar{b}$. These are equal if and only if the mapping $\bar{a} \mapsto \bar{b}$ is a partial isomorphism, which in turn is equivalent to $(G,\bar{a}) \cong_0^{\mathrm{local}} (G,\bar{b})$.

Let us assume that the statement holds for $k-1$ and show that it also holds for k. Consider the local game at position $(G,\bar{a},\bar{b},k)$. Without loss of generality, Spoiler starts the game by an a-move $a_k \in N_{2^{k-1}}(\bar{a})$. Let

$$\tau(\bar{y},x) = \bigwedge_{\varphi(\bar{y}x)\in \mathrm{tp}_{k-1}^{\mathrm{local}}(G,\bar{a}a_k)} \varphi(\bar{y}x) \wedge \bigwedge_{\varphi(\bar{y}x)\in \overline{\mathrm{tp}}_{k-1}^{\mathrm{local}}(G,\bar{a}a_k)} \neg\varphi(\bar{y}x)$$

be the local formula that exactly captures $\mathrm{tp}_{k-1}^{\mathrm{local}}(G,\bar{a}a_k)$. More precisely, for every tuple $\bar{a}'a_k' \in V(G)^{|a|+1}$ we have $\mathrm{tp}_{k-1}^{\mathrm{local}}(G,\bar{a}'a_k') = \mathrm{tp}_{k-1}^{\mathrm{local}}(G,\bar{a}a_k)$ if and only if $G \models \tau(\bar{a}',a_k')$.

The formula $\psi(\bar{y}) = \exists x \in N_{2^{k-1}}[\bar{y}]\ \tau(\bar{y}x)$ is a local formula with quantifier rank k, and is therefore contained in $\mathrm{tp}_k^{\mathrm{local}}(G,\bar{a})$ as witnessed by instantiating x with a_k. By assumption on equality of local k-types we then also have that $\psi(\bar{y}) \in \mathrm{tp}_k^{\mathrm{local}}(G,\bar{b})$ and hence there exists an element $b_k \in N_{2^{k-1}}(\bar{b})$ such that $\mathrm{tp}_{k-1}^{\mathrm{local}}(G,\bar{b}b_k) = \mathrm{tp}_{k-1}^{\mathrm{local}}(G,\bar{a}a_k)$. Duplicator chooses b_k as her response. The remaining game continues from position $(G,\bar{a}a_k,\bar{b}b_k,k-1)$. Since $\mathrm{tp}_{k-1}^{\mathrm{local}}(G,\bar{b}b_k) = \mathrm{tp}_{k-1}^{\mathrm{local}}(G,\bar{a}a_k)$ Duplicator wins by induction hypothesis. $\square$

It is not difficult to prove that also the converse of the lemma is true, however, we refrain from giving the proof as it is not needed for our further argumentation.

12.2.6 Games and Types with Guards

Theorem 12.5 and Lemma 12.8 together already show that for tuples of sufficiently large distance equality of local types implies equality of global types. We will need a stronger statement for graphs where a specific set of vertices is highlighted. To this end, we introduce special starting positions (G,A,B,k), where $A,B \subseteq V(G)$ are sets of vertices (for the global and local EF-game), which we call *guards*. Spoiler and Duplicator select elements a_k and b_k in the usual way with the constraints $a_k \in A$ and $b_k \in B$, and afterwards the (global or local) game continues at position $(G,a_k,b_k,k-1)$ as usual. Hence, both for the global and the local game, the role of the sets A and B is merely to constrain (guard) the choices for the first round. We write $(G,A) \cong_k (G,B)$ or $(G,A) \cong_k^{\mathrm{local}} (G,B)$ if Duplicator has a winning strategy for the global or local game starting from position (G,A,B,k).

First, we extend Lemma 12.8 to our new starting positions. Note that the following theorem no longer mentions local types, but global types of neighborhoods. Recall that $G\langle X \mapsto W\rangle$ denotes the graph G with the additional unary predicate X interpreted as the vertex set $W \subseteq V(G)$.

Lemma 12.9. *If $\mathrm{tp}_k(G\langle X \mapsto A\rangle\,[N_{2^{k-1}-1}[A]]) = \mathrm{tp}_k(G\langle X \mapsto B\rangle\,[N_{2^{k-1}-1}[B]])$, then*

$$(G,A) \cong_k^{\mathrm{local}} (G,B).$$

Proof. Fix G, A and B with $\mathrm{tp}_k(G\langle X \mapsto A\rangle\,[N_{2^{k-1}-1}[A]]) = \mathrm{tp}_k(G\langle X \mapsto B\rangle\,[N_{2^{k-1}-1}[B]])$. For brevity, let $G_A := G\langle X \mapsto A\rangle$ and $G_B := G\langle X \mapsto B\rangle$. To prove the statement, we need the following observation about local formulas.

Claim 12.10. *Let $\varphi(x)$ be a local formula with quantifier rank at most $k-1$. Then*

$$G_A \models \exists x \in X\ \varphi(x) \text{ if and only if } G_B \models \exists x \in X\ \varphi(x).$$

Proof. Since $G_A\left[N_{2^{k-1}-1}[A]\right]$ and $G_B\left[N_{2^{k-1}-1}[B]\right]$ have the same k-type, they agree in their evaluation of the sentence $\exists x \in X\ \varphi(x)$ with quantifier rank at most k.

Since $\varphi(\bar{x})$ is local and has quantifier rank $k-1$, all the quantified variables in $\varphi(x)$ can only lie within distance at most $\sum_{i=1}^{k-1} 2^{i-1} = \sum_{i=0}^{k-2} 2^i = 2^{k-1} - 1$ from x. Hence, all variables quantified in $\exists x \in X\ \varphi(x)$ must lie within distance at most $2^{k-1} - 1$ from X. Therefore, evaluating it on $G_A\left[N_{2^{k-1}-1}[A]\right]$ and G_A yields the same answer. The same holds for $G_B\left[N_{2^{k-1}-1}[B]\right]$ and G_B. Hence, also G_A and G_B agree in their evaluation of $\exists x \in X\ \varphi(x)$. ■

Consider the local game at position (G, A, B, k). Without loss of generality, Spoiler starts the game by an a-move $a_k \in A$. Let

$$\tau(x) = \bigwedge_{\varphi(x) \in \mathrm{tp}_{k-1}^{\mathrm{local}}(G_A, a_k)} \varphi(x) \wedge \bigwedge_{\varphi(x) \in \overline{\mathrm{tp}}_{k-1}^{\mathrm{local}}(G_A, a_k)} \neg\varphi(x)$$

be the local formula that defines $\mathrm{tp}_{k-1}^{\mathrm{local}}(G_A, a_k)$.

The sentence $\exists x \in X\ \tau(x)$ holds on G_A, as witnessed by instantiating x with a_k. By Claim 12.10, as $\tau(x)$ has quantifier rank $k-1$, $\exists x \in X\ \tau(x)$ is also true on G_B. Hence, there exists an element $b_k \in B$ such that $\mathrm{tp}_{k-1}^{\mathrm{local}}(G_A, a_k) = \mathrm{tp}_{k-1}^{\mathrm{local}}(G_B, b_k)$. Then in particular $\mathrm{tp}_{k-1}^{\mathrm{local}}(G, a_k) = \mathrm{tp}_{k-1}^{\mathrm{local}}(G, b_k)$. Duplicator chooses b_k as his next element. Then by Lemma 12.8 we have $(G, a_k) \cong_{k-1}^{\mathrm{local}} (G, b_k)$. Since we made no assumptions on Spoiler's first move $a_k \in A$, and Duplicator's response always yields a $b_k \in B$, we now have $(G, A) \cong_k^{\mathrm{local}} (G, B)$ as desired. □

Since for starting positions (G, A, B, k), the local and global game allow the same first moves, we get the following simple consequence of Theorem 12.5.

Lemma 12.11. *Consider a graph G with sets $A, B \subseteq V(G)$ such that* $\mathrm{dist}(A, B) > 2^k$. *Then* $(G, A) \cong_k (G, B)$ *if and only if* $(G, A) \cong_k^{\mathrm{local}} (G, B)$.

We can use these new starting positions to determine the truth values of formulas in graphs where the starting sets are highlighted.

Lemma 12.12. *Assume* $(G, A) \cong_k (G, B)$. *Then for every formula $\varphi(x)$ of quantifier rank at most $k-1$ in the signature of G we have* $G\langle A\rangle \models \exists x \in A\ \varphi(x) \Leftrightarrow G\langle B\rangle \models \exists x \in B\ \varphi(x)$.

Proof. Assume $G\langle A\rangle \models \exists x \in A\ \varphi(x)$, that is, there exists $a_k \in A$ with $G \models \varphi(a_k)$. Spoiler chooses $a_k \in A$ and Duplicator responds with $b_k \in B$ such that $(G, a_k) \cong_{k-1} (G, b_k)$. Hence, $\mathrm{tp}_{k-1}(G, a_k) = \mathrm{tp}_{k-1}(G, b_k)$, and in particular $G \models \varphi(b_k)$. We have $G\langle B\rangle \models \exists x \in B\ \varphi(x)$. The converse holds by symmetry. □

We combine Lemma 12.9, Lemma 12.11 and Lemma 12.12 into the following statement.

Lemma 12.13. *Assume* $\mathrm{dist}(A, B) > 2^k$ *and*

$$\mathrm{tp}_k(G\langle X \mapsto A\rangle\left[N_{2^{k-1}-1}[A]\right]) = \mathrm{tp}_k(G\langle X \mapsto B\rangle\left[N_{2^{k-1}-1}[B]\right]).$$

Then for every formula $\varphi(x)$ of quantifier rank at most $k-1$ in the signature of G we have $G\langle A\rangle \models \exists x \in A\ \varphi(x) \Leftrightarrow G\langle B\rangle \models \exists x \in B\ \varphi(x)$.

Again, Figure 12.2 illustrates that the distance constraint is necessary. For $k = 2$, $A = \{a\}$ and $B = \{b\}$, we have that both $G\langle X \mapsto A\rangle\, [N_{2^{k-1}-1}[A]]$ and $G\langle X \mapsto B\rangle\, [N_{2^{k-1}-1}[B]]$ have the same type: both are a star, whose center is marked X and whose leaves are marked red. However, for $\varphi(x) := \forall y\ \mathrm{Red}(y) \to E(x, y)$ we have

$$G\langle A\rangle \models \exists x \in A\ \varphi(x) \text{ and } G\langle B\rangle \not\models \exists x \in B\ \varphi(x).$$

Extending the statement to accommodate for further free variables in φ, will yield Theorem 12.14, which we restate for convenience.

Theorem 12.14. *Let G be a graph and let $A, B \subseteq V(G)$ be vertex sets such that $\mathrm{dist}(A, B) > 2^k$ and $\mathrm{tp}_k(G\langle X \mapsto A\rangle\, [N_{2^{k-1}-1}[A]]) = \mathrm{tp}_k(G\langle X \mapsto B\rangle\, [N_{2^{k-1}-1}[B]])$. Let $\bar{w} \in V(G)^{|\bar{y}|}$ be vertices with $\mathrm{dist}(\bar{w}, A \cup B) \geqslant 2^k$. Then for every formula $\varphi(\bar{y}, x)$ of quantifier rank at most $k - 1$ in the signature of G we have $G\langle A\rangle \models \exists x \in A\ \varphi(\bar{w}, x) \Leftrightarrow G\langle B\rangle \models \exists x \in B\ \varphi(\bar{w}, x)$.*

Proof. Assume $\bar{w} = w_1, \ldots, w_\ell$. We define G' to be the unary expansion of G with 2ℓ new predicates W_i and N_i for $1 \leqslant i \leqslant \ell$. We interpret $W_i = \{w_i\}$ and $N_i = N[w_i] \setminus \{w_i\}$ for $1 \leqslant i \leqslant \ell$. We define $\varphi'(x)$ to be the formula obtained from $\varphi(\bar{y}, x)$ by replacing all atoms $E(w_i, z)$ and $E(z, w_i)$ with $N_i(z)$ and all atoms $(w_i = z)$ and $(z = w_i)$ with $W_i(z)$. Then for every $v \in V(G)$ we have $G \models \varphi(\bar{w}, v) \Leftrightarrow G' \models \varphi'(v)$. Thus, it is sufficient to show $G'\langle A\rangle \models \exists x \in A\ \varphi'(x) \Leftrightarrow G'\langle B\rangle \models \exists x \in B\ \varphi'(x)$. Since $\mathrm{dist}(\bar{w}, A) \geqslant 2^k$, we have $G\, [N_{2^{k-1}-1}[A]] = G'\, [N_{2^{k-1}-1}[A]]$. The same holds for B and thus $\mathrm{tp}_k(G'[N_{2^{k-1}-1}(A)]\langle A \mapsto X\rangle) = \mathrm{tp}_k(G'[N_{2^{k-1}-1}(B)]\langle B \mapsto X\rangle)$. The statement then follows from Lemma 12.13. □

12.3 The Algorithm

In this section, we present our model checking theorem for monadically stable graph classes.

12.3.1 Setup

Let us recall the main result about the Flipper game. See Chapter 10 for definitions.

Theorem 10.4. *There is a budget-2 Flipper strategy $\mathsf{flip}^\star$ with the following property.*
For every monadically stable graph class $\mathcal{C}$ and radius $r \in \mathbb{N}$ there is $\ell \in \mathbb{N}$ such that $\mathsf{flip}^\star$ is ℓ-winning and has runtime $O_{\mathcal{C},r}(n^2)$ in the radius-r game on $\mathcal{C}$.

We will use the strategy $\mathsf{flip}^\star$ for our model checking algorithm. For every monadically stable graph class $\mathcal{C}$ and $r \in \mathbb{N}$, we define game-depth$(\mathcal{C}, r)$ to be the bound on the number of rounds needed for Flipper to win the radius-r budget-2 Flipper game on any graph from $\mathcal{C}$ while following $\mathsf{flip}^\star$. We call a sequence of positions $\mathcal{H} = (G_0, \mathcal{I}_0), \ldots, (G_\ell, \mathcal{I}_\ell)$ a $(\mathcal{C}, \rho)$-*history of length* ℓ, if it is a prefix of the Flipper run $\mathcal{R}(\mathsf{loc}, \mathsf{flip}^\star, G_0)$ for some radius-ρ Localizer strategy loc and some graph $G_0 \in \mathcal{C}$.

Our model checking algorithm will use recursion guided by the strategy $\mathsf{flip}^\star$. Playing the Flipper game for many game-depth$(\mathcal{C}, \rho)$ rounds, we reach the last position in a $(\mathcal{C}, \rho)$-history of length game-depth$(\mathcal{C}, \rho)$. This position is guaranteed to be a winning position for Flipper, that is, a single vertex graph. Here, model checking is trivial. By induction, we assume an algorithm for graphs resulting from $\ell + 1$ rounds of play (with the precise definition given by the following Definition 12.15), and use it to also do model checking on for graphs with only ℓ rounds played. Repeating this procedure gives us an algorithm for graphs on which zero rounds have been played, that is, a model checking algorithm for all graphs from $\mathcal{C}$. The choice $\rho := (16q(2^q) + 1)(2^q + 1)$ for the radius of the game emerges from the details of our proofs.

Definition 12.15. Let $\mathcal{C}$ be a monadically stable graph class and let $q, \ell, c \in \mathbb{N}$. We choose a radius $\rho := (16q(2^q) + 1)(2^q + 1)$ for the Flipper game. Note that ρ depends only on q. Consider an algorithm that, given as input

- a $(\mathcal{C}, \rho)$-history $(G_0, \mathcal{I}_0), \ldots, (G_\ell, \mathcal{I}_\ell)$ of length ℓ from the Flipper game,
- a unary expansion G of G_ℓ with a signature of at most c unary predicates, and
- a sentence φ with quantifier rank at most q,

decides whether $G \models \varphi$. We say this is an *efficient* MC$(\mathcal{C}, q, \ell, c)$*-algorithm*, if there exists a function f_{MC} bounding the runtime for every $\varepsilon > 0$ by

$$f_{\mathrm{MC}}(q, \ell, c, \varepsilon) \cdot |V(G)|^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5,$$

where $d := \text{game-depth}(\mathcal{C}, \rho) - \ell$ bounds the number of rounds needed to win the remaining Flipper game.

12.3.2 Computing Guarded Formulas

As the central building block of our algorithm, the following theorem converts sentences into guarded sentences, assuming we already have an efficient model checking algorithm for graphs where the game has progressed by one extra round.

Theorem 12.16. *There is an algorithm with the following property. For every monadically stable graph class* $\mathcal{C}$ *and* $\rho = (16q(2^q) + 1)(2^q + 1)$*, given as input*

- *a* $\big(\mathcal{C}, \rho\big)$*-history* $\mathcal{H} = (G_0, \mathcal{I}_0), \ldots, (G_\ell, \mathcal{I}_\ell)$ *of length* ℓ*,*
- *a unary expansion* G *of* G_ℓ *with a signature of at most* c *unary predicates,*
- *a sentence* φ *with quantifier rank at most* q*, and*
- *an efficient* MC$(\mathcal{C}, q, \ell + 1, c + 3)$*-algorithm,*

the algorithm computes sets $U_1, \ldots, U_t \subseteq V(G)$*, for some constant* $t \leqslant \mathrm{const}(q, c)$*, as well as a* $(U_1, \ldots, U_t)$*-guarded sentence* ξ *of quantifier rank* q*. Each* U_i *is contained in an* $(8q \cdot 2^q)$*-neighborhood of* G *and*

$$G \models \varphi \quad \Leftrightarrow \quad G\langle U_1, \ldots, U_t\rangle \models \xi.$$

For every class $\mathcal{C}$*, there exists a function* $f(q, c, \ell, \varepsilon)$ *such that for every* $\varepsilon > 0$*, the running time of this procedure is bounded by*

$$f(q, c, \ell, \varepsilon) \cdot |V(G)|^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5,$$

where $d := \text{game-depth}(\mathcal{C}, \rho) - \ell$ *bounds the number of rounds needed to win the remaining Flipper game.*

Instead of guarding all quantifiers at once, we start with guarding only one outermost quantifier. The following theorem will be the central step of our construction.

Theorem 12.17. *There is an algorithm with the following property. For every monadically stable graph class $\mathcal{C}$ and $\rho = (16q(2^q)+1)(2^q+1)$, given as input*

- *a $(\mathcal{C},\rho)$-history $\mathcal{H} = (G_0,\mathcal{I}_0),\dots,(G_\ell,\mathcal{I}_\ell)$ of length ℓ,*
- *a unary expansion G of G_ℓ with a signature of at most c unary predicates,*
- *a formula $\exists x\ \varphi(\bar{y},x)$ of quantifier rank at most q,*
- *sets $W_1,\dots,W_{|\bar{y}|}$, each contained in an r-neighborhood of G, and*
- *an efficient $\mathrm{MC}(\mathcal{C},q,\ell+1,c+3)$-algorithm,*

the algorithm computes sets $U_1,\dots,U_t \subseteq V(G)$, for some constant $t \leqslant \mathrm{const}(q,c,|\bar{y}|)$. Each U_i is contained in an $(r+8(2^q))$-neighborhood of G and for all tuples $\bar{w} \in W_1 \times \ldots \times W_{|y|}$, we have

$$G \models \exists x\ \varphi(\bar{w},x) \quad \Leftrightarrow \quad \bigvee_{i=1}^{t} G\langle U_i\rangle \models \exists x \in U_i\ \varphi(\bar{w},x).$$

For every class $\mathcal{C}$, there exists a function $f(q,c,\ell,\varepsilon,|\bar{y}|)$ such that for every $\varepsilon > 0$, the running time of this procedure is bounded by

$$f(q,c,\ell,\varepsilon,|\bar{y}|)\cdot|V(G)|^{((1+\varepsilon)^d)}\cdot|V(G_0)|^5,$$

where $d :=$ game-depth$(\mathcal{C},\rho)-\ell$ bounds the number of rounds needed to win the remaining Flipper game.

Proof. Let n be the number of vertices of G. Our goal is to compute the set of guards $\mathcal{U} = \{U_1,\dots,U_t\}$.

Neighborhood Cover Computation. As G_0 is from a monadically stable class $\mathcal{C}$ and G is obtained from G_0, by performing ℓ many 2-flips, we know that G is also from a monadically stable class $\mathcal{C}_\ell$ depending only on $\mathcal{C}$ and ℓ. We use Theorem 11.2 to compute a 2^q-neighborhood cover of diameter $4\cdot 2^q$. For every $\varepsilon > 0$ the overlap of the cover is bounded by $O_{\mathcal{C},\ell,q,\varepsilon}(n^\varepsilon)$ and the procedure takes time $O(n^5)$. Let $\{C_1,\dots,C_m\}$ be the computed cover. Without loss of generality, we can assume $m \leqslant n$, since otherwise redundant sets can be removed. We partition the vertices of G into sets $V_1,\dots,V_m$ such that for all $v \in V_i$, $N_{2^q}[v] \subseteq C_i$. Ties are broken arbitrarily.

Splitting the Existential Quantifier. It will be useful to partition the existential quantification of x in our input formula $\exists x\ \varphi(\bar{y},x)$ into a quantification over sets that are near and that are far from $W_1,\dots,W_{|\bar{y}|}$. To this end, let $V_i' := V_i \setminus N_{2^q}\left[\bigcup_{k=1}^{|\bar{y}|} W_k\right]$. Since every vertex of G is in some V_i, for all tuples $\bar{w} \in W_1 \times \ldots \times W_{|y|}$

$$G \models \exists x\ \varphi(\bar{w},x) \quad \Leftrightarrow$$
$$\bigvee_{i=1}^{|\bar{y}|} G\langle N_{2^q}[W_i]\rangle \models \exists x \in N_{2^q}[W_i]\ \varphi(\bar{w},x)\ \vee\ \bigvee_{i=1}^{m} G\langle V_i'\rangle \models \exists x \in V_i'\ \varphi(\bar{w},x). \quad (12.1)$$

Remember that the size of our solution $\mathcal{U}$ may depend only on q, c, and $|\bar{y}|$. Adding the sets $N_{2^q}[W_1],\dots,N_{2^q}[W_{|\bar{y}|}]$ to $\mathcal{U}$ would respect this size constraint. However, since m may depend on n, we are not allowed to add all sets $V_1',\dots,V_m'$ to $\mathcal{U}$. In the remainder of this proof, we will use Theorem 12.14 and the fact that each V_i' is sufficiently far away from $W_1,\dots,W_{|y|}$ to construct a set $S \subseteq [m]$ with the following property.

Property 12.18. The size of $S \subseteq [m]$ depends only on q and c and for all tuples $\bar{w} \in W_1 \times \ldots \times W_{|y|}$

$$\bigvee_{i=1}^{m} G\langle V_i' \rangle \models \exists x \in V_i' \; \varphi(\bar{w}, x) \quad \Rightarrow \quad \bigvee_{i \in S} G\langle X \mapsto N_{5(2^q)}[V_i'] \rangle \models \exists x \in X \; \varphi(\bar{w}, x).$$

After we found such a set S, we set

$$\mathcal{U} = \{N_{2^q}[W_1], \ldots, N_{2^q}[W_{|\bar{y}|}]\} \cup \{N_{5(2^q)}[V_i'] : i \in S\}.$$

Note that $|\mathcal{U}|$ depends only on q, c, and $|\bar{y}|$. Combining (12.1) and Property 12.18, it holds for all tuples $\bar{w} \in W_1 \times \ldots \times W_{|y|}$ that

$$G \models \exists x \; \varphi(\bar{w}, x) \quad \Rightarrow \quad \bigvee_{U \in \mathcal{U}} G\langle U \rangle \models \exists x \in U \; \varphi(\bar{w}, x).$$

The backwards implication of this statement holds obviously, since the right-hand side merely restricts the quantification of x. This yields the central statement

$$G \models \exists x \; \varphi(\bar{w}, x) \quad \Leftrightarrow \quad \bigvee_{U \in \mathcal{U}} G\langle U \rangle \models \exists x \in U \; \varphi(\bar{w}, x).$$

Since each W_i is contained in an r-neighborhood of G, each $N_{2^q}[W_i]$ is contained in an $(r + 2^q)$-neighborhood. Each set $N_{2^q}[V_i']$ is contained in C_i, which by construction is contained in an $4(2^q)$-neighborhood of G. It follows that $N_{5(2^q)}[V_i'] = N_{4(2^q)}[N_{2^q}[V_i']]$ is contained in a $8(2^q)$-neighborhood in G. Hence, each $U \in \mathcal{U}$ is contained in an $(r + 8(2^q))$-neighborhood of G. To finish the proof, we have to compute a small representative set S with Property 12.18.

Flip and Type Computation. As a first step towards computing S, we show how to use our given efficient MC$(\mathcal{C}, q, \ell + 1, c + 3)$-algorithm to compute $\text{tp}_q(G[N_{2^{q-1}-1}[V_i']]\langle X \mapsto V_i' \rangle)$ for all $i \in [m]$. To this end, we do for every $i \in [m]$ the following computations. Let $G_\ell^{\text{loc}} := G[N_{2^{q-1}-1}[V_i']]$. Note that this corresponds to a Localizer move in the radius-$4(2^q) \leqslant \rho$ Flipper game. We apply the Flipper strategy flip* to the graph G_ℓ^{loc} and internal state $\mathcal{I}_\ell$, yielding a 2-flip $G_{\ell+1}$ of G_ℓ^{loc} and a new internal state $\mathcal{I}_{\ell+1}$. By Theorem 10.4, this takes time

$$g_2(q) \cdot |V(G_0)|^2,$$

for some function $g_2(q)$ depending on $\mathcal{C}$. We can now extend $\mathcal{H}$ to a $(\mathcal{C}, \rho)$-history of length $\ell + 1$ by appending the new pair $(G_{\ell+1}, \mathcal{I}_{\ell+1})$. Using Lemma 4.1, we can recover a partition $\mathcal{P}$ and relation $F \subseteq \mathcal{P}^2$ witnessing that $G_{\ell+1}$ is a 2-flip of G_ℓ^{loc} in time $O(|V(G)|^2)$. We spend 3 additional unary predicates to construct $G_{\ell+1}^+$ by marking in $G_{\ell+1}$ the parts of $\mathcal{P}$ and the vertices from V_i'. Next, we enumerate the set Φ of normalized first-order sentences with quantifier rank at most q over the signature of $G_{\ell+1}^+$. Recall that $|\Phi|$ is bounded by a function of q and c. We use the given efficient MC$(\mathcal{C}, q, \ell + 1, c + 3)$-algorithm to evaluate every formula from Φ on $G_{\ell+1}^+$ and therefore compute $\text{tp}_q(G_{\ell+1}^+)$ in time

$$g_3(q, c) \cdot f_{\text{MC}}(q, \ell + 1, c + 3, \varepsilon) \cdot |V(G_{\ell+1}^+)|^{((1+\varepsilon)^{d-1})} \cdot |V(G_0)|^5,$$

for some function $g_3(q, c)$. Let us now argue how to derive $\text{tp}_q(G[N_{2^{q-1}-1}[V_i']]\langle V_i' \rangle)$ from $\text{tp}_q(G_{\ell+1}^+)$. This is easy to do by observing that for every sentence ψ, we have

$$\psi \in \text{tp}_q(G[N_{2^{q-1}-1}[V_i']]\langle V_i' \rangle) \quad \text{if and only if} \quad \psi' \in \text{tp}_q(G_{\ell+1}^+),$$

where ψ' is obtained from ψ by substituting every occurrence of the edge relation $E(x,y)$ with

$$E(x,y) \text{ XOR } \left(\bigvee_{(A,B)\in F} x \in A \wedge y \in B \right).$$

Similarly, we can derive $\mathrm{tp}_q(G[N_{2^{q-1}-1}[V_i']]\langle X \mapsto V_i'\rangle)$.

Computing a Representative Set. Now we use the previously computed q-types to pick S as a minimal subset of $[m]$ such that

$$\{\mathrm{tp}_q(G[N_{2^{q-1}-1}[V_i']]\langle X \mapsto V_i'\rangle) : i \in [m]\} = \{\mathrm{tp}_q(G[N_{2^{q-1}-1}[V_i']]\langle X \mapsto V_i'\rangle) : i \in S\}.$$

The size of S is at most the number of possible q-types on graphs with $c+3$ unary predicates, and thus can be bounded as a function of q and c. In order to show that S satisfies Property 12.18, let us fix $\bar{w} \in W_1 \times \cdots \times W_{|\bar{y}|}$ and argue that

$$\bigvee_{i=1}^{m} G\langle V_i'\rangle \models \exists x \in V_i' \; \varphi(\bar{w},x) \quad \Rightarrow \quad \bigvee_{i\in S} G\langle X \mapsto N_{5(2^q)}(V_i')\rangle \models \exists x \in X \; \varphi(\bar{w},x).$$

Assume $G\langle V_i'\rangle \models \exists x \in V_i' \; \varphi(\bar{w},x)$ for some i. If $V_i' \subseteq \bigcup_{j\in S} N_{5(2^q)}[V_j']$ for some $j \in S$, then the right-hand side follows immediately, so we can assume $V_i' \not\subseteq \bigcup_{j\in S} N_{5(2^q)}[V_j']$ for all $j \in S$. Fix some $j \in S$ and let us show that $\mathrm{dist}(V_i', V_j') > 2^q$ and $\mathrm{dist}(\bar{w}, V_i' \cup V_j') > 2^q$. Since we have $V_i' \not\subseteq \bigcup_{j\in S} N_{5(2^q)}[V_j']$, there exists a vertex in V_i' that has distance greater than $5(2^q)$ from every vertex in V_j'. Since V_i' embeds in a subgraph of G with diameter at most $4(2^q)$, every vertex in V_i' has distance greater than 2^q from every vertex in V_j'. This means $\mathrm{dist}(V_i', V_j') > 2^q$. We finally establish $\mathrm{dist}(\bar{w}, V_i' \cup V_j') > 2^q$ by combining

$$V_i' := V_i \setminus N_{2^q}\left[\bigcup_{k=1}^{|\bar{y}|} W_k\right], \quad V_j' := V_j \setminus N_{2^q}\left[\bigcup_{k=1}^{|\bar{y}|} W_k\right], \quad \bar{w} \in W_1 \times \cdots \times W_{|\bar{y}|}.$$

The set S was chosen representative in the sense that there is some $j \in S$ with

$$\mathrm{tp}_q(G[N_{2^{q-1}-1}[V_i']]\langle V_i' \to X\rangle) = \mathrm{tp}_q(G[N_{2^{q-1}-1}[V_j']]\langle V_j' \to X\rangle).$$

Since $G\langle V_i'\rangle \models \exists x \in V_i' \; \varphi(\bar{w},x)$, by Theorem 12.14, also $G\langle V_j'\rangle \models \exists x \in V_j' \; \varphi(\bar{w},x)$ and the right-hand side holds. Hence, S satisfies Property 12.18.

Running Time Analysis. At first, we analyze the running time spent for the computations in the paragraph *Flip and Type Computation.* As stated there, the run time is (using $m \leqslant n$) bounded by

$$\sum_{i\in[m]} g_2(q) \cdot |V(G_0)|^2 \leqslant n \cdot g_2(q) \cdot |V(G_0)|^2 \tag{12.2}$$

for computing the flips, plus

$$\sum_{i\in[m]} g_3(q,c) \cdot f_{\mathrm{MC}}(q, \ell+1, c+3, \varepsilon) \cdot |N_{2^q}[V_i']|^{((1+\varepsilon)^{d-1})} \cdot |V(G_0)|^5$$

for computing the q-types. Note that for all $\alpha \geqslant 1$ and non-negative numbers $n_1, \ldots, n_m$ we have $\sum_{i\in[m]} n_i^\alpha \leqslant (\sum_{i\in[m]} n_i)^\alpha$, bounding the running time for the q-type computation by

$$g_3(q,c) \cdot f_{\mathrm{MC}}(q, \ell+1, c+3, \varepsilon) \cdot \Big(\sum_{i\in[m]} |N_{2^q}[V_i']| \Big)^{((1+\varepsilon)^{d-1})} \cdot |V(G_0)|^5.$$

For every $i \in m$, we have $N_{2^q}[V_i'] \subseteq N_{2^q}[V_i] \subseteq C_i$, yielding

$$\sum_{i \in [m]} |N_{2^q}[V_i']| \leqslant \sum_{i \in [m]} |C_i| \leqslant g_1(q, \varepsilon, \ell) \cdot n^{1+\varepsilon},$$

where the last bound follows from the fact that we have n vertices, each occurring in at most $g_1(q, \varepsilon, \ell) \cdot n^\varepsilon$ clusters of the cover $\{C_1, \ldots, C_m\}$. Combining the previous two inequalities bounds the running time of the type computation by

$$g_3(q, c) \cdot f_{\mathrm{MC}}(q, \ell + 1, c + 3, \varepsilon) \cdot \left(g_1(q, \varepsilon, \ell) \cdot n^{1+\varepsilon}\right)^{((1+\varepsilon)^{d-1})} \cdot |V(G_0)|^5,$$

which is equal to

$$g_3(q, c) \cdot f_{\mathrm{MC}}(q, \ell + 1, c + 3, \varepsilon) \cdot g_1(q, \varepsilon, \ell)^{((1+\varepsilon)^{d-1})} \cdot n^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5. \tag{12.3}$$

The total running time spent in this paragraph, as given by the sum of (12.2) and (12.3) can be bounded by $g_4(q, \varepsilon, \ell) \cdot n^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5$, for some function $g_4(q, \varepsilon, \ell)$.

The computation in the paragraph *Neighborhood Cover Computation* takes time $O(n^5)$. Since the size of the representative set is bounded by a function of q and c, we can bound the computation time for the paragraphs *Splitting the Existential Quantifier* and *Computing a Representative Set* by $g_5(q, c, |\bar{y}|) \cdot n^2$, for some function $g_5(q, c, |\bar{y}|)$. Since $n \leqslant |V(G_0)|$, we can choose a function $f(q, c, \ell, \varepsilon, |\bar{y}|)$ such that the total running time is bounded by

$$f(q, c, \ell, \varepsilon, |\bar{y}|) \cdot n^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5. \qquad \square$$

Now we obtain Theorem 12.16 by simply applying Theorem 12.17 repeatedly, once for each quantifier. This will require no new insights, but will be a bit tedious to analyze. To help our inductive proof, we prove the following stronger statement. Then Theorem 12.16 follows as a special case when φ has no free variables, $p = q$ and $r = 0$.

Lemma 12.19. *There is an algorithm with the following property. For every monadically stable graph class $\mathcal{C}$ and $\rho = (16q(2^q) + 1)(2^q + 1)$, given as input*

- *a $(\mathcal{C}, \rho)$-history $\mathcal{H} = (G_0, \mathcal{I}_0), \ldots, (G_\ell, \mathcal{I}_\ell)$ of length ℓ,*
- *a unary expansion G of G_ℓ with a signature of at most c unary predicates,*
- *a formula $\varphi(\bar{y})$ with quantifier rank at most $p \leqslant q$,*
- *sets $W_1, \ldots, W_{|\bar{y}|}$, each contained in an r-neighborhood of G, and*
- *an efficient $\mathrm{MC}(\mathcal{C}, q, \ell + 1, c + 3)$-algorithm,*

the algorithm computes sets $U_1, \ldots, U_t \subseteq V(G)$, for some constant $t \leqslant \mathrm{const}(p, q, c, |\bar{y}|)$, as well as a $(U_1, \ldots, U_t)$-guarded formula $\xi(\bar{y})$ of quantifier rank p. Each U_i is contained in an $(r + q \cdot 8(2^q))$-neighborhood of G and for all tuples $\bar{w} \in W_1 \times \ldots \times W_{|y|}$ we have

$$G \models \varphi(\bar{w}) \quad \Leftrightarrow \quad G\langle U_1, \ldots, U_t \rangle \models \xi(\bar{w}).$$

For every class $\mathcal{C}$, there exists a function $f(p, q, c, \ell, \varepsilon)$ such that for every $\varepsilon > 0$, the run time of this procedure is bounded by

$$f(p, q, c, \ell, \varepsilon) \cdot |V(G)|^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5,$$

where $d :=$ game-depth$(\mathcal{C}, \rho) - \ell$ bounds the number of rounds needed to win the remaining Flipper game.

Proof. Let $\varepsilon > 0$. We prove the lemma by induction on p. For $p = 0$, note that every quantifier-free formula is $\emptyset$-guarded, and thus we can set $\xi(\bar{y}) := \varphi(\bar{y})$ and there is nothing more to show. Thus assume $p > 0$ and that the statement holds for $p-1$. We will construct an algorithm for p using the assumed algorithm for $p-1$ as a subroutine.

By normalization, $|\varphi|$ depends only on p, c and $|\bar{y}|$. Furthermore, $\varphi(\bar{y})$ is a boolean combination of formulas of the form $\exists x\ \psi(\bar{y}, x)$ of quantifier rank at most p. Thus, it is sufficient to prove the theorem for a single such formula $\exists x\ \psi(\bar{y}, x)$.

We apply Theorem 12.17 giving it as input

- the history $\mathcal{H} = (G_0, \mathcal{I}_0), \ldots, (G_\ell, \mathcal{I}_\ell)$,
- a unary expansion G of G_ℓ with a signature of at most c unary predicates,
- the formula $\exists x\ \psi(\bar{y}, x)$ of quantifier rank at most $p \leqslant q$,
- the sets $W_1, \ldots, W_{|\bar{y}|}$, each contained in an r-neighborhood of G, and
- the given $\mathrm{MC}(\mathcal{C}, q, \ell + 1, c + 3)$-algorithm.

In time

$$f'(q, c, \ell, \varepsilon, |\bar{y}|) \cdot |V(G)|^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5 \tag{12.4}$$

this yields sets $R_1, \ldots, R_{t'} \subseteq V(G)$ for some constant t' depending only on q, c, and $|\bar{y}|$. Each R_i is contained in an $(r + 8(2^q))$-neighborhood of G, such that for all tuples $\bar{w} \in W_1 \times \ldots \times W_{|y|}$, we have

$$G \models \exists x\ \psi(\bar{w}, x) \quad \Leftrightarrow \quad \bigvee_{i=1}^{t'} G\langle R_i \rangle \models \exists x \in R_i\ \psi(\bar{w}, x). \tag{12.5}$$

For each $i \in [t']$ we apply the algorithm for $p-1$ given by the induction hypothesis on

- the history $\mathcal{H}$, graph G, and $\mathrm{MC}(\mathcal{C}, q, \ell + 1, c + 3)$-algorithm,
- the formula $\psi(\bar{y}, x)$ of quantifier rank at most $p - 1 \leqslant q - 1$,
- the sets $W_1, \ldots, W_{|y|}, R_i \subseteq V(G)$, each contained in an $(r + 8(2^q))$-neighborhood of G.

In time

$$f(p-1, q, c, \ell, \varepsilon, |\bar{y}| + 1) \cdot |V(G)|^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5 \tag{12.6}$$

this yields a family of guarding sets $\mathcal{U}_i$ with $|\mathcal{U}_i|$ depending on $p-1$, q, c, and $|\bar{y}|$, as well as a $\mathcal{U}_i$-guarded formula $\xi_i(\bar{y})$ of quantifier rank $q-1$. Each $U \in \mathcal{U}_i$ is contained in an $\big(r + 8(2^q) + (q-1) \cdot 8(2^{q-1})\big)$-neighborhood of G (and thus in an $(r + q \cdot 8(2^q))$-neighborhood of G). For all tuples $\bar{w}v \in W_1 \times \ldots \times W_{|\bar{y}|} \times R_i$ we have

$$G \models \psi(\bar{w}, v) \quad \Leftrightarrow \quad G\langle \mathcal{U}_i \rangle \models \xi_i(\bar{w}, v).$$

Since the above statement holds no matter how $v \in R_i$ is chosen, existentially quantifying $v \in R_i$ preserves the equivalence. Hence, for all tuples $\bar{w} \in W_1 \times \ldots \times W_{|y|}$

$$G\langle R_i \rangle \models \exists x \in R_i\ \psi(\bar{w}, x) \quad \Leftrightarrow \quad G\langle R_i \rangle\langle \mathcal{U}_i \rangle \models \exists x \in R_i\ \xi_i(\bar{w}, x). \tag{12.7}$$

Combining (12.5) and (12.7) yields for every $\bar{w} \in W_1 \times \ldots \times W_{|y|}$,

$$G \models \exists x\ \psi(\bar{w}, x) \quad \Leftrightarrow \quad \bigvee_{i=1}^{t'} G\langle R_i \rangle\langle \mathcal{U}_i \rangle \models \exists x \in R_i\ \xi_i(\bar{w}, x),$$

which is equivalent to

$$G\langle R_1 \rangle\langle \mathcal{U}_1 \rangle \ldots \langle R_{t'} \rangle\langle \mathcal{U}_{t'} \rangle \models \bigvee_{i=1}^{t'} \exists x \in R_i\ \xi_i(\bar{w}, x).$$

Thus, we can define our guarding sets $\mathcal{U} = \{U_1, \dots, U_t\}$ as $\mathcal{U} := \{R_1, \dots, R_{t'}\} \cup \bigcup_{i=1}^{t'} \mathcal{U}_i$.

The running time is bounded by the bound (12.4) for the invocation of Theorem 12.17, plus t' times the bound (12.6) for the recursive calls with $p-1$, plus some minor bookkeeping overhead. We can choose $f(p, q, c, \ell, \varepsilon)$ such that this is at most

$$f(p, q, c, \ell, \varepsilon) \cdot |V(G)|^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5. \qquad \square$$

12.3.3 Reducing the Evaluation Radius

Our overall goal is to evaluate a sentence with quantifier rank q on a graph G. In the previous section, we have rewritten the sentence into an equivalent $\mathcal{U}$-guarded sentence of the same quantifier rank using guards $\mathcal{U} = \{U_1, \dots, U_t\}$. Each of the sets $U_i \subseteq V(G)$ is contained in an $r := q \cdot 8(2^q)$-neighborhood of G and thus the induced graph $G[U_1 \cup \dots \cup U_t]$ consists of components, which are contained in neighborhoods with radius at most $(2r+1)t$ in G.

One could imagine evaluating the $\mathcal{U}$-guarded sentence on G by recursing into each of these components and to compute flips using the strategy for the radius-$(2r+1)t$ Flipper game. Let us argue that this cannot work. The radius of the Flipper game is not allowed to grow over time, since otherwise the game is not guaranteed to terminate in a fixed number of rounds. In our construction, however, the number t of guards depends on the number c of unary predicates added over time and thus grows with the number of rounds of the Flipper game played so far. Thus, we are not allowed to recurse into components with radius $(2r+1)t$. We have to choose a fixed radius ρ for the Flipper game, depending only on q and $\mathcal{C}$.

In this section, we show that we can evaluate the $\mathcal{U}$-guarded sentence by only looking at subgraphs of G that are contained in neighborhoods of radius $\rho := (2r+1)(2^q+1)$, a quantity that depends only on q and $\mathcal{C}$ and does not grow over time. This is a consequence of the following Proposition 12.20. Remember that, to avoid lengthy additional notation, $\mathcal{U}$ will refer both to unary predicates guarding a formula and to the corresponding vertex sets in a graph G that interpret these predicates. It will be clear from the context which one is meant.

Proposition 12.20. *For a given $\mathcal{U}$-guarded sentence φ with quantifier rank at most q and symmetric relation $\mathcal{R} \subseteq \mathcal{U} \times \mathcal{U}$, one can compute a sentence $\varphi_{\mathcal{R}}$ such that for every graph G and set $\mathcal{U} \subseteq \mathcal{P}(V(G))$ satisfying*

$$\mathcal{R} = \{(U, W) \in \mathcal{U} \times \mathcal{U} : U \text{ and } W \text{ share a vertex or a connecting edge in } G\},$$

we have

$$G\langle\mathcal{U}\rangle \models \varphi \quad \Leftrightarrow \quad G\langle\mathcal{U}\rangle \models \varphi_{\mathcal{R}}.$$

Moreover, $\varphi_{\mathcal{R}}$ is a boolean combination of sentences with quantifier rank at most q and each sentence mentioned in $\varphi_{\mathcal{R}}$ is $\mathcal{U}'$-guarded for some $\mathcal{U}' \subseteq \mathcal{U}$ such that the graph $(\mathcal{U}, \mathcal{R})[\mathcal{U}']$ has diameter at most 2^q.

In particular, if each set of $\mathcal{U} \subseteq \mathcal{P}(V(G))$ is contained in an r-neighborhood of G, then $\bigcup \mathcal{U}'$ is contained in a subgraph of G with diameter at most $(2r+1)(2^q+1)$.

To see that the final "In particular, …" part follows from the central part of the statement, assume each set of $\mathcal{U} \subseteq \mathcal{P}(V(G))$ is contained in an r-neighborhood of G. For a $\mathcal{U}'$-guarded sentence ξ in $\psi_{\mathcal{R}}$, the graph $(\mathcal{U}, \mathcal{R})[\mathcal{U}']$ has diameter at most 2^q and all $(U_1, U_2) \in \mathcal{R}$ share a vertex or a connecting edge in G. As can be seen in the figure below, $\bigcup \mathcal{U}'$ is contained in a subgraph of G with diameter at most $(2r+1)(2^q+1)$.

We prove the central part of Proposition 12.20 inductively using a stronger statement involving formulas with free variables.

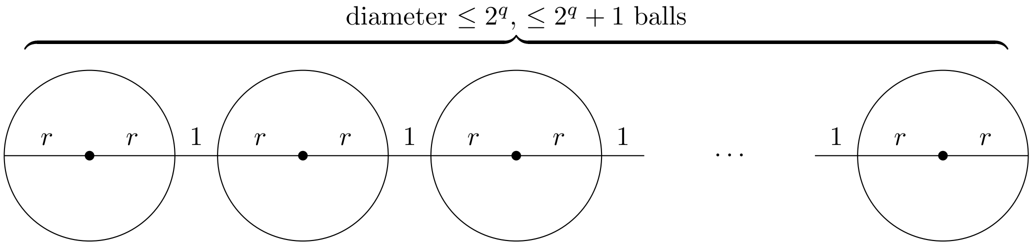

Lemma 12.21. *For a given $\mathcal{U}$-guarded formula $\varphi(\bar{y})$ with quantifier rank at most q, symmetric relation $\mathcal{R} \subseteq \mathcal{U} \times \mathcal{U}$ and sequence $U_1, \ldots, U_{|\bar{y}|} \in \mathcal{U}$ one can compute a formula $\varphi_{\mathcal{R}}(\bar{y})$ such that for every graph G, set $\mathcal{U} \subseteq \mathcal{P}(V(G))$ associated with the predicates $\mathcal{U}$ satisfying*

$$\mathcal{R} = \{(U, W) \in \mathcal{U} \times \mathcal{U} : U \text{ and } W \text{ share a vertex or a connecting edge in } G\},$$

and every $\bar{w} \in U_1 \times \cdots \times U_{|\bar{y}|}$ we have

$$G\langle\mathcal{U}\rangle \models \varphi(\bar{w}) \quad \Leftrightarrow \quad G\langle\mathcal{U}\rangle \models \varphi_{\mathcal{R}}(\bar{w}).$$

Moreover, $\varphi_{\mathcal{R}}(\bar{y})$ is a boolean combination of formulas with quantifier rank at most q and for each formula ξ mentioned in $\varphi_{\mathcal{R}}$ there exists $\mathcal{U}' \subseteq \mathcal{U}$ such that ξ is $\mathcal{U}'$-guarded, $(\mathcal{U}, \mathcal{R})[\mathcal{U}']$ has diameter at most 2^q, and $\{U_i : y_i \in \text{free}(\xi)\} \subseteq \mathcal{U}'$.

Proof. We consider an arbitrary graph G and $\mathcal{U} \subseteq \mathcal{P}(V(G))$ associated with the predicates $\mathcal{U}$ such that for all $U, W \in \mathcal{U}$ we have $(U, W) \in \mathcal{R}$ if and only if U, W share a vertex or a connecting edge in G. Let us also fix a sequence $U_1, \ldots, U_{|\bar{y}|} \in \mathcal{U}$. We prove the claim by induction over the structure of φ.

Atoms. Since φ is an atom, it is $\emptyset$-guarded and has either one or two free variables. Assume $\varphi(y_1)$ has a single free variable. We set $\varphi_{\mathcal{R}} := \varphi$ and $\mathcal{U}' := \{U_1\}$. Then $\varphi_{\mathcal{R}}$ itself is $\mathcal{U}'$-guarded and $(\mathcal{U}, \mathcal{R})[\mathcal{U}']$ is the single vertex graph that has diameter $0 \leqslant 2^0$. Assume now $\varphi(y_1, y_2)$ is a binary atom, that is, without loss of generality either $E(y_1, y_2)$ or $(y_1 = y_2)$. If $(U_1, U_2) \in \mathcal{R}$ we set $\varphi_{\mathcal{R}} := \varphi$ and $\mathcal{U}' := \{U_1, U_2\}$. Again, φ_R is $\mathcal{U}'$-guarded and $(\mathcal{U}, \mathcal{R})[\mathcal{U}']$ has diameter $1 = 2^0$. Otherwise, $(U_1, U_2) \notin \mathcal{R}$ and U_1, U_2 neither share a vertex nor a connecting edge in G. This implies $G \not\models E(w_1, w_2)$ and $G \not\models (w_1 = w_2)$ for all $w_1 \in U_x, w_2 \in U_y$. We set $\varphi_{\mathcal{R}}$ to be the false atom $\bot$ and $\mathcal{U}' = \emptyset$.

Boolean Combinations. If φ is of the form $\psi^1 \wedge \psi^2$ or $\neg\psi^1$ the construction is obvious: We obtain $\varphi^1_{\mathcal{R}}$ and $\varphi^2_{\mathcal{R}}$ via induction and set either $\varphi_{\mathcal{R}} := \psi^1_{\mathcal{R}} \wedge \psi^2_{\mathcal{R}}$ or $\varphi_{\mathcal{R}} := \neg\psi^1_{\mathcal{R}}$.

Existential Quantifiers. Assume $\varphi(\bar{y}) = \exists x \in U\ \psi(\bar{y}x)$. We apply the statement inductively on $\psi(\bar{y}x)$ (extending the sequence $U_1, \ldots, U_{|\bar{y}|}$ with U) and obtain a boolean combination $\psi_{\mathcal{R}}(\bar{y}x)$ of formulas with quantifier rank at most $q-1$ such that for every $\bar{w}v \in U_{y_1} \times \cdots \times U_{y_{|\bar{y}|}} \times U$

$$G\langle\mathcal{U}\rangle \models \psi(\bar{w}v) \quad \Leftrightarrow \quad G\langle\mathcal{U}\rangle \models \psi_{\mathcal{R}}(\bar{w}v).$$

For each formula ξ mentioned in $\psi_{\mathcal{R}}(\bar{y}x)$ we have, by induction, a set $\mathcal{U}_\xi \subseteq \mathcal{U}$ such that ξ is $\mathcal{U}_\xi$-guarded and $(\mathcal{U}, \mathcal{R})[\mathcal{U}_\xi]$ has diameter at most 2^{q-1}. The additional crucial property we obtain by induction is that $x \in \text{free}(\xi)$ implies $U \in \mathcal{U}_\xi$. We partition the formulas mentioned in

$\psi_{\mathcal{R}}(\bar{y}x)$ into sets X and $\overline{X}$, where X contains all formulas ξ with $x \in \text{free}(\xi)$, and $\overline{X}$ contains all formulas ξ with $x \notin \text{free}(\xi)$. The formulas in $\overline{X}$ are independent of x and we can thus write

$$\exists x \in U\ \psi_{\mathcal{R}}(\bar{y}x) \quad \equiv \quad \bigvee_{t:\overline{X}\to\{\bot,\top\}} \Big(\Big(\bigwedge_{\xi\in\overline{X}} (\xi(\bar{y}) \leftrightarrow t(\xi))\Big) \wedge \exists x \in U\ \psi_{\mathcal{R}}(\bar{y}x)\Big).$$

Now on the right-hand side, every occurrence of $\psi_{\mathcal{R}}$ is in a scope where the truth value of every $\xi \in \overline{X}$ is determined. Thus, we can replace every occurrence of ξ in $\psi_{\mathcal{R}}$ with said truth value $t(\xi) \in \{\bot, \top\}$. Let $\psi^t_{\mathcal{R}}$ be the formula obtained from $\psi_{\mathcal{R}}$ by replacing each occurrence of $\xi \in \overline{X}$ with $t(\xi)$. We obtain the equivalence

$$\exists x \in U\ \psi_{\mathcal{R}}(\bar{y}x) \quad \equiv \quad \varphi_{\mathcal{R}}(\bar{y}) := \bigvee_{t:\overline{X}\to\{\bot,\top\}} \Big(\Big(\bigwedge_{\xi\in\overline{X}} (\xi(\bar{y}) \leftrightarrow t(\xi))\Big) \wedge \exists x \in U\ \psi^t_{\mathcal{R}}(\bar{y}x)\Big).$$

Hence, for every $\bar{w} \in U_{y_1} \times \cdots \times U_{y_{|\bar{y}|}}$

$$G\langle\mathcal{U}\rangle \models \varphi(\bar{w}) \quad \Leftrightarrow \quad G\langle\mathcal{U}\rangle \models \varphi_{\mathcal{R}}(\bar{w}).$$

We observe that $\varphi_{\mathcal{R}}(\bar{y})$ is a boolean combination of old formulas from $\overline{X}$ and new formulas of the form $\exists x \in U\ \psi^t_{\mathcal{R}}$. All these formulas have quantifier rank at most q.

Consider now a new formula $\omega := \exists x \in U\ \psi^t_{\mathcal{R}}$ and let $\mathcal{U}_\omega := \bigcup_{\xi\in X} \mathcal{U}_\xi$. Since $\psi^t_{\mathcal{R}}$ eliminated all formulas from $\overline{X}$, we know that ω is $\mathcal{U}_\omega$-guarded. For all $\xi \in X$ we have $x \in \text{free}(\xi)$ and thus, as noted previously, $U \in \mathcal{U}_\xi$. By induction, each graph $(\mathcal{U}, \mathcal{R})[\mathcal{U}_\xi]$ has diameter at most 2^{q-1}. This means $(\mathcal{U}, \mathcal{R})[\mathcal{U}_\omega]$ is covered by graphs that all overlap in U and have diameter at most 2^{q-1}, implying that $(\mathcal{U}, \mathcal{R})[\mathcal{U}_\omega]$ has diameter at most 2^q. Since $\text{free}(\omega) \subseteq \bigcup_{\xi\in X} \text{free}(\xi)$, we also have $\{U_i : y_i \in \text{free}(\omega)\} \subseteq \mathcal{U}_\omega$. □

12.3.4 Main Result

By induction on the length of the Flipper game, we combine the observations from the previous two subsections into a single algorithm.

Proposition 12.22. *There is an algorithm that is an efficient* $\text{MC}(\mathcal{C}, q, \ell, c)$*-algorithm for every monadically stable class* $\mathcal{C}$ *and every* $q, \ell, c \in \mathbb{N}$.

Proof. Fix any monadically stable class $\mathcal{C}$ and $q, \ell, c \in \mathbb{N}$ and let $\rho := (16q(2^q) + 1)(2^q + 1)$. The algorithm mc we construct gets as input a $(\mathcal{C}, \rho)$-history $\mathcal{H} = (G_0, \mathcal{I}_0), \ldots, (G_\ell, \mathcal{I}_\ell)$, a unary expansion G of G_ℓ with at most c unary predicates, and a sentence φ with quantifier rank q. The goal is to decide whether $G \models \varphi$. We define mc recursively and prove its correctness and efficiency by induction on ℓ.

Base Case: If G contains only a single vertex, we trivially decide whether $G \models \varphi$ in time $O_{q,c}(1)$. This base case already establishes that mc is an efficient $\text{MC}(\mathcal{C}, q, \ell', c')$-algorithm for all $\ell', c' \in \mathbb{N}$ with $\ell' \geqslant \text{game-depth}(\mathcal{C}, \rho)$: By the guarantees of $\text{flip}^\star$, the last graph in every $(\mathcal{C}, \rho)$-history of length at least $\text{game-depth}(\mathcal{C}, \rho)$ is a winning position for Flipper and therefore only contains a single vertex, which is handled by this base case.

Inductive Case: If G is not a single vertex graph, we must have $\ell < \text{game-depth}(\mathcal{C}, \rho)$. We can now assume, by induction on ℓ, that for all $\ell', c' \in \mathbb{N}$ with $\ell' > \ell$ the algorithm mc is an efficient $\text{MC}(\mathcal{C}, q, \ell', c')$-algorithm which we can recursively call. We make use of this fact by calling the algorithm of Theorem 12.16 on $\mathcal{H}$, G, and φ, where we supply mc as an efficient $\text{MC}(\mathcal{C}, q, \ell+1, c+3)$-algorithm. We obtain a $\mathcal{U}$-guarded sentence ξ of quantifier rank q such that

$$G \models \varphi \quad \Leftrightarrow \quad G\langle\mathcal{U}\rangle \models \xi. \tag{12.8}$$

Here, $\mathcal{U} \subseteq \mathcal{P}(V(G))$ is a set with $|\mathcal{U}|$ depending only on q, c, such that each $U \in \mathcal{U}$ is contained in an $(8q \cdot 2^q)$-neighborhood of G.

In time $O(|\mathcal{U}|^2 \cdot |V(G)|^2)$, compute the relation

$$\mathcal{R} := \{(U, W) \in \mathcal{U} \times \mathcal{U} : U \text{ and } W \text{ share a vertex or a connecting edge in } G\}$$

Next, we invoke Proposition 12.20. This gives us a boolean combination ξ^* of sentences $\xi_1, \ldots, \xi_k$ with quantifier rank at most q such that

$$G\langle\mathcal{U}\rangle \models \xi \quad \Leftrightarrow \quad G\langle\mathcal{U}\rangle \models \xi^*. \tag{12.9}$$

Each sentence ξ_i is $\mathcal{U}_i$-guarded for some $\mathcal{U}_i \subseteq \mathcal{U}$ such that $\bigcup \mathcal{U}_i$ is contained in a subgraph of G with diameter at most $\rho = (16q(2^q) + 1)(2^q + 1)$, and thus also in a ρ-neighborhood of G. The running time of Proposition 12.20 is insignificant compared to the running time of Theorem 12.16. The time spent so far is bounded by

$$f(q, c, \ell, \varepsilon) \cdot |V(G)|^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5 \tag{12.10}$$

for some function $f(q, c, \ell, \varepsilon)$ and every $\varepsilon > 0$, where $d :=$ game-depth$(\mathcal{C}, \rho) - \ell$ bounds the number of rounds needed to win the remaining Flipper game.

Now for every $i \in [k]$, we proceed similarly as in the paragraph *Flip and Type Computation* of Theorem 12.17 to decide whether $G\langle\mathcal{U}\rangle \models \xi_i$. Let $G_\ell^{\text{loc}} := G[\bigcup \mathcal{U}_i]$ and as ξ_i is $\mathcal{U}_i$-guarded, we have

$$G\langle\mathcal{U}\rangle \models \xi_i \quad \Leftrightarrow \quad G_\ell^{\text{loc}}\langle\mathcal{U}\rangle \models \xi_i. \tag{12.11}$$

Note that since $\bigcup \mathcal{U}_i$ is contained in a ρ-neighborhood of G, the restriction to G_ℓ^{loc} corresponds to a Localizer move in the radius-ρ Flipper game. We apply the Flipper strategy $\mathsf{flip}^\star$ to the graph G_ℓ^{loc} and internal state $\mathcal{I}_\ell$, yielding a 2-flip $G_{\ell+1}$ of G_ℓ^{loc} and a new internal state $\mathcal{I}_{\ell+1}$. By Theorem 10.4, this takes time

$$g_2(q) \cdot |V(G_0)|^2, \tag{12.12}$$

for some function $g_2(q)$ depending on $\mathcal{C}$. Using Lemma 4.1, we compute in time $O(k \cdot |V(G)|^2)$ a partition $\mathcal{P}$ and a symmetric relation $F \subseteq \mathcal{P}$ witnessing that $G_{\ell+1}$ is a 2-flip of G_ℓ^{loc}. Let $G_{\ell+1}^+\langle\mathcal{U}\rangle$ be the monadic expansion of $G_{\ell+1}\langle\mathcal{U}\rangle$ with two unary predicates marking the parts of $\mathcal{P}$. We construct ξ_i' from ξ_i by substituting every occurrence of the edge relation $E(x, y)$ with

$$E(x, y) \ \text{XOR} \ \left(\bigvee_{(A,B) \in F} x \in A \wedge y \in B \right).$$

Then

$$G_\ell^{\text{loc}}\langle\mathcal{U}\rangle \models \xi_i \quad \Leftrightarrow \quad G_{\ell+1}^+\langle\mathcal{U}\rangle \models \xi_i'. \tag{12.13}$$

We now extend $\mathcal{H}$ to a $(\mathcal{C}, \rho)$-history of length $\ell + 1$ by appending the new pair $(G_{\ell+1}, \mathcal{I}_{\ell+1})$. We use a recursive call to mc (which is an efficient MC$(\mathcal{C}, q, \ell + 1, c + 2 + |\mathcal{U}|)$-algorithm by induction) to decide in time

$$f_2(q, \ell + 1, c + 2 + |\mathcal{U}|, \varepsilon) \cdot |V(G_{\ell+1}^+)|^{((1+\varepsilon)^{d-1})} \cdot |V(G_0)|^5 \tag{12.14}$$

whether $G_{\ell+1}^+\langle\mathcal{U}\rangle \models \xi_i'$, for some function f_2 and every $\varepsilon > 0$. By (12.11) and (12.13), this decides whether $G\langle\mathcal{U}\rangle \models \xi_i$.

Since we decided $G\langle\mathcal{U}\rangle \models \xi_i$ for all i, we can plug the truth values into the boolean combination ξ^*, telling us the answer to whether $G\langle\mathcal{U}\rangle \models \xi^*$. By (12.8) and (12.9), this finally gives us the answer whether $G \models \varphi$.

The total running time is bounded by (12.10) plus k times (12.12) and (12.14). Since both k and $|\mathcal{U}|$ are bounded by a function of q and c, we can choose $f_{\mathrm{MC}}(q, \ell, c, \varepsilon)$ such that for every $\varepsilon > 0$, the total running time is bounded by

$$f_{\mathrm{MC}}(q, \ell, c, \varepsilon) \cdot |V(G)|^{((1+\varepsilon)^d)} \cdot |V(G_0)|^5. \qquad \square$$

We obtain the final model checking result, by controlling the run time by choosing $\varepsilon > 0$ dependent on the game depth.

Theorem 2.2. *There is an algorithm that, given a graph G and a first-order sentence φ, decides whether $G \models \varphi$, and has the following property.*

For every monadically stable class $\mathcal{C}$, there exists a function $f : \mathbb{N} \times \mathbb{R} \to \mathbb{N}$ such that on any n-vertex graph $G \in \mathcal{C}$ and sentence φ the algorithm runs in time $f(|\varphi|, \varepsilon) \cdot n^{6+\varepsilon}$ for every $\varepsilon > 0$.

Proof. As input for the algorithm, we are given the graph G and a sentence φ. Let $\mathcal{H} := (G_0 := G, \mathcal{I}_0 := G)$ be the $(\mathcal{C}, \rho)$-history of length 0, for any monadically stable graph class $\mathcal{C} \ni G$ and any $\rho \in \mathbb{N}$. We call the algorithm from Proposition 12.22 on the history $\mathcal{H}$, the graph G, and the formula φ and return the result. This concludes the description of the algorithm.

We now show that for every monadically stable graph class $\mathcal{C}$, there exists a function $f : \mathbb{N} \times \mathbb{R} \to \mathbb{N}$ such that on any n-vertex graph $G \in \mathcal{C}$ and sentence φ the above algorithm runs in time $f(|\varphi|, \varepsilon) \cdot n^{6+\varepsilon}$ for every $\varepsilon > 0$.

Let q be the quantifier rank of φ and let $\rho := (16q(2^q) + 1)(2^q + 1)$. By Proposition 12.22 and the definition of an efficient $\mathrm{MC}(\mathcal{C}, q, 0, |\Sigma|)$-algorithm, the call to the algorithm from Proposition 12.22 runs in time

$$f_{\mathrm{MC}}(q, 0, |\Sigma|, \varepsilon) \cdot |V(G)|^{((1+\varepsilon)^{\text{game-depth}(\mathcal{C},\rho)})} \cdot |V(G)|^5$$

for some function f_{MC} that depends only on $\mathcal{C}$, and for any $\varepsilon > 0$. For any $\varepsilon' > 0$, we can choose $\varepsilon := (1 + \varepsilon')^{1/\text{game-depth}(\mathcal{C},\rho)} - 1 > 0$ in the above and get a running time of

$$f_{\mathrm{MC}}(q, 0, |\Sigma|, \varepsilon') \cdot |V(G)|^{1+\varepsilon'} \cdot |V(G)|^5 \leqslant f_{\mathrm{MC}}(q, 0, |\Sigma|, \varepsilon') \cdot |V(G)|^{6+\varepsilon'}.$$

As our choice of ε and all other parameters of f_{MC} depends only on $|\varphi|$ and $\mathcal{C}$, this means for every class $\mathcal{C}$ there is a function $f(|\varphi|, \varepsilon)$ such that the runtime is bounded by $f(|\varphi|, \varepsilon) \cdot |V(G)|^{6+\varepsilon}$ for every $\varepsilon > 0$.

Finally, note that this algorithm decides whether $G \models \varphi$ on any graph G: it always terminates and gives no wrong answers. This is due to the fact that each graph G is contained in the singleton class $\mathcal{C}_G := \{G\}$ that only contains G. As this class only contains a single finite graph, it is trivially monadically stable and the algorithm works. However, in this case where no stronger properties for the graph class $\mathcal{C}_G$ are assumed, the function f, that bounds the runtime, depends on G. This means the total running time can have an arbitrarily bad dependence on the number of vertices in G. $\square$

12.4 Strongly Uniform Fixed-Parameter Tractability

The textbook [21] classifies a parameterized problem as *uniform fixed-parameter tractable*, if it can be solved in time $f(k) \cdot n^c$ for every instance of size n with parameter k, where f is an arbitrary function and c is a constant. Here "uniform" refers to the fact that a single algorithm works for every value of the parameter k. Using this notation, Theorem 2.2 implies the following.

The first-order model checking problem (parameterized by formula length) is uniform fixed-parameter tractable on every monadically stable graph class.

We note that the above statement is strictly weaker than Theorem 2.2: It allows a different algorithm for each class $\mathcal{C}$, while the algorithm from Theorem 2.2 works for every class. In the same book [21], a second, stronger variant of fixed-parameter tractability is defined: a parameterized problem is *strongly uniform fixed-parameter tractable* if it is uniform fixed-parameter tractable and the function f is computable. This second definition is often used synonymous with the term *fixed-parameter tractable* (see e.g. the textbooks [21, 15, 34]).

We point out that our Theorem 2.2 does not establish *strongly* uniform fixed-parameter tractability for all monadically stable classes. The algorithm crucially exploits the fact that in every monadically stable class $\mathcal{C}$ for every radius r there is a bound game-depth$(\mathcal{C}, r)$ on how long it takes for Flipper to win the Flipper game. However, there are monadically stable classes $\mathcal{C}$, where there is no computable upper bound for the function game-depth$(\mathcal{C}, \cdot)$, as the following example shows.

Example 12.23. Consider the class $\mathcal{C}_\Sigma$ containing for every $k \geqslant 1$ the k-subdivided clique of order $\Sigma(k)$, where $\Sigma(\cdot)$ is the *busy beaver function*: $\Sigma(k)$ is the maximum number of 1s a Turing machine with k states and alphabet $\{0, 1\}$ can write and still halt [73]. The class $\mathcal{C}_\Sigma$ is monadically stable (and even nowhere dense). However, there is no computable function which upper bounds $\Sigma(\cdot)$, and one can verify that the same holds for the game-depth$(\mathcal{C}_\Sigma, \cdot)$.

The given example is quite artificial. Indeed, for all natural monadically stable classes $\mathcal{C}$ we know of, game-depth$(\mathcal{C}, \cdot)$ can be bounded by a computable function. In this case we also want the runtime of the algorithm to be bounded by a computable function. We take inspiration from Grohe, Kreutzer, and Siebertz, who faced the same issue in their model checking algorithm for nowhere dense classes. Their solution was to tie the computability of the runtime bound for their algorithm to the computability of certain nowhere denseness parameters. In order to do so, they introduced *effectively nowhere dense classes* [48]. We define those classes and contribute corresponding definitions for the stable and dependent case.

Definition 12.24. Let $f : \mathbb{N} \to \mathbb{N}$ be a function. A graph G is

- *f-nowhere-dense* if for every $k \in \mathbb{N}$:
  the k-subdivided clique of size $f(k)$ not a subgraph of G;
- *monadically f-stable* if for every formula $\varphi(x, y)$:
  the half-graph of order $f(|\varphi|)$ is not contained in $T_\varphi(G)$;
- *monadically f-dependent* if for every formula $\varphi(x, y)$:
  the powerset graph of order $f(|\varphi|)$ is not contained in $T_\varphi(G)$.

We call f a nowhere denseness / stability / dependence function of G. A graph class $\mathcal{C}$ is *f-nowhere-dense* / *monadically f-stable* / *monadically f-dependent* if every graph in $\mathcal{C}$ is. In this case we say $\mathcal{C}$ is *effectively nowhere-dense* / *effectively monadically stable* / *effectively monadically dependent*, if the function f can be chosen to be computable.

It is easy to see that these parameterized definitions match the unparameterized ones in the following sense: A graph class is nowhere dense if and only if it is f-nowhere-dense for some function f. The same equivalence holds for the stable and the dependent case. The goal of this section is to prove following extension of Theorem 2.2.

Theorem 12.25. *There is an algorithm that, given a graph G and a first-order sentence φ, decides whether $G \models \varphi$, and has the following property.*

For every monadically stable graph class $\mathcal{C}$, there exists a function $f : \mathbb{N} \times \mathbb{R} \to \mathbb{N}$ such that on any n-vertex graph $G \in \mathcal{C}$ and sentence φ the algorithm runs in time $f(|\varphi|, \varepsilon) \cdot n^{6+\varepsilon}$ for every $\varepsilon > 0$. If $\mathcal{C}$ is effectively monadically stable, then f is computable.

The statement implies strongly uniform fixed-parameter first-order model checking for every effectively monadically stable graph class.

We can simplify notation and avoid conditionals by extending our asymptotic notation. Let $p_1, \ldots, p_k$ be (not necessarily computable) functions on the rationals[1] with arbitrary arity. For each $t \in \mathbb{N}$, we denote by $\mathbf{C}^t_{p_1,\ldots,p_k}$ the set of functions $f : \mathbb{Q}^t \to \mathbb{Q}$ computable by a Turing machine with oracle access to $p_1, \ldots, p_k$. We write

- $U_{p_1,\ldots,p_k}(\cdot)$ to denote an anonymous, monotone, and unbounded function

$$f : \mathbb{N} \to \mathbb{N} \text{ with } f \in \mathbf{C}^1_{p_1,\ldots,p_k},$$

- $\mathrm{const}(p_1, \ldots, p_k)$ to denote an anonymous natural number $c \in \mathbf{C}^0_{p_1,\ldots,p_k}$,
- $O_{p_1,\ldots,p_k}(\cdot)$ to denote an anonymous function $f : \mathbb{N} \to \mathbb{N}$ upper bounded by

$$h(x) := \mathrm{const}(p_1, \ldots, p_k) \cdot x + \mathrm{const}(p_1, \ldots, p_k).$$

Using this notation, we strengthen Theorem 12.25 as follows.

Theorem 12.26. *There is an algorithm that, given a graph G and a first-order sentence φ, decides whether $G \models \varphi$, and has the following property.*

For every monadically f-stable graph G and sentence φ, the algorithm runs in time $O_{f,|\varphi|,\varepsilon}(n^{6+\varepsilon})$ for every $\varepsilon > 0$ and $n := |V(G)|$.

It follows from our asymptotic notation that if in the above f is computable, then the anonymous function bounding the factor of the runtime of the algorithm is also computable.

To prove Theorem 12.26, we have to show that the running time of the model checking algorithm depends in a computable way on the stability function f of the input graph. This is mostly straightforward but requires tracing the whole construction of this thesis (and additionally a few sources that were used). It would have been possible to do this analysis "in place". However, due to the notational overhead involved, we have decided to confine this finer-grained analysis of the construction to this section, where we will only sketch the process.

Computable Bounds for the Flipper Game

In this subsection we relate the bounds of the Flipper game winning strategy (Theorem 10.4) to the stability function of the graph on which the game is played, as follows.

Proposition 12.27. *There is a budget-2 Flipper strategy $\mathsf{flip}^\star$ with the following property.*

For every monadically f-stable graph G and radius $r \in \mathbb{N}$, $\mathsf{flip}^\star$ is $\mathrm{const}(f, r)$-winning and has runtime $O_{f,r}(n^2)$ in the radius-r game on G.

By definition, every monadically f-stable class is also monadically f-dependent. The bounds proven for the Flipper game in Chapter 10 crucially use the fact that monadically dependent classes have the insulation property. To preserve the dependence on f, we need a parameterized definition of the insulation property, similar to the definition of monadic f-stability.

[1] We restrict ourselves to rationals instead of reals, to ensure that every number has a finite representation that can be handled by a Turing machine.

Definition 12.28. Fix functions $f : \mathbb{N}^2 \to \mathbb{N}$ and $g : \mathbb{N} \to \mathbb{N}$ and a graph G. We say:

- G is *g-pattern-free*, if for every $r \geqslant 1$, G excludes as induced subgraphs
  - all flipped star r-crossing of order $g(r)$, or
  - all flipped clique r-crossing of order $g(r)$, or
  - all flipped half-graph r-crossing of order $g(r)$, or
  - the comparability grid of order $g(r)$.
- G is *f-prepattern-free*, if for every $h, k \in \mathbb{N}$, G contains no prepattern of order $f(h, k)$ on an insulator of cost at most k and height at most h in G.
- G has the *(f, g)-insulation-property*, if for every $r, m \in \mathbb{N}$, $W \subseteq V(G)$ with $|W| \geqslant f(r, m)$, there is a subset $W_\star \subseteq W$ of size at least m that is $(r, g(r))$-insulated in G.
- G is *(f, g)-flip-breakable*, if for every $r, m \in \mathbb{N}$, $W \subseteq V(G)$ with $|W| \geqslant f(r, m)$, there exist subsets $A, B \subseteq W$ with $|A|, |B| \geqslant m$ and a $g(r)$-flip H of G such that $\mathrm{dist}_H(A, B) > r$.

We show that the parameters for (almost) all the characterizations of monadic dependence given in Part II, are computable from each other. We leave out the characterization by efficient interpretations.

Lemma 12.29. *We have the following:*

1. *Every monadically f-dependent graph is g-pattern-free,*
   for some $g \in \mathbf{C}^1_f$ depending only on f.
2. *Every f-pattern-free graph is g-prepattern-free,*
   for some $g \in \mathbf{C}^2_f$ depending only on f.
3. *Every f-prepattern-free graph has the (g, h)-insulation-property,*
   for some $g \in \mathbf{C}^2_f$ and $h \in \mathbf{C}^1_f$ depending only on f.
4. *Every graph with the (f, g)-insulation-property is (h, i)-flip-breakable,*
   for some $h \in \mathbf{C}^2_{f,g}$ and $i \in \mathbf{C}^1_{f,g}$ depending only on f and g.
5. *Every (f, g)-flip-breakable graph is monadically h-dependent,*
   for some $h \in \mathbf{C}^1_{f,g}$ depending only on f and g.

Proof. We revisit the proofs from Part II.

1. We show that every monadically f-dependent graph is g-pattern-free for some $g \in \mathbf{C}^1_f$ depending only on f. It follows from the proof of Proposition 6.46 that there is a constant c such that for every $r \in \mathbb{N}$, there is a formula $\varphi_r(x, y)$ of length at most $c \cdot r$ with the following property. For every graph G containing a flipped r-crossing or comparability grid of order at least n, $T_{\varphi_r}(G)$ contains all bipartite graphs with sides of size at most $n - 1$. We can therefore set $g(r) := 2^{f(c \cdot r)} + 1 \in \mathbf{C}^1_f$. Suppose G is monadically f-dependent but not g-pattern-free. Then there exists some $r \in \mathbb{N}$ for which G contains a flipped r-crossing or comparability grid of order at least $2^{f(c \cdot r)} + 1$. This means a formula of length at most $c \cdot r$ transduces all bipartite graphs of size $2^{f(c \cdot r)} + 1 - 1$ from G. This includes the powerset graph of order $f(c \cdot r)$; a contradiction.
2. We show that every f-pattern-free graph is g-prepattern-free for some $g \in \mathbf{C}^2_f$ depending only on f. By Proposition 6.45, there is a function $p : \mathbb{N}^3 \to \mathbb{N}$ such that every graph containing a prepattern of order $p(h, k, m)$ on an insulator of height h and cost k contains an r-crossing or comparability grid of order m for some $r \geqslant 8h$. It follows that every f-pattern-free graph must be g-prepattern-free for $g(h, k) := p(h, k, f(8h))$. Inspecting the proofs of Chapter 6 reveals that the function p is computable: the crossings and comparability grids are extracted by repeatedly applying Bipartite Ramsey (Lemma 4.15), which yields computable bounds. We therefore have $g \in \mathbf{C}^1_f$.

3. We show that every f-prepattern-free graph has the (g, h)-insulation-property for some $g \in \mathbf{C}^2_f$ and $h \in \mathbf{C}^1_f$ depending only on f. The proof of Proposition 5.23 shows that g and h can be recursively constructed using f and the functions that bound the length and cost of the orderless and ordered insulator growing lemmas (Lemmas 5.19 and 5.20). An analysis of these lemmas (and their dependencies) reveals these functions to be computable. In particular, all the tools we use yield computable, bounds including Ramsey's Theorem (Fact 4.13), the Ramsey-type result for matchings, co-matchings, and half-graphs (Fact 5.49), and the pigeonhole principle. We therefore have $g \in \mathbf{C}^2_f$ and $h \in \mathbf{C}^1_f$.
4. Follows directly by Lemma 5.25.
5. Follows from the proof of Lemma 5.33 and the computable bounds from Lemma 4.10. □

We obtain the following corollary.

Corollary 12.30. *Every monadically f-stable graph has the (g, h)-insulation-property, for some $g \in \mathbf{C}^2_f$ and $h \in \mathbf{C}^1_f$ depending only on f.*

Using this corollary in the proof of Lemma 10.23, we obtain the following parameterized version of Lemma 10.23.

Lemma 12.31. *For every well-ordered, monadically f-stable graph G, $r \in \mathbb{N}$, and family $\mathcal{B}_0$ of pairwise disjoint r-balls in G, there exists a canonical classifier $\mathfrak{B} = (\mathcal{B} \subseteq \mathcal{B}_0, S, \mathrm{ex}, \mathrm{rep})$ of size $U_{f,r}(|\mathcal{B}_0|)$ and order $\mathrm{const}(f, r)$ in G.*

Similarly, we obtain a parameterized version of predictable flip-flatness (Proposition 10.7).

Lemma 12.32. *There is an algorithm that takes as input $r, k \in \mathbb{N}$, a well-ordered graph G, and a size five set $Z \subseteq V(G)$, and computes in time $O_{r,k}(|V(G)|^2)$ a k-flip $\mathrm{Predict}(r, k, G, Z)$ of G with the following properties:*

For every function $f : \mathbb{N} \to \mathbb{N}$ and radius $r \in \mathbb{N}$ there is a bound $k_{f,r} \leqslant \mathrm{const}(f, r)$ and functions $\mathrm{Flip}_{f,r}$ and $\mathrm{Flat}_{f,r}$ such that for every well-ordered, monadically f-stable graph G, sets $X, Z \subseteq V(G)$ and parameter $k \geqslant k_{f,r}$ we have

- $\mathrm{Flip}_{f,r}(G, X)$ *is a $k_{f,r}$-flip of G,*
- $\mathrm{Flat}_{f,r}(G, X)$ *is a size $U_{f,r}(|X|)$ subset of X,*
- $\mathrm{Flat}_{f,r}(G, X)$ *is distance-r independent in* $\mathrm{Flip}_{f,r}(G, X)$*, and*
- *if Z is a size 5 subset of* $\mathrm{Flat}_{f,r}(G, X)$ *then* $\mathrm{Predict}(r, k, G, Z) = \mathrm{Flip}_{f,r}(G, X)$.

It is now straightforward to derive Proposition 12.27, which we restate here for convenience.

Proposition 12.27. *There is a budget-2 Flipper strategy* $\mathsf{flip}^\star$ *with the following property.*

For every monadically f-stable graph G and radius $r \in \mathbb{N}$, $\mathsf{flip}^\star$ *is $\mathrm{const}(f, r)$-winning and has runtime $O_{f,r}(n^2)$ in the radius-r game on G.*

Computable Overlap for Neighborhood Covers

In this subsection we relate the overlap of the neighborhood covers to the stability function of the graph as follows.

Proposition 12.33. *There is an algorithm that, given an n-vertex graph G and a radius $r \in \mathbb{N}$, computes a distance-r neighborhood cover of G with diameter at most $4r$ in time $O(n^5)$.*

If G is monadically f-stable, then the overlap of the cover is bounded by $O_{f,r,\varepsilon}(n^\varepsilon)$ for every $\varepsilon > 0$.

In the proof of Theorem 11.2, the overlap of the neighborhood cover depends in a computable way on the neighborhood complexity of the rth power G^r of the input graph G. As the rth power of a graph is generated by a transduction of length $O(r)$, the following lemma can be used to lift the computability of the bound on the neighborhood complexity from G to G^r.

Lemma 12.34. *For every monadically f-stable (f-dependent) graph G and formula $\varphi(x, y)$, every graph in $T_\varphi(G)$ is monadically g-stable (g-dependent) for some $g \in \mathbf{C}^1_{f,|\varphi|}$ which depends only on f and $|\varphi|$.*

Proof. Follows by the construction of Fact 4.4, which shows that transductions are transitive. □

Therefore, in order to prove Proposition 12.33, it suffices to link the neighborhood complexity to the stability function of G as follows.

Lemma 12.35. *For every monadically f-stable G, $\varepsilon > 0$, and $A \subseteq V(G)$, we have*

$$|\{N[v] \cap A : v \in V(G)\}| \leqslant O_{f,\varepsilon}(|A|^{1+\varepsilon}).$$

We first consult the literature to verify that the neighborhood complexity bounds of a graph G depends on the nowhere denseness function of G in a computable way. The following is a parameterized version of Fact 11.4.

Lemma 12.36. *For every f-nowhere-dense graph G, $\varepsilon > 0$, and $A \subseteq V(G)$, we have*

$$|\{N[v] \cap A : v \in V(G)\}| \leqslant O_{f,\varepsilon}(|A|^{1+\varepsilon}).$$

Proof. It is proven in [23, Lem 4.7] that we have

$$|\{N[v] \cap A : v \in V(G)\}| \leqslant f_{\text{nei}}(\varepsilon) \cdot |A|^{1+\varepsilon}$$

for some function f_{nei} constructed from

- the aforementioned nowhere density function f (denoted as f_ω in [23]),
- a function $f_\nabla : \mathbb{N} \times \mathbb{Q} \to \mathbb{N}$ bounding the edge density of *shallow minors*, and
- a function $f_{\text{wcol}} : \mathbb{N} \times \mathbb{Q} \to \mathbb{N}$ bounding the *weak coloring numbers* of G.

By inspecting the proofs of [61, Thm. 3.2] and [29, Thm. 3.15] we conclude that $f_\nabla \in \mathbf{C}^2_f$. From [86, Lem. 3.4 and Cor. 3.5] we get that $f_{\text{wcol}} \in \mathbf{C}^2_{f_\nabla} \subseteq \mathbf{C}^2_f$. Hence, $f_{\text{nei}} \in \mathbf{C}^1_f$. □

To make use of this statement, we also need the following:

Lemma 12.37. *Every monadically f-dependent graph that excludes a biclique of size k as a subgraph, is g-nowhere-dense for some $g \in \mathbf{C}^1_{f,k}$ depending only on f and k.*

Proof. By the proof of the upcoming Lemma 13.7 and Lemma 12.29. □

Using Lemmas 12.34, 12.36 and 12.37 in Section 11.1, we obtain the parameterized almost linear neighborhood complexity for monadically f-stable graphs (Lemma 12.35). As we argued before, this is sufficient to prove the existence of the desired neighborhood covers (Proposition 12.33).

Computable Bounds for the Model Checking Algorithm

We have linked the bounds of the Flipper game (Proposition 12.27) and the overlap of the neighborhood covers (Proposition 12.33) to the stability function of the input graph.

By tracing the algorithm presented in Section 12.3, it is now straightforward to prove Theorem 12.26, which we restate here for convenience.

Theorem 12.26. *There is an algorithm that, given a graph G and a first-order sentence φ, decides whether $G \models \varphi$, and has the following property.*

For every monadically f-stable graph G and sentence φ, the algorithm runs in time $O_{f,|\varphi|,\varepsilon}(n^{6+\varepsilon})$ for every $\varepsilon > 0$ and $n := |V(G)|$.

In particular, for every monadically f-stable graph and formula φ with quantifier rank q, the recursion depth of the algorithm is bounded by $\mathrm{const}(f, q)$. Note that during the algorithm, neighborhood covers are constructed for flips of the input graph G. Again, as flips are expressible by transductions, the computability of the bounds can be lifted from G using Lemma 12.34.

This concludes the sketch of the proof of Theorem 12.25, which we restate here.

Theorem 12.25. *There is an algorithm that, given a graph G and a first-order sentence φ, decides whether $G \models \varphi$, and has the following property.*

For every monadically stable graph class $\mathcal{C}$, there exists a function $f : \mathbb{N} \times \mathbb{R} \to \mathbb{N}$ such that on any n-vertex graph $G \in \mathcal{C}$ and sentence φ the algorithm runs in time $f(|\varphi|, \varepsilon) \cdot n^{6+\varepsilon}$ for every $\varepsilon > 0$. If $\mathcal{C}$ is effectively monadically stable, then f is computable.

Part IV

The Breakability Framework

Outline Part IV

In this part, we show that natural restrictions of flip-flatness and flip-breakability can be used to characterize nowhere denseness, bounded clique- and tree-width, and bounded shrub- and tree-depth. The following theorem summarizes our results.

Theorem 2.5. *For every graph class $\mathcal{C}$, the following holds.*

(1) $\mathcal{C}$ is flip-breakable	*if and only if it is*	*monadically dependent.*
(2) $\mathcal{C}$ is flip-flat	*if and only if it is*	*monadically stable.*
(3) $\mathcal{C}$ is deletion-breakable	*if and only if it is*	*nowhere dense.*
(4) $\mathcal{C}$ is deletion-flat	*if and only if it is*	*nowhere dense.*
(5) $\mathcal{C}$ is dist.-∞ flip-breakable	*if and only if it has*	*bounded clique-width.*
(6) $\mathcal{C}$ is dist.-∞ flip-flat	*if and only if it has*	*bounded shrub-depth.*
(7) $\mathcal{C}$ is dist.-∞ deletion-breakable	*if and only if it has*	*bounded tree-width.*
(8) $\mathcal{C}$ is dist.-∞ deletion-flat	*if and only if it has*	*bounded tree-depth.*

In Parts II and III, we have already shown the equivalences (1) and (2) (Theorems 2.1 and 2.3). In Chapter 13 we study nowhere dense classes and show equivalence (3). Additionally, we show that many of the known results connecting nowhere denseness and monadic dependence can now easily be proved from our results in Parts II and III. In particular, we give a new proof of the uniform quasi-wideness characterization [16, 61, 60] that uses flip-flatness as a black box. This corresponds to the equivalence (4) in the above theorem.

The remaining equivalences (5) to (8) are shown in Chapters 14 to 17.

Chapter 13

Nowhere Denseness

Definition 13.1. A graph class $\mathcal{C}$ is *nowhere dense* if for every radius $r \in \mathbb{N}$, there exists a bound $N_r \in \mathbb{N}$ such that no graph from $\mathcal{C}$ contains an r-subdivided clique of order N_r as a subgraph.

Nowhere dense classes act as the appropriate restriction of monadic stability and dependence to sparse graph classes. They have been extensively studied and many different characterizations are known. Building on our characterizations for monadic stability and dependence, we (re)prove the following characterizations of nowhere dense classes that highlight their connection with stability and dependence.

Theorem 13.2. *Let $\mathcal{C}$ be a graph class. The following are equivalent.*

(1) $\mathcal{C}$ is nowhere dense.

(2) $\mathcal{C}$ is monadically stable and weakly sparse.

(3) $\mathcal{C}$ is monadically dependent and weakly sparse.

(4) The monotone closure of $\mathcal{C}$ is stable.

(5) The monotone closure of $\mathcal{C}$ is dependent.

(6) $\mathcal{C}$ is deletion-flat (i.e. uniformly quasi-wide).

(7) $\mathcal{C}$ is deletion-breakable.

The equivalence (1) $\Leftrightarrow$ (2) $\Leftrightarrow$ (3) is implied by [30, Thm. 6]. For the sake of completeness, we give a standalone proof here. The equivalence (1) $\Leftrightarrow$ (4) $\Leftrightarrow$ (5) was proven in [72, 1]. We show that it also follows from our characterizations of monadic stability and dependence. The notion of *deletion-flatness* was introduced in [16] under the name *uniform quasi-wideness*, and the equivalence (1) $\Leftrightarrow$ (6) was proven in [60, 61]. We give a new proof, which remarkably uses the flip-flatness characterization for monadically stable classes as a black box. We introduce the new notion of *deletion-breakability* as a sparse counterpart to *flip-breakability* and show the equivalence (1) $\Leftrightarrow$ (7). We define the involved notions.

Definition 13.3 (Deletion-Flatness)**.** A graph class $\mathcal{C}$ is *deletion-flat* if for every radius $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $k_r \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N_r(m)$ there exist sets $S \subseteq V(G)$ with $|S| \leqslant k_r$ and $A \subseteq W \setminus S$ with $|A| \geqslant m$ such that for every two distinct vertices $u, v \in A$:

$$\mathrm{dist}_{G-S}(u, v) > r.$$

Definition 13.4 (Deletion-Breakability)**.** A graph class $\mathcal{C}$ is *deletion-breakable*, if for every radius $r \in \mathbb{N}$ there exists a function $N_r : \mathbb{N} \to \mathbb{N}$ and a constant $k_r \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N_r(m)$ there exist a set $S \subseteq V(G)$ with $|S| \leqslant k_r$ and subsets $A, B \subseteq W \setminus S$ with $|A|, |B| \geqslant m$ such that

$$\mathrm{dist}_{G-S}(A, B) > r.$$

Definition 13.5 (Weakly sparse classes)**.** A graph class $\mathcal{C}$ is *weakly sparse*, if there exists a bound k, such that no graph from $\mathcal{C}$ contains the biclique of order k as a subgraph.

We show the equivalences (1) ⇔ (2) ⇔ (3) and (1) ⇔ (4) ⇔ (5) and (1) ⇔ (6) ⇔ (7) separately.

Weakly Sparse Monadic Stability and Dependence

Lemma 13.6. *Every nowhere dense class $\mathcal{C}$ is weakly sparse and monadically stable.*

Proof. It is easy to see that $\mathcal{C}$ is weakly sparse. Assume towards a contradiction that $\mathcal{C}$ is not monadically stable. Then by Theorem 2.1, there is some $r \in \mathbb{N}$ such that $\mathcal{C}$ contains as induced subgraphs flipped *star r-crossings / clique r-crossings / half-graphs* of arbitrarily order n for any $n \in \mathbb{N}$. For the last two patterns we get a contradiction, as they contain bicliques of order $U(n)$ as subgraphs, and are therefore not weakly sparse. For the flipped star r-crossings of order n, there are three cases:

- No layers of the star r-crossings were flipped: Then it contains a $(2r + 1)$-subdivided clique of order $\lfloor \sqrt{n} \rfloor$ as an induced subgraph; a contradiction to $\mathcal{C}$ being nowhere dense.
- A layer of the star r-crossings was flipped with itself: Then the layer forms a clique of size at least n; a contradiction to $\mathcal{C}$ being weakly sparse.
- Two distinct layers of the star r-crossings were flipped with each other: We find a biclique of order at least $\lfloor n/2 \rfloor$ as a subgraph between the two layers; a contradiction to $\mathcal{C}$ being weakly sparse. □

Lemma 13.7. *Every weakly sparse, monadically dependent graph class is nowhere dense.*

The lemma will follow easily from the following Ramsey-type result.

Lemma 13.8. *Every graph G that contains an r-subdivided biclique of order n as a subgraph contains either a (non-subdivided) biclique order $U_r(n)$ as a subgraph, or an r'-subdivided biclique of order $U_r(n)$ for some $r' \in [r]$ as an induced subgraph.*

Proof. If $r = 0$, then we trivially have a biclique of order n as a subgraph. So assume $r \geqslant 1$. Let $\{a_i : i \in [n]\}$ and $\{b_i : i \in [n]\}$ be the principal vertices of the r-subdivided biclique in G and let $\{p_{i,j,k} : i, j \in [n], k \in [r]\}$ be the subdivision vertices, such that $\bar{p}_{i,j} = (a_i, p_{i,j,1}, \ldots, p_{i,j,r}, b_j)$ is an $(r + 2)$-vertex path in G and the inner vertices of all the $p_{i,j}$ are distinct. By Bipartite Ramsey (Lemma 4.15), there exists sets $I, J \subseteq [n]$ of size $U_r(n)$ such that for all $i, i' \in I$ and $j, j' \in J$

$$\mathrm{atp}_G(\bar{p}_{i,j}, \bar{p}_{i',j'}) \text{ depends only on } \mathrm{otp}(i, i') \text{ and } \mathrm{otp}(j, j'). \tag{$*$}$$

Both $\{a_i : i \in I\}$ and $\{b_j : j \in J\}$ are either a clique or an independent set. In the absence of large semi-induced bicliques, we can therefore assume.

$$\text{Both } \{a_i : i \in I\} \text{ and } \{b_j : j \in J\} \text{ are independent sets.} \tag{X}$$

If there is an edge between a_i and b_j for any $i \in I$ and $j \in J$, then by (∗), G contains a large semi-induced biclique on the vertices $\{a_i : i \in I\}$ and $\{b_j : j \in J\}$ and we are done. Otherwise, there exist $r' \in [r]$ and numbers $1 \leqslant k_1 < \ldots < k_{r'} \leqslant r$ such that for all $i \in I$ and $j \in J$

$$\bar{q}_{i,j} = (a_i, q_{i,j,1} := p_{i,j,k_1}, \ldots, q_{i,j,r'} := p_{i,j,k_{r'}}, b_j)$$

is a shortest path between a_i and b_j in $G[\bar{p}_{i,j}]$. The vertices $\bigcup_{i \in I, j \in J} \bar{q}_{i,j}$ form a (supergraph of) an r'-subdivided biclique in G with the following property.

$$\bar{q}_{i,j} \text{ is an induced path in } G \text{ for all } i \in I, j \in J. \qquad \text{(X)}$$

Now assume there are indices $i, i' \in I$, $j, j' \in J$, $k, k' \in [r']$ with $(i, j) \neq (i', j')$ such that $q_{i,j,k}$ and $q_{i',j',k'}$ are adjacent. By symmetry and up to swapping the roles of I and J, we can assume that $i < i'$. Let $I_<$ and $I_>$ be the first and last $\lfloor |I|/2 \rfloor$ elements of I, and define $J_<$ and $J_>$ in the same way for J. By (∗), the vertices $\{q_{i,j,k} : i \in I_<\}$ and $\{q_{i,j,k'} : i \in I_>\}$ form a large semi-induced biclique in G, and we are done. We can therefore assume the following.

$$q_{i,j,k} \text{ and } q_{i',j',k'} \text{ are non-adjacent } \forall i, i' \in I, j, j' \in J, k, k' \in [r'] \text{ with } (i, j) \neq (i', j'). \qquad \text{(X)}$$

In the same way, we can show that in the absence of large bicliques we have the following.

$$a_i \text{ and } q_{i',j',k'} \text{ are non-adjacent } \forall i, i' \in I, j' \in J, k' \in [r'] \text{ with } i \neq i'. \qquad \text{(X)}$$

$$b_j \text{ and } q_{i',j',k'} \text{ are non-adjacent } \forall i' \in I, j, j' \in J, k' \in [r'] \text{ with } j \neq j'. \qquad \text{(X)}$$

Combining the facts marked with (X) proves that, in the absence of a large semi-induced biclique, there is a large induced r'-subdivided biclique with principal vertices $\{a_i : i \in I\}$ and $\{b_j : j \in J\}$ in G. This concludes the proof. □

Monotone Stability and Dependence

Lemma 13.9. *Let $\mathcal{C}$ be a graph class that is not nowhere dense. Then the monotone closure of $\mathcal{C}$ is independent.*

Proof. By Lemma 13.7, $\mathcal{C}$ is either not weakly sparse or monadically independent. In the first case, the monotone closure of $\mathcal{C}$ contains all bipartite graphs and is therefore independent. Otherwise, $\mathcal{C}$ is monadically independent, and already the hereditary closure of $\mathcal{C}$ interprets the class of all graphs by Theorem 2.3, and is therefore independent. □

Deletion-Flatness and Deletion-Breakability

Lemma 13.10. *Every weakly sparse, flip-flat graph class $\mathcal{C}$ is deletion-flat.*

Proof. Since $\mathcal{C}$ is weakly sparse, there is $\ell \in \mathbb{N}$ such that no graph from $\mathcal{C}$ contains a biclique of order ℓ as a subgraph. Let $G \in \mathcal{C}$ and $W \subseteq V(G)$. By flip-flatness, there is a k-flip H for some $k \leqslant \text{const}(\mathcal{C}, r)$ and a set $A \subseteq W$ of size $U_{\mathcal{C},r}(|W|)$ whose vertices have pairwise disjoint r-neighborhoods in H. We will show that there is an induced subgraph G' of G, obtained by deleting at most $k \cdot \ell$ vertices, in which the r-neighborhoods of the vertices in A are pairwise disjoint, too. Since there are no restrictions on the choice of G and W, this will prove the lemma.

Let $\mathcal{P} \subseteq 2^{V(G)}$ and $F \subseteq \mathcal{P}^2$ be the partition of size k and symmetric relation witnessing that H is a k-flip of G. We can assume $|A| \geqslant 2\ell$ and additionally by the pigeonhole principle:

1. Every part $P \in \mathcal{P}$ intersects $N_{r-1}[v]$ either for all or for no $v \in A$.
2. There is a single part in $\mathcal{P}$ that contains all of A.

Claim 13.11. *Let $(P, Q) \in F$ be a pair of flipped parts, such that P overlaps with $N^H_{r-1}[v]$ for some vertex $v \in A$. Then Q has size less than ℓ.*

Proof. As argued before, P overlaps with $N^H_{r-1}[v]$ for every $v \in A$. We choose a set $P' \subseteq P$ to contain exactly one vertex from $N^H_{r-1}[v]$ for every $v \in A$. Assume towards a contradiction that Q has size at least ℓ. Because the r-neighborhoods around A in H are disjoint, each vertex in Q can be adjacent to at most one vertex from P' in H. Since $(P, Q) \in F$ this means that in G, each vertex in Q is adjacent to all but at most one vertex from P' in G. Let $Q_\star$ be a size ℓ subset of Q. Let $P_\star \subseteq P'$ be the set obtained by removing for each $v \in Q_\star$ the at most one vertex in P' that is not adjacent to v in G. Since $|P'| \geqslant 2\ell$, we have $|P_\star| \geqslant \ell$. As $Q_\star$ and $P_\star$ are fully adjacent in G, G must contain a biclique of order ℓ as a subgraph; a contradiction. ■

Let G' be the induced subgraph of G obtained by deleting every part $Q \in \mathcal{P}$ of size at most ℓ. In total, we delete at most $k \cdot \ell$ vertices, as desired. Note that no vertex from A was deleted since A is completely contained inside a single part of size greater than ℓ. For a graph G, $v \in V(G)$, $i \in \mathbb{N}$, we denote by $X^G_i(v)$ the set of vertices of distance exactly i from v in G.

Claim 13.12. *$X^{G'}_i(v) \subseteq X^H_i(v)$ for every $v \in A$ and $0 \leqslant i \leqslant r$.*

Proof. We prove the claim by induction on i. The base case $r = 0$ is vacuously true. For the inductive step with $0 < i \leqslant r$, assume towards a contradiction that there is a vertex $u \in X^{G'}_i[v] - X^H_i[v]$. This is witnessed by an edge $uw \in E(G') \subseteq E(G)$ with $w \in X^{G'}_{i-1}[v]$. By induction also $w \in X^H_{i-1}[v]$, but $uw \notin E(H)$. This means there exists a flipped pair of parts $(P, Q) \in F$ with $w \in P$ and $u \in Q$. As P overlaps with $N^H_{i-1}(v)$ in i, the part Q must have size less than ℓ by Claim 13.11. The part Q was therefore deleted during the construction of G'. A contradiction to the assumption that u is contained in G'. ■

The claim implies that the r-neighborhoods of the vertices in A are disjoint in G'. □

Proof of Lemma 13.7. Assume $\mathcal{C}$ is monadically dependent and weakly sparse. Suppose $\mathcal{C}$ is not nowhere dense. Then there is some $r \in \mathbb{N}$ such that $\mathcal{C}$ contains arbitrarily large r-subdivided cliques as subgraphs. It is easy to see that $\mathcal{C}$ must also contain arbitrarily large r-subdivided bicliques as subgraphs. Then by Lemma 13.8 (and the pigeonhole principle) $\mathcal{C}$ also contains arbitrarily large bicliques as subgraphs or arbitrarily r'-subdivided as induced subgraphs. In the first case, we get a contradiction to the assumption that $\mathcal{C}$ is weakly sparse. In the second case, we get a contradiction to the assumption that $\mathcal{C}$ is monadically dependent: r'-subdivided cliques are the same as star r'-crossings, which are forbidden in monadically dependent classes by Theorem 2.3. □

Lemma 13.13. *Every deletion-flat class is deletion-breakable.*

Proof. By deletion-flatness, in every huge set, we find a large set of vertices pairwise of distance greater than r after removing few vertices. We can partition them into two halves A and B to obtain deletion-breakability. □

Lemma 13.14. *Every deletion-breakable class is nowhere dense.*

Proof. Assume towards a contradiction that $\mathcal{C}$ is not nowhere dense but deletion-breakable with bounds $N_r(\cdot)$ and k_r for every $r \in \mathbb{N}$. By definition, there exists a radius $r > 1$ such that $\mathcal{C}$ contains arbitrarily large $(r-1)$-subdivided cliques as subgraphs. Let $G \in \mathcal{C}$ be a graph containing an $(r-1)$-subdivided clique of size $N_r(k_r + 1)$, whose principal vertices we denote with W. By deletion-breakability, W contains two subsets A and B, each of size $k_r + 1$, such that

$$\mathrm{dist}_{G-S}(A, B) > r$$

for some vertex set S of size at most k_r. Since W is an $(r-1)$-subdivided clique in G, there exist $k_r + 1$ disjoint paths of length r, that each start in A and end in B. As $|S| \leqslant k_r$, at most one of those paths must survive in $G - S$, witnessing that $\mathrm{dist}_{G-S}(A, B) \leqslant r$; a contradiction. □

Wrapping Up

We are finally ready to prove Theorem 13.2, which we restate for convenience.

Theorem 13.2. *Let $\mathcal{C}$ be a graph class. The following are equivalent.*

(1) $\mathcal{C}$ is nowhere dense.

(2) $\mathcal{C}$ is monadically stable and weakly sparse.

(3) $\mathcal{C}$ is monadically dependent and weakly sparse.

(4) The monotone closure of $\mathcal{C}$ is stable.

(5) The monotone closure of $\mathcal{C}$ is dependent.

(6) $\mathcal{C}$ is deletion-flat (i.e. uniformly quasi-wide).

(7) $\mathcal{C}$ is deletion-breakable.

Proof. We collect the lemmas proven so far.

The implications (1) ⇒ (2) ⇒ (3) ⇒ (1) are given by Lemma 13.6, the fact that monadic stability implies monadic dependence, and Lemma 13.7.

The implications (1) ⇒ (4) ⇒ (5) ⇒ (1) are proven as follows. If $\mathcal{C}$ is nowhere dense, then by definition, also the monotone closure of $\mathcal{C}$ is nowhere dense. By the previously shown implications, this monotone closure is also monadically stable and in particular stable. Stability then implies dependence. Finally, if the monotone closure of $\mathcal{C}$ is dependent, then $\mathcal{C}$ must be nowhere dense by Lemma 13.9.

The implications (1) ⇒ (2) ⇒ (6) ⇒ (7) ⇒ (1) are given by Lemmas 13.6, 13.10, 13.13 and 13.14, where we additionally use the fact that monadically stable classes are flip-flat (Theorem 2.1). □

Chapter 14

Bounded Clique-Width

Definition 14.1. A graph class $\mathcal{C}$ is *distance-∞ flip-breakable*, if there exists a function $N : \mathbb{N} \to \mathbb{N}$ and a constant $k \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N(m)$ there exist subsets $A, B \subseteq W$ with $|A|, |B| \geqslant m$ and a k-flip H of G such that in H, no two vertices $a \in A$ and $b \in B$ are in the same connected component.

The goal of this chapter is to prove the following.

Theorem 14.2. *A graph class has bounded clique-width if and only if it is distance-∞ flip-breakable.*

To prove the theorem, we work with *rank-width*, a parameter that is functionally equivalent to clique-width. A graph G has rank-width at most k if there is a tree T whose leaves are the vertices of G, and inner nodes have degree at most 3, such that for every edge e of the tree, the bipartition $A \uplus B$ of the leaves of T into the leaves on either side of e, has *cut-rank* at most k. The cut-rank of a bipartition $A \uplus B$ of the vertex set of a graph G, denoted $\mathrm{rk}_G(A, B)$, is defined as the rank, over the two-element field, of the $(0, 1)$-matrix with rows A and columns B, where the entry at row $a \in A$ and column $b \in B$ is 1 if $ab \in E(G)$ and 0 otherwise.

Fact 14.3 ([51, Proposition 6.3]). *A graph class has bounded clique-width if and only if it has bounded rank-width.*

Lemma 14.4. *Let T be a rooted subtree of a binary tree and let W be a subset of the leaves of T. There exists an edge $e \in E(T)$ such that the two subtrees T_1 and T_2 obtained by removing e from T each contain at least $\frac{1}{4}|W|$ vertices from W.*

Proof. For a vertex $v \in V(T)$, denote by $T(v)$ the subtree rooted at v. Let $(v_1, \ldots, v_m)$ be a root-to-leaf path in T such that for all $1 \leqslant i < m$, the vertex v_{i+1} is the child of v_i whose subtree contains the most elements from W, where ties are broken arbitrarily. Let $i \in [m]$ be the largest index such that $T(v_i)$ contains at least $\frac{1}{4}|W|$ vertices from W. By construction, $T(v_2)$ contains at least $\frac{1}{2}|W|$ elements from W, and therefore $i > 1$, i.e. v_i has a parent v_{i-1}. $T(v_i)$ contains less than $\frac{1}{2}|W|$ vertices from W, as both of its at most two children contain less than $\frac{1}{4}|W|$ elements from W. Therefore, the edge connecting v_i and v_{i-1} is the desired edge. □

Lemma 14.5. *Every graph class with bounded rank-width is distance-∞ flip-breakable.*

Proof. Fix a number k and let $\mathcal{C}$ be a graph class of rank-width at most k. We will show that $\mathcal{C}$ is distance-∞ flip-breakable for $N(m) := 4m$ using $2^k + 2^{2^k}$ flips. For every graph $G \in \mathcal{C}$ there is a rooted subtree T of a binary tree with leaves $V(G)$, such that for every edge $e \in E(T)$, the bipartition $X \uplus Y$ of the leaves of T into the leaves on either side of e, has cut-rank at most k. Let $W \subseteq V(G)$ be a set of size $4m$. By Lemma 14.4, there exists an edge e such that in the corresponding bipartition $X \uplus Y$ of $V(G)$, both X and Y each contain at least m many vertices of W. Observe that since $\mathrm{rk}_G(X,Y) \leqslant k$, X induces at most 2^k distinct neighborhoods over Y. Then there is a $(2^k + 2^{2^k})$-flip H of G in which there are no edges between X and Y: the corresponding partition of $V(G)$ partitions the vertices of X into 2^k parts depending on their neighborhood in Y and partitions the vertices of Y in 2^{2^k} parts depending on their neighborhood in X. □

A set W of vertices of G is *well-linked*, if for every bipartition $A \uplus B$ of $V(G)$, the cut-rank of $A \uplus B$ satisfies $\mathrm{rk}_G(A,B) \geqslant \min(|A \cap W|, |B \cap W|)$. We use the following two facts.

Fact 14.6 ([51, Thm. 5.2])**.** *Every graph of rank-width greater than k contains well-linked set of size k.*

Fact 14.7 ([83, Lem. D.2])**.** *Let G be a graph and $A \uplus B$ a bipartition of $V(G)$ with $\mathrm{rk}_G(A,B) > k$. Then for every k-flip H of G there is some edge $ab \in E(H)$ with $a \in A$ and $b \in B$.*

Lemma 14.8. *Every graph class with unbounded rank-width is not distance-∞ flip-breakable.*

Proof. Let $\mathcal{C}$ be a graph class with unbounded rank-width. Assume towards a contradiction that $\mathcal{C}$ is distance-∞ flip-breakable with bounds $N(\cdot)$ and k. By Fact 14.6, there exists a graph $G \in \mathcal{C}$ that contains a well-linked set W of size at least $N(k+1)$. By distance-∞ flip-breakability, there exists a k-flip H of G and two sets $A, B \subseteq W$ of size $k+1$ each, such that no two vertices $a \in A$ and $b \in B$ are in the same component in H. We can therefore find a bipartition of $\mathcal{X} \uplus \mathcal{Y}$ of the connected components of H, such that $\mathcal{X}$ contains all the components containing a vertex of A and $\mathcal{Y}$ contains all the components containing a vertex of B. Components containing neither a vertex of A nor of B can be distributed arbitrarily among $\mathcal{X}$ and $\mathcal{Y}$. Let $X := \bigcup \mathcal{X}$ and $Y := \bigcup \mathcal{Y}$. Then $X \uplus Y$ is a bipartition of $V(H)$, and there is no edge between X and Y in H. Since W is well-linked we have that the cut-rank $X \uplus Y$ satisfies

$$\mathrm{rk}_G(X,Y) \geqslant \min(|X \cap W|, |Y \cap W|) \geqslant \min(|A|, |B|) = k+1.$$

By Fact 14.7, there must be an edge $ab \in E(H)$ with $a \in X$ and $b \in Y$; a contradiction. □

Combining Fact 14.3, Lemma 14.5, and Lemma 14.8 now yields Theorem 14.2.

Chapter 15

Bounded Tree-Width

Definition 15.1. A graph class $\mathcal{C}$ is *distance-∞ deletion-breakable*, if there exists a function $N : \mathbb{N} \to \mathbb{N}$ and a constant $k \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N(m)$ there exist a set $S \subseteq V(G)$ with $|S| \leqslant k$ and subsets $A, B \subseteq W \setminus S$ with $|A|, |B| \geqslant m$ such that in $G - S$, no two vertices $a \in A$ and $b \in B$ are in the same connected component.

The goal of this chapter is to prove the following.

Theorem 15.2. *A graph class has bounded tree-width if and only if it is distance-∞ deletion-breakable.*

We start with the forward direction. We assume familiarity with tree-width and (nice) tree decompositions. See for example [15] for an introduction.

Lemma 15.3. *Every graph class with bounded tree-width is distance-∞ deletion-breakable.*

Proof. Let $\mathcal{C}$ be a graph class of tree-width at most $k - 1$. We show that $\mathcal{C}$ is distance-∞ deletion-breakable with bounds $N(m) := 4(m + k)$ and $k := k$.

Consider a graph $G \in \mathcal{C}$ of tree-width at most $k - 1$ and a subset W containing at least $N(m) = 4(m + k)$ vertices of G. We fix a nice tree decomposition of G of width k, and associate with every bag t the set $V(t) \subseteq V(G)$ consisting of all vertices that are contained either in t or in a descendant of t. We consider a walk that starts at the root of the nice tree decomposition and walks downwards towards the leaves. Whenever the walk reaches a join node, it proceeds towards the child t that maximizes $|V(t) \cap W|$. As each node has at most two children, the cardinality $|V(t) \cap W|$ of the current node t of the walk can decrease at most by a factor $\frac{1}{2}$ with each step. Hence, as in the proof of Lemma 14.4, we reach at some point a node t with

$$m + k \leqslant \frac{1}{4}|W| \leqslant |V(t) \cap W| < \frac{1}{2}|W| \leqslant 2(m + k).$$

Let $S \subseteq V(G)$ be the vertices in this bag t and choose $A = (V(t) \cap W) \setminus S$ and $B = W \setminus V(t)$. Note that $|S| \leqslant k$ and $|A|, |B| \geqslant m$. By the definition of tree decompositions, the vertices S act as a separator in the desired sense: In $G - S$, no two vertices $a \in A$ and $b \in B$ are in the same connected component. $\square$

The backwards direction will follow easily from the Grid-Minor theorem by Robertson and Seymour. First some notation. A graph H is a *minor* of a graph G if there exists a *minor model* μ of H in G. A minor model is a map μ that assigns to every vertex $v \in V(H)$ a connected subgraph $\mu(v)$ of G and to every edge $e \in E(H)$ an edge $\mu(e) \in E(G)$ satisfying

- for all $u, v \in V(H)$ with $u \neq v$: $V(\mu(u)) \cap V(\mu(v)) = \emptyset$;
- for every $(u, v) \in E(H)$: $\mu((u, v)) = (u', v')$ for vertices $u' \in V(\mu(u))$ and $v' \in V(\mu(v))$.

Possibly deviating from our notation in previous sections, in this section the *k-grid* is the graph on the vertex set $[k] \times [k]$ where two vertices (i, j) and (i', j') are adjacent if and only if $|i - i'| + |j - j'| = 1$. We can now state the Grid-Minor theorem.

Fact 15.4 ([75, Thm. 1.5])**.** *Let $\mathcal{C}$ be a graph class with unbounded tree-width. Then for every $k \in \mathbb{N}$, $\mathcal{C}$ contains a graph which contains the k-grid as a minor.*

Lemma 15.5. *Every graph class with unbounded tree-width is not distance-∞ deletion-breakable.*

Proof. Assume towards a contradiction that $\mathcal{C}$ has unbounded tree-width but is distance-∞ deletion-breakable with bounds $N(\cdot)$ and k. Let $t := N(2k + 2)$. By Fact 15.4, there is a graph $G \in \mathcal{C}$ such that there exists a minor model μ of the t-grid in G. Let $W := \{v_1, \ldots, v_t\} \subseteq V(G)$ be a set *representing* the bottom row of the t-grid: we pick one vertex v_i from the subgraph $\mu((i, 1))$ for each $i \in [t]$. We apply distance-∞ flip-breakability to the set W in G, which yields a set $S \subseteq V(G)$ of size k and disjoint sets $A, B \subseteq W \setminus S$, each of size $2k + 2$, such that in $G - S$ there is no path from a vertex in A to a vertex in B. Let $i_A \in [t]$ be an index such that each of the sets

$$A_1 := \{v_1, \ldots, v_{i_A}\} \cap A \quad \text{and} \quad A_2 := \{v_{i_A+1}, \ldots, v_t\} \cap A$$

contains $k + 1$ elements of A. Pick i_B, B_1, and B_2 symmetrically. Assume first $i_A \leqslant i_B$. Then $i < j$ for all $v_i \in A_1$ and $v_j \in B_2$. Let $A_\star := \{(i, 1) : v_i \in A_1\}$ be the vertices from the bottom row of the t-grid represented by A_1 and likewise let $B_\star := \{(j, 1) : v_j \in B_1\}$. In Figure 15.1 it is easy to see that the vertices from $A_\star$ and $B_\star$ can be matched by $k + 1$ disjoint paths in the t-grid.

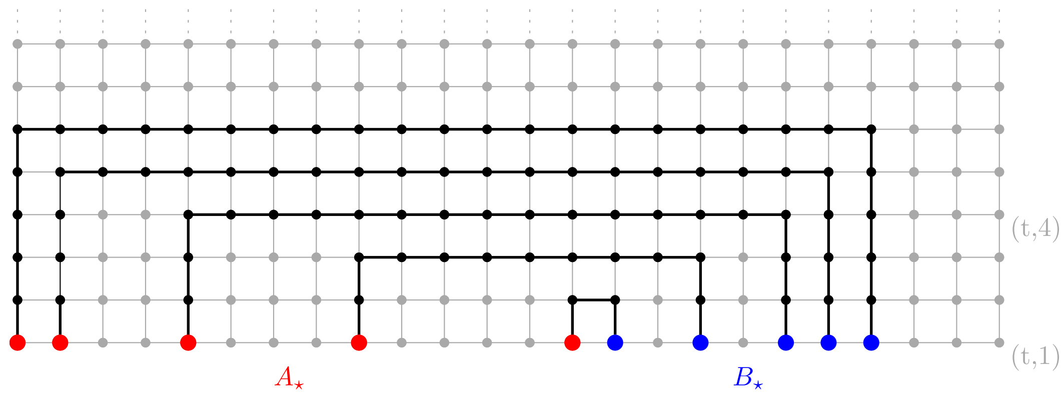

Figure 15.1: Pairing the vertices of $A_\star$ and $B_\star$ with disjoint paths in the t-grid.

By the definition of minor, also A_1 and B_1 can be matched by $k + 1$ disjoint paths in G. We now reach the desired contradiction, as we assumed that no paths run between A and B in $G - S$, but removing the at most k vertices from S can destroy at most k paths.

In the case where $i_A > i_B$, we have $i < j$ for all $v_i \in B_1$ and $v_j \in A_2$ and argue symmetrically. $\square$

Combining Lemma 15.3 and Lemma 15.5 now yields Theorem 15.2.

Chapter 16

Bounded Shrub-Depth

Definition 16.1. A graph class $\mathcal{C}$ is *distance-∞ flip-flat*, if there exists a function $N : \mathbb{N} \to \mathbb{N}$ and a constant $k \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N(m)$ there exists a subset $W_\star \subseteq W$ with $|W_\star| \geqslant m$ and a k-flip H of G such that in H, no two vertices $u, v \in W_\star$ are in the same connected component.

The goal of this chapter is to prove the following.

Theorem 16.2. *A graph class has bounded shrub-depth if and only if it is distance-∞ flip-flat.*

To prove that bounded shrub-depth implies distance-∞ flip-flatness, we work with *flip-depth*, a parameter that is functionally equivalent to shrub-depth. It is defined as follows. The single vertex graph K_1 has flip-depth 0. For $k > 0$, a graph G has flip-depth at most k, if it is a 2-flip of a disjoint union of (arbitrarily many) graphs of flip-depth at most $k - 1$.

Fact 16.3 ([43, Thm. 3.6])**.** *A graph class has bounded shrub-depth if and only if it has bounded flip-depth.*

More precisely, [43, Thm. 3.6] shows the functional equivalence of shrub-depth and the graph parameter *SC-depth*. The definition of SC-depth is obtained from the definition of flip-depth by replacing the 2-flip with a *set complementation*, that is, the operation of complementing all the edges in an arbitrary subset of the vertices. 2-flips generalize set complementations, but any 2-flip can be simulated by performing at most three set complementations. Therefore, flip-depth and SC-depth are functionally equivalent.

Lemma 16.4. *For every graph G of flip-depth at most k and every set $W \subseteq V(G)$, there exists a subset $W_\star \subseteq W$ with $|W_\star| \geqslant |W|^{\frac{1}{2^k}}$ and a 4^k-flip H of G, such that in H, no two vertices from $W_\star$ are in the same connected component.*

Proof. We prove the lemma by induction on k. If $k = 0$, we have $G = K_1$ where the statement holds. For the inductive step assume G has flip-depth at most $k + 1$. Then G is 2-flip of a graph H that is a disjoint union of graphs of flip-depth at most k. If in H at least $|W|^{\frac{1}{2}}$ vertices of W are in pairwise different components, we are done. Otherwise, there exists a component C of H which contains at least $|W|^{\frac{1}{2}}$ vertices of W. By assumption, $H[C]$ has flip-depth at most k. By induction there is a 4^k-flip $H_\star$ of $H[C]$ and a set $W_\star \subseteq W$ of size at least $|W_\star| \geqslant |W|^{\frac{1}{2} \cdot \frac{1}{2^k}}$ such that all vertices from $W_\star$ are in pairwise different components in $H_\star$. Refining

- the size 2 partition of the flip that produced H from G,
- the size 2 partition which marks the component C in H,
- the size 4^k partition of the flip that produced $H_\star$ from $H[C]$,

yields a partition witnessing a 4^{k+1}-flip of G in which all vertices from $W_\star$ are in pairwise different components as desired. □

Corollary 16.5. *Every class of bounded shrub-depth is distance-∞ flip-flat.*

We will use the following proof strategy for the other direction.

Fact 16.6 ([69, Thm. 1.1])**.** *Every class of unbounded shrub-depth transduces the class of all paths.*

Lemma 16.7. *Let $\mathcal{C}$ be a graph class that is distance-∞ flip-flat. Every class that is transducible from $\mathcal{C}$ is also distance-∞ flip-flat.*

Lemma 16.8. *The class of all paths is not distance-∞ flip-flat.*

It is easy to see that combining Fact 16.6, Lemma 16.7, Lemma 16.8 yields the following.

Lemma 16.9. *Every graph class with unbounded shrub-depth is not distance-∞ flip-flat.*

Together with Corollary 16.5, the above lemma proves Theorem 16.2. It remains to prove Lemma 16.7 and Lemma 16.8. The former is an immediate consequence of the following fact where $\varphi(x, y)$ is a formula in the language of colored graphs and $\varphi(G)$ denotes the graph with vertex set $V(G)$ and edge set $\{uv : G \models \varphi(u, v) \vee \varphi(v, u)\}$.

Fact 16.10 ([83, Lem. H.3])**.** *For every formula $\varphi(x, y)$ and $k \in \mathbb{N}$ there exists $s, \ell \in \mathbb{N}$ such that for every colored graph G and for every k-flip G' of G, there exists an ℓ-flip H' of $H := \varphi(G)$ such that for every two vertices u and v adjacent in H' we have that $\mathrm{dist}_{G'}(u, v) \leqslant s$.*

Proof of Lemma 16.7. Assume the class $\mathcal{C}$ is distance-∞ flip-flat with bounds $N(\cdot)$ and k and transduces the class $\mathcal{D}$ using the formula $\varphi(x, y)$. We will show that also $\mathcal{D}$ is flip-flat with bounds $N(\cdot)$ and ℓ, where ℓ is the bound obtained from Fact 16.10 for φ and k. Let $H \in \mathcal{D}$, $m \in \mathbb{N}$, and $W \subseteq V(G)$ be a set of size at least $N(m)$. Since $\mathcal{C}$ transduces $\mathcal{D}$, we have $H = \varphi(G)[V(H)]$ for some colored graph $G \in \mathcal{C}$. By distance-∞ flip-flatness, there is a k-flip G' of G and a subset $W_\star \subseteq W$ of size at least m whose vertices are pairwise in different components in G'. By Fact 16.10, there is also an ℓ-flip H' of $\varphi(G)$ in which the vertices of $W_\star$ are in pairwise different components. It follows that $H'[V(H)]$ is the desired ℓ-flip of H. □

Proof of Lemma 16.8. Assume towards a contradiction that the class of all paths is distance-∞ flip-flat with bounds $N(\cdot)$ and k. Let G be the path containing $N(8k+2)$ vertices. By flip-flatness there exists a k-coloring $\mathcal{K}$ of G and $F \subseteq \mathcal{K}^2$ such that $G_\oplus := G \oplus F$ contains at least $8k + 2$ components. Let $\mathcal{B} := \{C \in \mathcal{K} : |C| \geqslant 5\}$ be the *big* color classes and $W := \bigcup(\mathcal{K} \setminus \mathcal{B})$ be the at most $4k$ vertices contained in *small* color classes. Let G' be the subgraph of G obtained by isolating W and let $F' := F \cap \mathcal{B}^2$ be the restriction of F to $\mathcal{B}$. $G'_\oplus := G' \oplus F'$ is a subgraph of $G_\oplus$: every edge uv in $G'_\oplus$ has no endpoint in W and is therefore also present in $G_\oplus$. Since $G'_\oplus$ and $G_\oplus$ share the same vertex set, $G'_\oplus$ has at least as many components as $G_\oplus$. In order to arrive at the desired contradiction, it remains to bound the number of components in $G'_\oplus$. Towards this goal, we first bound the number of components in G'. As G' is obtained from a path by isolating at most $4k$ vertices, G' contains at most $8k + 1$ components: the path is cut in at most $4k$ places leading to $4k + 1$ components plus the additional at most $4k$ isolated vertices.

Claim 16.11. *If two vertices u and v are adjacent in G', then they are connected in $G'_\oplus$.*

Proof. Since u and v are adjacent, they are from big color classes $\mathcal{K}(u)$ and $\mathcal{K}(v)$. Assume the adjacency between $\mathcal{K}(u)$ and $\mathcal{K}(v)$ was flipped, as otherwise we are done. As G' has maximum degree two, we have $|N_1^{G'}[u] \cup N_1^{G'}[v]| \leqslant 4$. If $\mathcal{K}(u) = \mathcal{K}(v)$ then there exists a vertex in that class that is adjacent to none of u and v in G' and therefore adjacent to both of them in $G'_\oplus$ and we are done. Otherwise, there are three vertices $U \subseteq \mathcal{K}(u)$ non-adjacent to v in G' and three vertices $V \subseteq \mathcal{K}(v)$ non-adjacent to u. Again using the fact that G' has maximum degree two, we find $u' \in U$ and $v' \in V$ that are non-adjacent in G'. It follows that (u, v', u', v) is a path in $G'_\oplus$. ■

It follows that $G'_\oplus$ (and also $G_\oplus$) contains at most $8k + 1$ components; a contradiction. □

Chapter 17

Bounded Tree-Depth

Definition 17.1. A graph class $\mathcal{C}$ is *distance-∞ deletion-flat*, if there exists a function $N : \mathbb{N} \to \mathbb{N}$ and a constant $k \in \mathbb{N}$ such that for all $m \in \mathbb{N}$, $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| \geqslant N(m)$ there exists a set $S \subseteq V(G)$ with $|S| \leqslant k$ and a subset $W_\star \subseteq W \setminus S$ with $|W_\star| \geqslant m$ such that in $G - S$, no two vertices $u, v \in W_\star$ are in the same connected component.

In this chapter we relate distance-∞ deletion-flatness to the graph parameter *tree-depth.* The single vertex graph K_1 has tree-depth 1. For $k > 1$, a graph G has tree-depth at most k if there exists a vertex whose deletion splits G into a disjoint union of (arbitrarily many) graphs of tree-depth at most $k - 1$.

Theorem 17.2. *A graph class has bounded tree-depth if and only if it is distance-∞ deletion-flat.*

Proof. Essentially, the definitions of tree-depth and distance-∞ deletion-flatness are obtained from the definitions of flip-depth and distance-∞ flip-flatness by replacing flips with vertex deletions. Proving that every class of bounded tree-depth is distance-∞ flip-flat is therefore analogous to the proof of Lemma 16.4. The other direction follows by combining the following easy facts.

1. If a class has bounded tree-depth, then so does its closure under taking subgraphs.
2. Every class of unbounded tree-depth contains all paths as subgraphs. [62, Prop. 6.1]
3. The class of all paths is not distance-∞ flip-flat. □

Discussing Dense Versions of Bounded Tree-Depth

For convenience, we reproduce the overview of graph classes from the motivational Chapter 1 in Figure 17.1. Recall that a class property is a set of graph classes. The reader might wonder whether there is a class property that fits into the empty bottom right corner of the hierarchy. The answer depends on what restrictions we impose on the rightmost column. One reasonable demand is that every class property P in the right column satisfies the following three conditions:

(C1) P is transduction-closed,

(C2) the restriction of P to weakly sparse classes is equivalent to the leftmost property in the same row as P,

(C3) P is not subsumed by monadic stability.

Under these constraints, the bottom right corner should remain empty. Every property P that is transduction closed but not monadically stable contains the class of all paths, which is weakly sparse but has unbounded tree-depth.

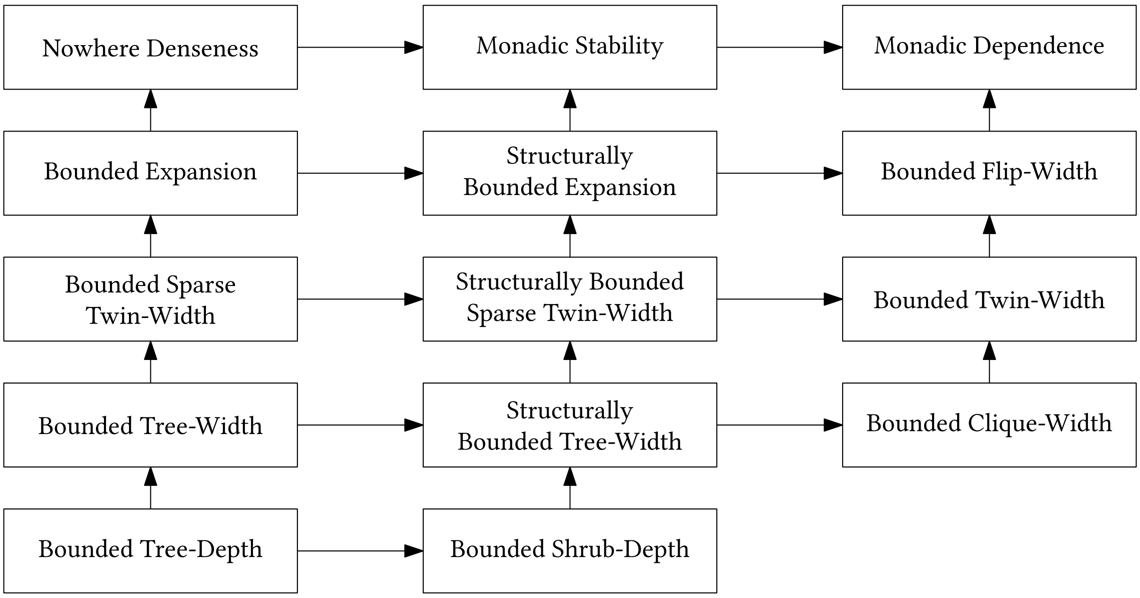

Figure 17.1: A hierarchy of the class properties discussed in this thesis.

A different point of view is given by Gajarský, Pilipczuk, and Toruńczyk in [41], where they provide an in depth discussion about the relations between the columns of (a variant of) the hierarchy in Figure 17.1. Regarding the rightmost column, they argue to drop condition (C3) and conjecture that the rightmost can be defined from the leftmost column as the largest properties that are transduction-closed and whose restriction to weakly sparse classes yields the leftmost column. We illustrate this conjecture in the case of monadic dependence. Monadic dependence is transduction-closed and weakly sparse monadic dependence is equivalent to nowhere denseness. Moreover, the only class property that is transduction-closed but contains monadically independent classes is the trivial class property that contains all graph classes. Its weakly sparse restriction is the property of being weakly sparse, which strictly generalizes nowhere denseness. Therefore, monadic dependence is exactly the largest transduction-closed class property whose restriction to weakly sparse classes is equivalent to nowhere denseness. Adopting this point of view, the fitting property to put into the bottom right corner of Figure 17.1 is again bounded shrub-depth, which can be easily argued using the results from [69].

Chapter 18

Conclusion

In this thesis we have presented several characterizations of monadically stable and monadically dependent graph classes, that bridge the gap between structural graph theory and model theory. One of the driving forces behind the study of monadic stability and monadic dependence is their conjectured relation to the first-order model checking problem [2, 7].

Conjecture 1.1. *A hereditary graph class admits fixed-parameter tractable model checking if and only if it is monadically dependent (assuming* $\mathrm{FPT} \neq \mathrm{AW}[*]$*).*

Building on our characterizations, we have resolved the hardness side of this conjecture, by showing that model checking is $\mathrm{AW}[*]$-hard on every hereditary, monadically independent graph class. On the tractability side we have made strong progress by developing a fixed-parameter tractable first-order model checking algorithm for monadically stable graph classes. Whether model checking is fixed-parameter tractable also on the more general monadically dependent classes remains an open question. How should one approach this question? Our model checking algorithm for monadically stable classes builds on two main ingredients: the Flipper game and sparse neighborhood covers. The Flipper game has a *qualitative* character. For every radius r, the length of the Flipper game is bounded by a constant. Flip-flatness and the characterization by forbidden induced subgraphs also fall into this category. On the other hand, the neighborhood covers have a *quantitative* character. Instead of constant bounds, we have bounds of the form n^{ε} or $n^{1+\varepsilon}$ for the overlap of the cover and the neighborhood complexity. Our model checking algorithm combines both the qualitative and the quantitative aspect of monadic stability. We predict that model checking monadically dependent classes requires further advancements in both directions. On the qualitative side, we made initial progress with our flip-breakability characterization. The obvious next question is the following.

Question 18.1. Is there a game characterization for monadically dependent classes?

The quantitative theory of monadic dependence is less developed. A first step here would be to prove the following conjecture stated in [24].

Conjecture 18.2. *Every monadically dependent class has almost linear neighborhood complexity.*

For monadically stable classes, the following conjecture [68, 41] still remains open.

Conjecture 18.3. *Every monadically stable graph class is structurally nowhere dense (i.e., transducible from a nowhere dense class).*

This conjecture is currently open even for classes of *bounded local shrub-depth*. A graph class $\mathcal{C}$ has bounded local shrub-depth if there is a function $f : \mathbb{N} \to \mathbb{N}$ such that every graph induced by an every r-neighborhood in $\mathcal{C}$ has shrub-depth at most $f(r)$. Using the Flipper game, it can be easily shown that these classes are monadically stable. However, it is not known whether they are transducible from nowhere dense classes. Could they even be transducible from classes of bounded local tree-depth?

While we have shown that all first-order definable problems are fixed-parameter tractable on monadically stable classes, a finer grained analysis of the parameterized complexity is missing.

Question 18.4. Do monadically stable (monadically dependent) classes admit polynomial kernels for k-Independent-Set and k-Dominating-Set?

For some subclasses this question has been answered positively in [25]. Here, a possible angle of attack is to use the fact that flip-flatness yields distance-r independent sets of polynomial size [27]. These bounds are obtained by replacing the bounds from Ramsey's theorem with the polynomial bounds from [58]. Recent results for graphs of bounded VC dimension [67] suggest that polynomial bounds can also be obtained for flip-breakability.

Bibliography

[1] 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.

[2] Algorithms, Logic and Structure Workshop in Warwick. Open Problems. *https://warwick.ac.uk/fac/sci/maths/people/staff/daniel_kral/alglogstr/openproblems.pdf*, 2016. [Online; accessed 10-Jun-2024].

[3] John T Baldwin. *Fundamentals of stability theory*, volume 12. Cambridge University Press, 2017.

[4] 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.

[5] Édouard Bonnet, Jan Dreier, Jakub Gajarský, Stephan Kreutzer, Nikolas Mählmann, Pierre Simon, and Szymon Toruńczyk. 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.

[6] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: Small classes. *Combinatorial Theory*, 2(2), 2022.

[7] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. Twin-width iv: Ordered graphs and matrices. *J. ACM*, 71(3), jun 2024.

[8] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width i: Tractable fo model checking. *J. ACM*, 69(1), nov 2021.

[9] Samuel Braunfeld and Michael C. Laskowski. Existential characterizations of monadic nip. *arXiv preprint arXiv:2209.05120*, 2022.

[10] Bernard Chazelle and Emo Welzl. Quasi-optimal range searching in space of finite VC-dimension. *Discret. Comput. Geom.*, 4:467–489, 1989.

[11] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. Strong computational lower bounds via parameterized complexity. *Journal of Computer and System Sciences*, 72(8):1346–1367, 2006.

[12] Zhi-Zhong Chen, Michelangelo Grigni, and Christos H. Papadimitriou. Planar map graphs. In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 514–523. ACM, 1998.

[13] Bruno Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Information and computation*, 85(1):12–75, 1990.

[14] Bruno Courcelle, Johann A Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems*, 33(2):125–150, 2000.

[15] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Daniel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms.* Springer Publishing Company, Incorporated, 1st edition, 2015.

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

[17] Anuj Dawar, Martin Grohe, and Stephan Kreutzer. Locally excluding a minor. In *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*, pages 270–279. IEEE, 2007.

[18] Anuj Dawar and Stephan Kreutzer. Domination Problems in Nowhere-Dense Classes. In Ravi Kannan and K. Narayan Kumar, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 4 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 157–168, Dagstuhl, Germany, 2009. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[19] Erik D. Demaine, Felix Reidl, Peter Rossmanith, Fernando Sánchez Villaamil, Somnath Sikdar, and Blair D. Sullivan. Structural sparsity of complex networks: Bounded expansion in random models and real-world graphs. *Journal of Computer and System Sciences*, 105:199–241, 2019.

[20] 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.

[21] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity.* Texts in Computer Science. Springer, London, 2013.

[22] Rodney G. Downey, Michael R. Fellows, and Udayan Taylor. The parameterized complexity of relational database queries and an improved characterization of w[1]. In *International Conference on Discrete Mathematics and Theoretical Computer Science*, 1996.

[23] Pål Grønås Drange, Markus S. Dregi, Fedor V. Fomin, Stephan Kreutzer, Daniel Lokshtanov, Marcin Pilipczuk, Michal Pilipczuk, Felix Reidl, Saket Saurabh, Fernando Sánchez Villaamil, and Somnath Sikdar. Kernelization and sparseness: the case of dominating set. *CoRR*, abs/1411.4575, 2014.

[24] Jan Dreier, Ioannis Eleftheriadis, Nikolas Mählmann, Rose M. McCarty, Michał Pilipczuk, and Szymon Toruńczyk. First-order model checking on monadically stable graph classes. *arXiv preprint arXiv:2311.18740*, 2023.

[25] 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.

[26] Jan Dreier, Nikolas Mählmann, and Sebastian Siebertz. First-order model checking on structurally sparse graph classes. In *Proceedings of the 55th Annual ACM Symposium on*

Theory of Computing, STOC 2023, pages 567–580, New York, NY, USA, 2023. Association for Computing Machinery.

[27] Jan Dreier, Nikolas Mählmann, Sebastian Siebertz, and Szymon Toruńczyk. Indiscernibles and flatness in monadically stable and monadically NIP classes. In Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors, *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023, July 10-14, 2023, Paderborn, Germany*, volume 261 of *LIPIcs*, pages 125:1–125:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

[28] Jan Dreier, Nikolas Mählmann, and Szymon Toruńczyk. Flip-breakability: A combinatorial dichotomy for monadically dependent graph classes. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, page 1550–1560, New York, NY, USA, 2024. Association for Computing Machinery.

[29] Zdeněk Dvořák. *Asymptotical structure of combinatorial objects.* PhD thesis, Charles University, Faculty of Mathematics and Physics, 2007.

[30] Zdeněk Dvořák. Induced subdivisions and bounded expansion. *European Journal of Combinatorics*, 69:143–148, 2018.

[31] Zdeněk Dvořák, Daniel Král, and Robin Thomas. Testing first-order properties for subclasses of sparse graphs. *J. ACM*, 60(5):36:1–36:24, 2013.

[32] Kord Eickmeyer, Archontia C. Giannopoulou, Stephan Kreutzer, O-joung Kwon, Michał Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. Neighborhood complexity and kernelization for nowhere dense classes of graphs. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017*, volume 80 of *LIPIcs*, pages 63:1–63:14. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2017.

[33] Jörg Flum and Martin Grohe. Fixed-parameter tractability, definability, and model-checking. *SIAM Journal on Computing*, 31(1):113–145, 2001.

[34] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series).* Springer-Verlag, Berlin, Heidelberg, 2006.

[35] Markus Frick and Martin Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM (JACM)*, 48(6):1184–1206, 2001.

[36] 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.

[37] Jakub Gajarský, Petr Hliněný, Jan Obdržálek, Daniel Lokshtanov, and M. S. Ramanujan. A new perspective on FO model checking of dense graph classes. *ACM Trans. Comput. Logic*, 21(4), July 2020.

[38] Jakub Gajarský, 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.

[39] 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. In Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors, *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261 of *Leibniz International Proceedings in Informatics*

(LIPIcs), pages 128:1–128:16, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[40] Jakub Gajarský, Michał Pilipczuk, Wojciech Przybyszewski, and Szymon Toruńczyk. Twin-Width and Types. In *49th International Colloquium on Automata, Languages, and Programming*, pages 123:1–123:21, 2022.

[41] 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*, pages 1–12, 2022.

[42] Robert Ganian, Petr Hliněný, Alexander Langer, Jan Obdržálek, Peter Rossmanith, and Somnath Sikdar. Lower bounds on the complexity of mso1 model-checking. *Journal of Computer and System Sciences*, 80(1):180–194, 2014.

[43] Robert Ganian, Petr Hliněný, Jaroslav Nešetřil, Jan Obdržálek, Patrice Ossona de Mendez, and Reshma Ramadurai. When trees grow low: Shrubs and fast MSO1. In Branislav Rovan, Vladimiro Sassone, and Peter Widmayer, editors, *Mathematical Foundations of Computer Science 2012*, pages 419–430, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[44] William Gasarch, Andy Parrish, and Sandow Sinai. Three proofs of the hypergraph ramsey theorem (an exposition). *arXiv preprint arXiv:1206.4257*, 2012.

[45] Sylvain Gravier, Frédéric Maffray, Jérôme Renault, and Nicolas Trotignon. Ramsey-type results on singletons, co-singletons and monotone sequences in large collections of sets. *European Journal of Combinatorics*, 25(5):719–734, 2004.

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

[47] Martin Grohe and Stephan Kreutzer. Methods for algorithmic meta theorems. *Model Theoretic Methods in Finite Combinatorics*, 558:181–206, 2011.

[48] 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.

[49] Wilfrid Hodges. *Model Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993.

[50] Wilfrid Hodges. *A Shorter Model Theory*. Cambridge University Press, 1997.

[51] Sang il Oum and Paul Seymour. Approximating clique-width and branch-width. *Journal of Combinatorial Theory, Series B*, 96(4):514–528, 2006.

[52] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.

[53] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001.

[54] Stephan Kreutzer. *Algorithmic meta-theorems*, pages 177–270. London Mathematical Society Lecture Note Series. Cambridge University Press, 2011.

[55] 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.

[56] Stephan Kreutzer and Siamak Tazari. Lower bounds for the complexity of monadic second-order logic. In *2010 25th Annual IEEE Symposium on Logic in Computer Science*, pages 189–198. IEEE, 2010.

[57] Leonid Libkin. *Elements of finite model theory*, volume 41. Springer, 2004.

[58] Maryanthe Malliaris and Saharon Shelah. Regularity lemmas for stable graphs. *Transactions of the American Mathematical Society*, 366(3):1551–1585, 2014.

[59] Jiří Matoušek. Spanning trees with low crossing number. *RAIRO - Theoretical Informatics and Applications - Informatique Théorique et Applications*, 25(2):103–123, 1991.

[60] Jaroslav Nešetřil and Patrice Ossona de Mendez. First order properties on nowhere dense structures. *The Journal of Symbolic Logic*, 75(3):868–887, 2010.

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

[62] Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity: graphs, structures, and algorithms*, volume 28. Springer Science & Business Media, 2012.

[63] Jaroslav Nešetřil and Patrice Ossona de Mendez. Structural sparsity. *Russian Mathematical Surveys*, 71(1):79, 2016.

[64] 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.

[65] Jaroslav Nešetřil and Patrice Ossona de Mendez. Grad and classes with bounded expansion i. decompositions. *European Journal of Combinatorics*, 29(3):760–776, 2008.

[66] Jaroslav Nešetřil, Patrice Ossona de Mendez, Roman Rabinovich, and Sebastian Siebertz. Classes of graphs with low complexity: The case of classes with bounded linear rankwidth. *European Journal of Combinatorics*, 91:103223, 2021. Colorings and structural graph theory in context (a tribute to Xuding Zhu).

[67] Tung Nguyen, Alex Scott, and Paul Seymour. Induced subgraph density. vi. bounded vc-dimension. *arXiv preprint arXiv:2312.15572*, 2023.

[68] Patrice Ossona de Mendez. First-order transductions of graphs (invited talk). In *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021*, volume 187 of *LIPIcs*, pages 2:1–2:7. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2021.

[69] Patrice Ossona de Mendez, Michał Pilipczuk, and Sebastian Siebertz. Transducing paths in graph classes with unbounded shrubdepth. *European Journal of Combinatorics*, page 103660, 2022.

[70] Michał Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. On the number of types in sparse graphs. In *33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018*, pages 799–808. ACM, 2018.

[71] Anand Pillay. *An introduction to stability theory.* Courier Corporation, 2008.

[72] Klaus-Peter Podewski and Martin Ziegler. Stable graphs. *Fundamenta Mathematicae*, 100(2):101–107, 1978.

[73] Tibor Radó. On non-computable functions. *The Bell System Technical Journal*, 41(3):877–884, 1962.

[74] Frank P. Ramsey. On a Problem of Formal Logic. *Proceedings of the London Mathematical Society*, s2-30(1):264–286, 01 1930.

[75] Neil Robertson and Paul D Seymour. Graph minors. v. excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.

[76] Norbert Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, 1972.

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

[78] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41(1):247–261, 1972.

[79] Saharon Shelah. *Classification theory: and the number of non-isomorphic models.* Elsevier, 1990.

[80] Pierre Simon. *A Guide to NIP Theories.* Lecture Notes in Logic. Cambridge University Press, 2015.

[81] Katrin Tent and Martin Ziegler. *A Course in Model Theory.* Lecture Notes in Logic. Cambridge University Press, 2012.

[82] Szymon Toruńczyk. Lectures on finite model theory. *https://duch.mimuw.edu.pl/~szymtor/fmt.pdf*, 2022. [Online; accessed 10-Jun-2024].

[83] Szymon Toruńczyk. Flip-width: Cops and robber on dense graphs. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 663–700, Santa Cruz, CA, USA, nov 2023. IEEE Computer Society. Full version available at https://arxiv.org/abs/2302.00352.

[84] Vladimir Naumovitch Vapnik and Alexei Iakovlevitch Červonenkis. On the uniform convergence of relative sequences of events to their probabilities. *Theory Probab. Appl.*, 16:264–280, 1971.

[85] Emo Welzl. Partition trees for triangle counting and other range searching problems. In *Fourth Annual Symposium on Computational Geometry, SoCG 1988*, pages 23–33. ACM, 1988.

[86] Xuding Zhu. Colouring graphs with bounded generalized colouring number. *Discrete Mathematics*, 309(18):5562–5568, 2009. Combinatorics 2006, A Meeting in Celebration of Pavol Hell's 60th Birthday (May 1–5, 2006).