Low Rank MSO

Mikołaj Bojańczyk
University of Warsaw, Poland

Michał Pilipczuk
University of Warsaw, Poland

Wojciech Przybyszewski
University of Warsaw, Poland

Marek Sokołowski
Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

Giannos Stamoulis
Université Paris Cité, CNRS, IRIF, F-75013, Paris, France

Abstract

We introduce a new logic for describing properties of graphs, which we call *low rank MSO*. This is the fragment of monadic second-order logic in which set quantification is restricted to vertex sets of bounded cutrank. We prove the following statements about the expressive power of low rank MSO.

- Over any class of graphs that is weakly sparse, low rank MSO has the same expressive power as separator logic. This equivalence does not hold over all graphs.
- Over any class of graphs that has bounded VC dimension, low rank MSO has the same expressive power as flip-connectivity logic. This equivalence does not hold over all graphs.
- Over all graphs, low rank MSO has the same expressive power as flip-reachability logic.

Here, separator logic is an extension of first-order logic by basic predicates for checking connectivity, which was proposed by Bojańczyk [ArXiv 2107.13953] and by Schirrmacher, Siebertz, and Vigny [ACM ToCL 2023]. Flip-connectivity logic and flip-reachability logic are analogues of separator logic suited for non-sparse graphs, which we propose in this work. In particular, the last statement above implies that every property of undirected graphs expressible in low rank MSO can be decided in polynomial time.

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

Keywords and phrases First Order logic, Monadic Second Order logic, cutrank, flips

Digital Object Identifier 10.4230/LIPIcs.LICS.2026.22

Related Version *Full Version*: `https://arxiv.org/abs/2502.08476` [2]

Funding The research of Mi. Pilipczuk, W. Przybyszewski, and G. Stamoulis has been supported by the project BOBR that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement No. 948057. In particular, a majority of work on this manuscript was done while G. Stamoulis was affiliated with University of Warsaw.
Mikołaj Bojańczyk: Mikołaj Bojańczyk was supported by the Polish National Science Centre (NCN) grant "Polynomial finite state computation" (2022/46/A/ST6/00072).

Acknowledgements We would like to thank the anonymous reviewers for their valuable comments and remarks.

1 Introduction

Much of the contemporary work on logic on graphs revolves around the first-order logic FO and variants of the monadic second-order logic MSO. The general understanding is that MSO behaves well, in terms of computational aspects, on graphs that have a tree-like structure [4, 5]. The tameness of FO is the subject of an ongoing line of work and is expected to be connected with the non-existence of complicated local structures [6, 7, 8, 9, 23]. A general overview of this area can be found in the recent survey of Pilipczuk [17].

Naturally, there are interesting concepts of logic on graphs besides FO, MSO, and their variants. One such logic was proposed independently by Bojańczyk [1] (under the name *separator logic*) and by Schirrmacher, Siebertz, and Vigny [20] (under the name FO+conn). The idea is to extend FO by a simple mechanism for expressing connectivity: we allow the usage of predicates $\mathsf{conn}_k(s, t, a_1, \ldots, a_k)$, for every $k \in \mathbb{N}$, that verify the existence of a path connecting vertices s and t that avoids vertices $a_1, \ldots, a_k$. Thus, the expressive power of separator logic lies strictly between FO and MSO: the logic is not bound only to local properties, like FO, but full set quantification is not available. Bojańczyk [1] argued that on classes of graphs of bounded pathwidth, separator logic characterizes a suitable analogue of star-free languages. Pilipczuk, Schirrmacher, Siebertz, Toruńczyk, and Vigny [18] showed that computational tameness of separator logic (precisely, fixed-parameter tractability of the model-checking problem) coincides on subgraph-closed classes with a natural dividing line in structural graph theory: exclusion of a fixed topological minor.

The drawback of separator logic is that the connectivity predicates are designed with separators consisting of a bounded number of vertices in mind. Such separators are a fundamental concept in the structural theory of sparse graphs, but become ill-suited for the treatment of dense graphs. In contrast, both FO and MSO can be naturally deployed on graphs regardless of their sparsity, and much of the recent developments around these logics concern understanding their expressive power and computational aspects on classes of well-structured dense graphs. What would be then a dense analogue of separator logic?

Low rank MSO. We introduce a new logic on graphs called *low rank MSO* with expressive power lying strictly between FO and MSO. We argue that low rank MSO is a natural dense analogue of separator logic. The main idea is the following: as separator logic allows quantification over small separators[1] as understood in the theory of sparse graphs, its dense analogue should allow quantification over small separators as understood in the theory of dense graphs. The latter concept is delivered by the notion of *cutrank*.

Let G be a graph with vertex set V and X be a subset of V; denote $\widehat{X} \coloneqq V - X$ for brevity. The idea is to measure the complexity of the cut $(X, \widehat{X})$ by examining the adjacency matrix $\mathsf{Adj}_G[X, \widehat{X}]$: the $\{0, 1\}$-matrix with rows indexed by X and columns indexed by $\widehat{X}$, where the entry at the intersection of the row of $u \in X$ and the column of $v \in \widehat{X}$ is 1 if u and v are adjacent, and 0 otherwise. The *cutrank* of X, denoted $\mathrm{rk}(X)$, is the rank of $\mathsf{Adj}_G[X, \widehat{X}]$ over the binary field $\mathbb{F}_2$. Cutrank was introduced by Oum and Seymour in their work on *rankwidth* [15, 16]. This graph parameter, together with the functionally equivalent *cliquewidth* [5] and *NLC-width* [24], is understood as a suitable notion of tree-likeness for dense graphs. In particular, in graph theory there is an ongoing work on constructing a theory of *vertex-minors*, which is expected to be a dense counterpart of the theory of graph minors. In this theory, rankwidth is the analogue of treewidth, and cuts of low cutrank play the same role as vertex separators of bounded size. See the survey of Kim and Oum [12] and the PhD thesis of McCarty [13] for an introduction to this area.

With the motivation behind cutrank understood, the definition of low rank MSO becomes natural: it is the fragment of MSO on (undirected) graphs, where every use of the set quantifier must be accompanied by an explicit bound on the cutrank of the quantified set.

[1] In fact, separator logic allows to query whether a small set of vertices is a separator.

The syntax uses the following constructors:

$$\underbrace{\forall x \quad \exists x}_{\substack{\text{quantification}\\ \text{over vertices}}} \qquad \underbrace{\forall X : r \quad \exists X : r}_{\substack{\text{quantification over}\\ \text{subsets of vertices}\\ \text{with cutrank at most } r}} \qquad \underbrace{\wedge \quad \vee \quad \neg}_{\substack{\text{Boolean}\\ \text{combinations}}} \qquad \underbrace{E(x,y)}_{\text{edge relation}} \qquad \underbrace{x \in X}_{\text{set membership}}$$

Here r denotes some fixed integer. The semantics are defined in the expected way, as a fragment of MSO with quantification over sets of vertices. (This variant of MSO is often called MSO_1 in the literature, as opposed to MSO_2, where quantification over edge sets is also allowed. We do not consider MSO_2 here.)

The goal of this work is to introduce low rank MSO, study its expressive power, and pose multiple questions related to it. While we are mostly consider undirected graphs[2], low rank MSO can be defined on any structures where a meaningful notion of rank can be considered. We expand on this in Section 7.

Our results. It is not hard to see that predicates conn_k can be expressed in low rank MSO, hence low rank MSO is at least as expressive as separator logic (see Proposition 9). We first prove that on classes of sparse graphs, the expressive power of the two logics actually coincide. Here, we say that a class of graphs $\mathscr{C}$ is *weakly sparse* if there is $t \in \mathbb{N}$ such that all graphs from $\mathscr{C}$ do not contain the biclique $K_{t,t}$ as a subgraph.

▶ **Theorem 1.** *Let $\mathscr{C}$ be a weakly sparse class of graphs. Then for every formula of low rank MSO there exists a formula of separator logic such that the two formulas are equivalent on all graphs from $\mathscr{C}$.*

Theorem 1 corroborates the claim that low rank MSO is a dense analogue of separator logic. There is, however, an important difference. Low rank MSO is defined as a fragment of MSO, and not as an extension of FO, like the separator logic. This has a particular impact on the computational aspects. For instance, observe that for every fixed sentence φ of separator logic, whether φ holds in a given a graph G can be decided in polynomial time. Indeed, it suffices to apply a brute-force recursive algorithm that, for every next vertex variable quantified, examines all possible evaluations of this variable in the graph; and in the leaves of the recursion, adjacency and connectivity predicates can be checked in polynomial time. This approach does not work for low rank MSO, for the number of sets of bounded rank in a graph can be exponential.

To bridge this gap, it would be useful to find a logic closer in spirit to FO that would be equivalent to low rank MSO. For this, we use the concept of *definable flips*, which has played an important role in the recent advances in understanding FO on dense graphs, see [3, 7, 8, 9, 23]. Formal definitions are given in Section 2, but in a nutshell, a *definable flip* of a graph G is any graph that can be obtained from G as follows:

- Select a tuple $\bar{a}$ of vertices of G. These are the *parameters* of the flip.
- Classify the vertices of G according to their adjacency to the vertices of $\bar{a}$. The vertices of $\bar{a}$ belong to their own singleton classes.
- For every pair of classes K and L, either leave the adjacency relation between K and L intact, or *flip* it: exchange all edges with non-edges, and vice versa. This can be applied also for $K = L$, which amounts to complementing the adjacency relation within K.

[2] Concretely, we consider vertex-colored undirected graphs, where the vertex set may be equipped with some unary predicates.

Note that a single flip applied in the last point can remove a large biclique, so a definable flip of a dense graph may be much sparser. In the aforementioned works [3, 7, 8, 9, 23], definable flips have been identified as a useful notion of separations in dense graphs that can be described succinctly, by specifying a constant number of parameters and a constant-size piece of information on which pairs of classes should be flipped.

This leads to the following candidate for an extension of FO that could be equivalent to low rank MSO. We define *flip-connectivity logic* as the extension of FO on undirected graphs obtained by additionally allowing the usage of predicates $\mathsf{flipconn}_{k,A}(s,t,a_1,\ldots,a_k)$. Such a predicate verifies whether s and t are in the same connected component of the flip of G defined by the tuple $\bar{a}=(a_1,\ldots,a_k)$ and the relation A specifying pairs of classes to be flipped. Again, it is not hard to prove that these predicates can be expressed in low rank MSO, implying that low rank MSO is at least as expressive as flip-connectivity logic (see Proposition 9). We show that the expressive power of low rank MSO is in fact strictly larger on all graphs, but the two logics are equivalent on every class of graphs with bounded Vapnik–Chervonenkis (VC) dimension (see Section 5.2 for the definition).

▶ **Theorem 2.** *There is a sentence of low rank MSO that cannot be expressed in flip-connectivity logic.*

▶ **Theorem 3.** *Let $\mathscr{C}$ be a graph class of bounded VC dimension. Then for every formula of low rank MSO there exists a formula of flip-connectivity logic such that the two formulas are equivalent on all graphs from $\mathscr{C}$.*

Finally, we show that flip-connectivity logic can be amended to a stronger logic so that it becomes equivalent to low rank MSO on all graphs. Concretely, in *flip-reachability logic*, we consider a more general notion of directed flips that work naturally in the setting of directed graphs. The definition is essentially the same as in the undirected setting, except that flipping the arc relation between classes K and L exchanges arcs with non-arcs in the set $K\times L$. Note that thus, directed flips work on *ordered pairs* of classes: exchanging arcs with non-arcs in $K\times L$ and in $L\times K$ are two different operations, and they can be applied or not independently. Flip-reachability logic is defined by extending FO by predicates $\mathsf{flipreach}_{k,A}(s,t,a_1,\ldots,a_k)$ that test whether t is *reachable* from s in the directed flip defined by $\bar{a}=(a_1,\ldots,a_k)$ and A.

While flip-reachability logic naturally works in the domain of directed graphs, we can deploy it also on undirected graphs by replacing every edge with a pair of arcs directed oppositely. Note, however, that even if we start with an undirected graph G, a definable directed flip of G is not necessarily undirected. Hence, access to the reachability relation in directed flips of G may provide a larger expressive power than access to the connectivity relation in undirected flips of G. We prove that this is indeed the case, and the expressive power of flip-reachability logic meets that of low rank MSO.

▶ **Theorem 4.** *For every formula of low rank MSO there exists a formula of flip-reachability logic such that the two formulas are equivalent on all undirected graphs.*

Similarly to flip-connectivity logic, also the flip-reachability predicates are expressible in low rank MSO (see Proposition 9), hence the expressive power of the two logics is indeed equal. We also remark that in Theorem 4 it is important that we consider low rank MSO and flip-reachability logic in the domain of undirected graphs. We currently do not know whether also on all directed graphs, the expressive power of flip-reachability logic coincides with that of (naturally defined) low rank MSO.

The proof of Theorem 4 relies on a combinatorial characterization of sets of low rank in undirected graphs. This characterization is captured in the Low Rank Structure Theorem (Theorem 25). This essentially states that the family of all sets of low rank admits a bounded-size representation. Very roughly, we show that each set X of low rank is described consicely by a bounded-size tuple of vertices, which determines a set X_+ to be included, a set X_- to be excluded, and a partition $\mathcal{X}$ of the remaining vertices into classes, so that X is the union of some members of $\mathcal{X}$ and X_+. A more detailed presentation of this theorem is given in Section 6.

Similarly to separator logic, both flip-connectivity logic and flip-reachability logic are defined as extensions of FO by predicates whose satisfaction on given arguments can be checked in polynomial time. Hence, for every fixed sentence φ of flip-reachability logic, it can be decided in polynomial time whether a given graph G satisfies φ. By combining this observation with Theorem 4 we conclude the following.

▶ **Corollary 5.** *Every graph property definable in low rank MSO can be decided in polynomial time.*

In the terminology of parameterized complexity, this proves that the model-checking problem for low rank MSO is *slice-wise polynomial*, that is, belongs[3] to the complexity class XP.

Organization. In Section 2 we formally introduce all the considered logics and establish some basic properties. In Section 3 we propose a simple framework for proving equivalence with low rank MSO. Main results comparing low rank MSO with separator logic (Theorem 1), flip-connectivity logic (Theorems 2 and 3), and flip-reachability logic (Theorem 4) are discussed in Sections 4, 5, and 6, respectively. In Section 7 we discuss several open questions and possible directions for future work. The proofs of statements marked with ♠ can be found in the full version of this work [2].

2 Definitions and basic properties

For $k \in \mathbb{N}$, we denote $[k] \coloneqq \{1, \ldots, k\}$. Unless explicitly stated, all graphs in this chapter are finite, simple (i.e., without loops or parallel edges), and undirected. They may also be *vertex-colored*: the vertex set is equipped with a number of distinguished subsets, called colors. We treat such graphs as relational structures consisting of the vertex set, a binary symmetric edge relation (denoted E), and a number of unary predicates signifying colors. Note that a vertex may belong to several different colors simultaneously, or to no color at all.

For a set A of vertices of G, we denote by $G[A]$ the subgraph of G induced by A, i.e. the graph with vertex set A, edge set consisting of all edges of G with both endpoints in A, and colors inherited naturally. For a vertex v of G and a set of vertices $A \subseteq V(G)$, we denote by $E(v, S)$ the set of neighbors of v in S, i.e. $E(v, S) = \{u \in S \mid uv \in E(G)\}$. The *neighborhood* of a vertex is $N(v) \coloneqq E(v, V(G))$.

We assume reader's familiarity with the first-order logic FO and monadic second-order logic MSO. For convenience, in all formulas of all the considered logics, including FO and its extensions, we allow free set variables. If X is such a set variable and x is a vertex

[3] Formally, Corollary 5 places the problem only in non-uniform XP, because for every sentence φ of low rank MSO we obtain a different algorithm for deciding the satisfaction of φ. However, our proof of Theorem 4 can be turned into an effective algorithm for translating φ into an equivalent sentence φ' of flip-reachability logic, which can be subsequently decided by brute force. This yields a uniform XP algorithm for model-checking low rank MSO. We omit the details.

variable, then we allow membership tests of the form $x \in X$. By $\bar{x}$ we denote a tuple of vertex variables and by $\overline{X}$ we denote a tuple of set variables. For a graph G and a tuple of vertices $\bar{x}$, we denote by $G^{\bar{x}}$ the set of evaluations of $\bar{x}$ in G, i.e. functions from the variables in $\bar{x}$ to the vertices of G. Similarly, for a tuple of set variables $\overline{X}$, we denote by $G^{\overline{X}}$ the set of evaluations of $\overline{X}$ in G. For brevity, we might identify tuples of vertices with their respective evaluations. We say that a logic $\mathcal{L}$ is an *extension* of FO if $\mathcal{L}$ contains FO as a fragment.

For a formula $\varphi(x, \overline{Y}, \bar{z})$ of a logic $\mathcal{L}$, a graph G, and evaluations $\overline{B}$ of $\overline{Y}$ and $\bar{c}$ of $\bar{z}$, we say that $\varphi(x, \overline{B}, \bar{c})$ *defines* a set $A \subseteq V(G)$ in G if

$$A = \{v \in V(G) \mid G \models \varphi(v, \overline{B}, \bar{c})\}.$$

We also denote by $\varphi(G, \overline{B}, \bar{c})$ the set defined by $\varphi(x, \overline{B}, \bar{c})$ in G.

Low rank MSO. We have already defined low rank MSO in Section 1. In the full version, we offer some simple remarks about the choice of rank over $\mathbb{F}_2$ as the measure of complexity of a matrix. In particular, we argue that replacing this choice with any other similar measure would yield a logic with the same expressive power. For the remainder of this paper, the cutrank of a set X is simply called the *rank* of X.

Separator logic, flip-connectivity logic, and flip-reachability logic. Separator logic has also been introduced in Section 1. Recall that it is defined as the extension of FO on graphs by predicates conn_k for $k \in \mathbb{N}$, each of arity $k+2$, with the following semantics: if G is a graph and $s, t, a_1, \ldots, a_k$ are vertices of G, then $\mathsf{conn}_k(s, t, a_1, \ldots, a_k)$ holds in G if and only if there is a path with endpoints s and t that does not pass through any of the vertices $a_1, \ldots, a_k$.

We now define flip-connectivity logic and flip-reachability logic. Flip-reachability logic works naturally on directed graphs (*digraphs*) and flip-connectivity logic will be a special case of the definition, hence we need to introduce some terminology on digraphs.

A directed graph (digraph) is a pair $G = (V, E)$ consisting of a set $V = V(G)$ of vertices and a set $E = E(G)$ of *arcs*. An arc from u to v is denoted $\vec{uv}$ and has tail u and head v. We specify that digraphs do not contain self-loops, so $u \neq v$ for each arc $\vec{uv}$, and neither do they contain multiple copies of an arc. However, we allow parallel arcs connecting two vertices in the opposite directions.

The *atomic type* of a k-tuple $\bar{v} = (v_1, \ldots, v_k)$ of vertices of an undirected graph G, denoted $\mathrm{atp}(\bar{v}) = \mathrm{atp}(v_1, \ldots, v_k)$, is the set of all atomic formulas of the form $x_i = x_j$, $E(x_i, x_j)$, and $U(x_i)$ for some $i, j \in \{1, \ldots, k\}$ and a unary predicate U in the language of G, satisfied by $\bar{v}$ in G. Naturally, two p-tuples $\bar{u}$, $\bar{v}$ of vertices satisfy the same quantifier-free formulas of first-order logic with no parameters if and only if $\mathrm{atp}(\bar{u}) = \mathrm{atp}(\bar{v})$. We define *edge type* of a k-tuple $\bar{v}$ similarly to the atomic type, but we do not consider the unary predicates.

Let atp^k denote the set of all possible atomic types of k-tuples of vertices of an undirected graph. Note that atp^k is finite and of size bounded by a function of k and the number of unary predicates interpreted in G. Atomic types could be also naturally defined for directed graphs, but we will use them only in the undirected context.

Next, we define flips of (undirected) graphs with respect to a tuple of parameters. Let $k \in \mathbb{N}$, $A \subseteq \mathrm{atp}^{k+1} \times \mathrm{atp}^{k+1}$, G be an undirected graph and $\bar{a} = (a_1, \ldots, a_k)$ be a k-tuple of vertices of G – the *parameters* of the flip. Then the *A-flip of G with parameters $\bar{a}$*, denoted $G \oplus_{\bar{a}} A$, is the directed graph with the same vertex set as G, where for distinct $u, v \in V(G)$, we have

$$\vec{uv} \in E(G \oplus_{\bar{a}} A) \qquad \text{if and only if} \qquad [uv \in E(G)] \text{ xor } [(\mathrm{atp}(u, \bar{a}), \mathrm{atp}(v, \bar{a})) \in A].$$

Note that if the relation A is symmetric, then the resulting graph is always undirected, in the sense that every arc $\vec{uv}$ is accompanied with the opposite arc $\vec{vu}$. In this case, we will say that the A-flip is *symmetric* and consider its result to be an undirected graph.

This definition now allows us to formally introduce the flip-reachability logic and its weaker flip-connectivity counterpart.

▶ **Definition 6** (Flip-reachability logic). *For each $k \in \mathbb{N}$ and each $A \subseteq \mathrm{atp}^{k+1} \times \mathrm{atp}^{k+1}$, we introduce the flip-reachability predicate*

$$\mathsf{flipreach}_{k,A}(s, t, a_1, \ldots, a_k)$$

as the relation on vertices in a graph that holds precisely when there exists a directed path from s to t in $G \oplus_{\bar{a}} A$. Flip-reachability logic is first-order logic over graphs in which the universe is formed by the vertices of a graph, and the available relations are the binary edge relation, the unary predicates in the language, as well as the flip-reachability predicates.

▶ **Definition 7** (Flip-connectivity logic). *For $k \in \mathbb{N}$ and symmetric $A \subseteq \mathrm{atp}^{k+1} \times \mathrm{atp}^{k+1}$, we introduce the flip-connectivity predicate*

$$\mathsf{flipconn}_{k,A}(s, t, a_1, \ldots, a_k)$$

as the relation on vertices in a graph that holds precisely when s and t are in the same connected component of $G \oplus_{\bar{a}} A$. Flip-connectivity logic is first-order logic over graphs in which the universe is formed by the vertices of a graph, and the available relations are the binary edge relation, the unary predicates in the language, as well as the flip-connectivity predicates.

Easy comparisons. Let us first establish the straightforward relations between the considered logics in terms of expressive power. We first note the following.

▶ **Lemma 8.** *Every flip-reachability predicate can be expressed in low rank MSO.*

Proof. Fix $k \in \mathbb{N}$, $A \subseteq \mathrm{atp}^{k+1} \times \mathrm{atp}^{k+1}$, a graph G, and vertices $s, t, a_1, \ldots, a_k$; denote $\bar{a} = (a_1, \ldots, a_k)$. Observe that the predicate $\mathsf{flipreach}_{k,A}(s, t, a_1, \ldots, a_k)$ is false if and only if there is a set X of vertices of G such that $s \in X$, $t \notin X$, and there are no arcs from X to $\widehat{X}$ in $G \oplus_{\bar{a}} A$, where we denote $\widehat{X} = V(G) - X$. Note that these conditions, for a given set X, can be encoded by a first-order formula taking $X, s, t, a_1, \ldots, a_k$ as free variables. So it remains to show that quantification over such sets X can be done using a low-rank quantifier, that is, that X has rank bounded by a function of k.

For this, we argue that the set X of vertices of G that are reachable from s in $G \oplus_{\bar{a}} A$ has rank at most $|\mathrm{atp}^{k+1}|$. To see this, consider two vertices $u, u' \in X$ such that $\mathrm{atp}(u, \bar{a}) = \mathrm{atp}(u', \bar{a})$ (in G), as well as a vertex $v \in \widehat{X}$. Note that v is adjacent in G to either both u and u' or to none of them, for otherwise either $\vec{uv}$ or $\vec{u'v}$ would appear as an arc in $G \oplus_{\bar{a}} A$, which would contradict the assumption that v is not reachable from s. Hence, whether there is an edge in G between a vertex u in X and a vertex v in $\widehat{X}$ depends only on the atomic type $\mathrm{atp}(u, \bar{a})$. It follows that the adjacency matrix $\mathsf{Adj}_G[X, \widehat{X}]$ has at most $|\mathrm{atp}^{k+1}|$ distinct rows, hence its rank is at most $|\mathrm{atp}^{k+1}|$. ◀

▶ **Proposition 9.** *The following holds:*

- *For every formula of separator logic there is an equivalent formula of flip-connectivity logic.*

- *For every formula of flip-connectivity logic there is an equivalent formula of flip-reachability logic.*
- *For every formula of flip-reachability logic there is an equivalent formula of low rank* MSO*.*

Proof. For the first point, it suffices to observe that there is a symmetric $A \subseteq \mathrm{atp}^{k+1} \times \mathrm{atp}^{k+1}$ such that for every graph G and $\bar{a} \in V(G)^k$, the flip $G \oplus_{\bar{a}} A$ is equal to G with all edges incident to the vertices of $\bar{a}$ removed; then the predicate $\mathsf{conn}_k(s, t, a_1, \ldots, a_k)$ is equivalent to $\mathsf{flipconn}_{k,A}(s, t, a_1, \ldots, a_k)$. The second point is obvious: flip-connectivity predicates are special cases of flip-reachability predicates obtained by restricting attention to symmetric relations A. The third point follows immediately from Lemma 8. ◀

Finally, we note that flip-connectivity logic has a strictly larger expressive power than separator logic. A distinguishing property is *co-connectivity*: connectivity of the complement of the graph.

▶ **Proposition 10.** *There is a sentence of flip-connectivity logic that verifies whether a graph is co-connected. However, there is no such sentence in separator logic.*

Proof. For the sentence of flip-connectivity logic verifying co-connectivity, we can take

$$\forall s\, \forall t\, \mathsf{flipconn}_{0,A}(s, t), \qquad \text{where } A = \mathrm{atp}^1 \times \mathrm{atp}^1\,.$$

We are left with proving that the property cannot be expressed in separator logic.

For the sake of contradiction, suppose there is a sentence φ of separator logic that holds exactly in graphs whose complements are connected. Since φ is finite, there is some number $k \in \mathbb{N}$ such that all connectivity predicates present in φ are of arity at most $k + 2$. Now, for every integer $n > k + 2$, consider the following two graphs: G_n is the complement of a cycle on n vertices, and H_n is the complement of the disjoint union of two cycles on n vertices. Note that G_n is co-connected while H_n is not, hence φ holds in G_n and does not hold in H_n. Further, it can be easily verified that both G_n and H_n are k-connected, that is, no pair of vertices can be disconnected by a k-tuple of other vertices. Hence, every predicate $\mathsf{conn}_\ell(s, t, a_1, \ldots, a_\ell)$, for $\ell \leqslant k$, is equivalent on $\{G_n, H_n : n > k + 2\}$ to the first-order formula

$$\bigwedge_{i \in [\ell]} (s \neq a_i \wedge t \neq a_i).$$

By replacing all the connectivity predicates with the formulas above, we turn φ into a first-order sentence φ' such that for all $n > k + 2$, φ' holds in G_n and does not hold in H_n. However, a standard argument based on Ehrenfeucht–Fraïsse games for first-order logic shows that such a sentence φ' distinguishing G_n from H_n does not exist. ◀

3 A general framework for proving equivalence with low rank MSO

In the upcoming sections we will prove that low rank MSO has the same expressive power as separator logic on weakly sparse graph classes (Theorem 1), as flip-connectivity logic on graph classes with bounded VC-dimension (Theorem 3), and as flip-reachability logic on all graphs (Theorem 4). All these proofs follow a similar scheme: we prove that quantification over sets of low rank can be emulated in the considered logic. The goal of this section is to provide a framework that captures common features of all three proofs, reducing the task to verifying a concrete property of the considered logic on the considered graph classes.

This key property is captured in the following definition.

▶ **Definition 11.** *We say that a logic $\mathcal{L}$ has* low rank definability property *on a class of (colored) graphs $\mathscr{C}$ if for every $r \in \mathbb{N}$ and formula $\varphi(X, \overline{Y}, \bar{z})$ of $\mathcal{L}$ with free set variables $X, \overline{Y}$ and free vertex variables $\bar{z}$, there is a formula $\psi(x, \bar{t}, \overline{Y}, \bar{z})$ of $\mathcal{L}$ with the following property: For every $G \in \mathscr{C}$, evaluation $\overline{B}$ of $\overline{Y}$ and evaluation $\bar{c}$ of $\bar{z}$; if there is a set $A \subseteq V(G)$ with $\mathrm{rk}(A) \leqslant r$ such that $G \models \varphi(A, \overline{B}, \bar{c})$, then there is a set $A' \subseteq V(G)$ and a tuple $\bar{d} \in V(G)^{\bar{t}}$ such that*

$$G \models \varphi(A', \overline{B}, \bar{c}) \qquad \textit{and} \qquad A' = \psi(G, \bar{d}, \overline{B}, \bar{c}).$$

▶ **Theorem 12.** *Suppose a logic $\mathcal{L}$ is an extension of* FO *and has low rank definability property on a class of (colored) graphs $\mathscr{C}$. Then for every formula of low rank* MSO *there is a formula of $\mathcal{L}$ so that the two formulas are equivalent on all graphs from $\mathscr{C}$.*

Proof. We show that over our class of graphs $\mathscr{C}$, every formula of low rank MSO is equivalent to some formula of $\mathcal{L}$ by induction on the structure of a formula of low rank MSO. The only interesting part of the induction is low rank set quantification, since all other formula constructors of low rank MSO are already present in FO. Consider a formula that begins with such a quantifier:

$$\exists X : r \quad \varphi(X, \overline{Y}, \bar{z}). \tag{1}$$

By induction assumption, the inner formula φ can be expressed in $\mathcal{L}$. We may also assume that for every graph G, a set $A \subseteq V(G)$, and evaluations $\overline{B}$ of $\overline{Y}$ and $\bar{c}$ of $\bar{z}$, the condition $G \models \varphi(A, \overline{B}, \bar{c})$ implies that the rank of A is at most r. Indeed, since $\mathcal{L}$ is an extension of FO, we can express in $\mathcal{L}$ that a definable set (such as A) has rank at most r. This can be done by expressing that there are at most r independent rows/columns in $\mathsf{Adj}_G[A, \widehat{A}]$ – see Lemma 2.2 in the full version. This assertion can always be added to φ. Then, by low rank definability property we have a formula $\psi(x, \bar{t}, \overline{Y}, \bar{z})$ of $\mathcal{L}$ such that for any graph $G \in \mathscr{C}$, whenever there exists a set $A \subseteq V(G)$ of rank at most r satisfying $G \models \varphi(A, \overline{B}, \bar{c})$ for some evaluation $\overline{B}$ of $\overline{Y}$ and $\bar{c}$ of $\bar{z}$, then we have a tuple $\bar{d} \in V(G)^{\bar{t}}$ and a set $A' = \psi(G, \bar{d}, \overline{B}, \bar{c})$ such that $G \models \varphi(A', \overline{B}, \bar{c})$. Since the satisfaction of φ implies that the rank is at most r, we can express the formula (1) in $\mathcal{L}$. ◀

Next, we propose tools for streamlining the verification of the low rank definability property for the considered extensions of FO. Intuitively, for each considered extension and relevant graph class $\mathscr{C}$, we will show a structure theorem showing that every set A of low rank in a graph $G \in \mathscr{C}$ admits a certain structure guarded by a small set of parameters $S \subseteq V(G)$. It will be useful to consider A in a suitable "simplification" of G, which we call the *S-operation* of G. The definition of S-operation depends on the considered logic, as explained formally below. In this definition, we use the notion of *atomic type over S in G*, which corresponds to the atomic type of an arbitrary fixed enumeration $\bar{s}$ of S in G.

▶ **Definition 13.** *Let $\mathcal{L}$ be either separator logic, or flip-connectivity logic, or flip-reachability logic. Let also G be a colored graph and $S \subseteq V(G)$ be a subset of its vertices. We say that a graph G' is an S-*operation *with respect to $\mathcal{L}$ of G if the following holds:*

- *If $\mathcal{L}$ is separator logic, then G' is the graph obtained from G by isolating vertices from S (i.e. removing all edges incident to vertices in S). Further, for each $s \in S$ add a unary predicate that marks the neighborhood of S in G, and a unary predicate that marks only s.*
- *If $\mathcal{L}$ is flip-connectivity logic, then G' is a symmetric flip of G with parameters S. Further, for every atomic type over S in G (with an arbitrary fixed enumeration of S), add a unary predicate that marks the vertices of G of this type.*
- *If $\mathcal{L}$ is flip-reachability logic, then G' is a flip of G with parameters S. Again, for every atomic type over S in G, add a unary predicate that marks the vertices of G of this type.*

Observe that if we fix an arbitrary enumeration of S, then the unary predicates that we add to G' (but not their interpretations) depend only on $|S|$. In other words, we extend the vocabulary of G by adding new unary predicates whose number depends only on $|S|$. Therefore, for a given class of graphs $\mathscr{C}$ we may talk about the language of k-operations, for a fixed $k \in \mathbb{N}$. This is the language that consists of all the relations used in graphs from $\mathscr{C}$ and all the unary predicates added to S-operations of graphs from $\mathscr{C}$ for S of size k; here we use the above observation that for these unary predicates only the size of S matters.

Next, for each considered extension $\mathcal{L}$ of FO, we show that we can freely translate $\mathcal{L}$-formulas working on the graph and on its S-operation. The two next lemmas differ on the place where G and G' appear in the "if and only if" part of their statements.

▶ **Lemma 14** (♠). *Let $\mathcal{L}$ be either separator logic, or flip-connectivity logic, or flip-reachability logic. Let also G be a graph, $S \subseteq V(G)$ be a subset of its vertices, G' be an S-operation with respect to $\mathcal{L}$ of G, and $\varphi(\overline{X}, \bar{y})$ be a formula of $\mathcal{L}$ with free set variables $\overline{X}$ and free vertex variables $\bar{y}$. Then there is a formula $\varphi'(\overline{X}, \bar{y})$ of $\mathcal{L}$, depending only on φ and $|S|$, such that for every evaluation $\overline{A}$ of $\overline{X}$ and $\bar{b}$ of $\bar{y}$ we have*

$$G \models \varphi(\overline{A}, \bar{b}) \qquad \textit{if and only if} \qquad G' \models \varphi'(\overline{A}, \bar{b}).$$

▶ **Lemma 15** (♠). *Let $\mathcal{L}$ be either separator logic, or flip-connectivity logic, or flip-reachability logic. Let also G be a graph, $S \subseteq V(G)$ be a subset of its vertices, $\bar{s}$ be a tuple enumerating all the elements of S, G' be an S-operation of G with respect to $\mathcal{L}$, and $\varphi(\overline{X}, \bar{y})$ be a formula of $\mathcal{L}$ with free set variables $\overline{X}$ and free vertex variables $\bar{y}$. Then there is a formula $\varphi'(\overline{X}, \bar{y}, \bar{z})$ of $\mathcal{L}$, depending only on φ and $|S|$ with $|\bar{z}| = |\bar{s}|$, such that for every evaluation $\overline{A}$ of $\overline{X}$ and $\bar{b}$ of $\bar{y}$ we have*

$$G' \models \varphi(\overline{A}, \bar{b}) \qquad \textit{if and only if} \qquad G \models \varphi'(\overline{A}, \bar{b}, \bar{s}).$$

Lemmas 14 and 15 allow us to reduce verification of the low rank definability property from the original graph to its S-operation, for any set of parameters S of bounded size. As the S-operation will be a much simpler graph, arguing definability there is significantly easier. We use this scheme in the proofs of Theorems 1 and 3. In the proof of Theorem 4, low rank definability property is argued directly.

4 Relation to separator logic

In this section we prove Theorem 1. By Theorem 12, it suffices to show that for every weakly sparse graph class $\mathscr{C}$, separator logic has low rank definability property on $\mathscr{C}$. The plan is as follows. We first prove a combinatorial characterization of low rank sets in weakly sparse graph classes. Next, we use this characterization to prove low rank definability property.

For the characterization, we prove that every low rank set in a graph from a weakly sparse graph class is very close to a (vertex) separation of small order. Let us first recall some graph-theoretic terminology. A *separation* of a graph G is a pair (L, R) of vertex subsets such that $L \cup R = V(G)$ and there is no edge with one endpoint in $L - R$ and the other in $R - L$. The *order* of the separation (L, R) is $|L \cap R|$, i.e. the number of vertices the two sides have in common. We shall say that a set of vertices $X \subseteq V(G)$ is *captured* by a separation (L, R) if we have $L - R \subseteq X \subseteq L$. That is, X has to contain all the vertices that lie strictly in the left side of the separation (set $L - R$) and may additionally contain a subset of the separator (set $L \cap R$). In this terminology, the combinatorial characterization reads as follows.

▶ **Lemma 16.** *Let G be a graph that does not contain the complete bipartite graph $K_{t,t}$ as a subgraph. Let X be a set of rank at most r in G. Then X is captured by some separation of G of order at most $2^{r+1}(t-1)$.*

Proof. Denote $\widehat{X} \coloneqq V(G)-X$ and consider the bipartite adjacency matrix $M \coloneqq \mathsf{Adj}_G[X,\widehat{X}]$. Call a row of M *frequent* if it appears in M at least t times, that is, there are at least $t-1$ other rows equal to it. Define *frequent columns* of M analogously. Note that the entry at the intersection of every frequent row and every frequent column must be 0, for otherwise the at least t equal rows and the at least t equal columns would induce a $t \times t$ submatrix of M entirely filled with 1s, which in turn would yield a $K_{t,t}$ subgraph in G. This proves that the following sets L and R form a separation of G:

- L comprises all the vertices of X and those vertices of $\widehat{X}$ whose columns are non-frequent; and
- R comprises all the vertices of $\widehat{X}$ and those vertices of X whose rows are non-frequent.

Clearly, (L,R) captures X. Further, $L \cap R$ comprises all the vertices of X whose rows are non-frequent and all the vertices of $\widehat{X}$ whose columns are non-frequent. Since $\mathrm{rk}(X) \leqslant r$, the rank of M over $\mathbb{F}_2$ is at most r, hence M has at most 2^r distinct rows and at most 2^r distinct columns. Since every non-frequent row occurs at most $t-1$ times in M, M has at most $2^r(t-1)$ non-frequent rows; and similarly M has at most $2^r(t-1)$ non-frequent columns. We conclude that $|L \cap R| \leqslant 2^{r+1}(t-1)$, as required. ◀

In the next lemma we verify the low rank definability property for $r=0$ and formulas without additional free variables. This is the key step of the proof, where we use the idempotence property of separator logic in order to turn set quantification into first-order quantification. Note here that if X is a set of rank 0 in a graph G, then every connected component of G is either entirely contained in X or is disjoint with X. Hence, X is the union of the vertex set of a collection of connected components of G.

▶ **Lemma 17.** *Let $\varphi(X)$ be a formula of separator logic with one free set variable X. Then there exists a formula $\varphi'(x,\bar{t})$ of separator logic, where x and $\bar{t}$ are free vertex variables, such that the following holds. For every graph G, if there exists a set $A \subseteq V(G)$ of rank 0 such that $G \models \varphi(A)$, then there exists $A' \subseteq V(G)$ of rank 0 and an evaluation $\bar{d}$ of variables $\bar{t}$ such that*

$$G \models \varphi(A') \qquad \text{and} \qquad A' = \varphi'(G,\bar{d}).$$

Proof. Let Σ be the set of unary predicates used by φ, and let q be the quantifier rank of φ. We will use the standard notion of *types* tailored to separator logic. Concretely, for a Σ-colored graph H, the *type* of H is the set of all sentences (up to logical equivalence) of separator logic in the language of Σ-colored graphs with quantifier rank at most $q+1$ that are satisfied in H. A standard induction argument shows that for fixed q and Σ, there are only boundedly many (in terms of q and Σ) non-equivalent such sentences, hence also only boundedly many types. Let $\mathcal{T}$ be the set of all types; then $|\mathcal{T}|$ is bounded by a function of q and Σ. Observe that for each $\tau \in \mathcal{T}$, there is a sentence ψ_τ such that $H \models \psi_\tau$ if and only if H has type τ. Indeed, it suffices to take ψ_τ to be the conjunction of all the sentences belonging to τ.

Let $\mathcal{C}$ be the set of connected components of G and let $\{\mathcal{C}_\tau : \tau \in \mathcal{T}\}$ be the partition of $\mathcal{C}$ according to the types of the individual components; that is, a component $C \in \mathcal{C}$ belongs to $\mathcal{C}_\tau$ iff the type of C is τ. The following claim follows from a standard argument involving Ehrenfeucht–Fraïssé games.

▷ Claim 18. There exists a constant p, depending only on q, such that the following holds. Suppose for some type $\tau \in \mathcal{T}$, A contains more than p components from $\mathcal{C}_\tau$, and A is also disjoint with more than p components from $\mathcal{C}_\tau$. Then for every $C \in \mathcal{C}_\tau$ contained in A, we have $G \models \varphi(A - C)$.

An immediate corollary of Claim 18 is the following.

▷ Claim 19. Suppose there exists $A \subseteq V(G)$ of rank 0 such that $G \models \varphi(A)$. Then there also exists $A' \subseteq V(G)$ of rank 0 such that the following holds:
- $G \models \varphi(A')$; and
- for each $\tau \in \mathcal{T}$, A' contains at most p components of $\mathcal{C}_\tau$ or A' contains all but at most p components of $\mathcal{C}_\tau$.

It remains to construct a sentence $\varphi'(x, \bar{t})$ that will define some set A' whose existence is asserted by Claim 19. For this, for each $\tau \in \mathcal{T}$ construct a p-tuple of variables $\bar{t}_\tau$, and let $\bar{t}$ be the union of $\bar{t}_\tau$ over all $\tau \in \mathcal{T}$. For every function $g \colon \mathcal{T} \to \{0, 1\}$, we may use the sentences $\{\psi_\tau : \tau \in \mathcal{T}\}$ to write a formula $\alpha_g(x, \bar{t})$ expressing the following assertion: supposing x belongs to a component C whose type is τ, we have that C contains a vertex of $\bar{t}_\tau$ if $g(\tau) = 0$, and C does not contain a vertex of $\bar{t}_\tau$ if $g(\tau) = 1$. Then, we may write $\varphi'(x, \bar{t})$ as the formula expressing the following:
- for some $g \colon \mathcal{T} \to \{0, 1\}$, we have $G \models \varphi(\alpha_g(G, \bar{t}))$; and
- $\alpha_{g^\circ}(G, \bar{t})$ is true, where g° is the lexicographically smallest function from $\mathcal{T}$ to $\{0, 1\}$ satisfying the above.

By Claim 19, for some evaluation $\bar{d}$ of $\bar{t}$ and some function $g \colon \mathcal{T} \to \{0, 1\}$, we have $A' = \alpha_g(G, \bar{d})$. As $G \models \varphi(A')$, the first point above holds. Then we have $\varphi'(G, \bar{d}) = \alpha_{g^\circ}(G, \bar{d})$, where g° is as in the second point. And by construction, this set has rank 0 and satisfies φ. ◀

We now generalize the conclusion of Lemma 17 to low rank definability property using Lemma 16 in combination with Lemmas 14 and 15.

▶ **Lemma 20.** *For every weakly sparse graph class $\mathscr{C}$, separator logic has low rank definability property on $\mathscr{C}$.*

Proof. Let $r \in \mathbb{N}$ be a bound on the rank and $\varphi(X, \bar{Y}, \bar{z})$ be a formula of separator logic, as in the definition of low rank definability property. We may assume that $\overline{Y} = \emptyset$ and $\bar{z} = \emptyset$. Indeed, otherwise we may expand the signature by $|\overline{Y}| + |\bar{z}|$ unary predicates marking the free variables of φ, and perform the reasoning on the modified formula φ that treats the free variables of $\overline{Y}$ and $\bar{z}$ as unary predicates present in the structure. Once we obtain a suitable formula φ', the additional unary predicates can be interpreted back as free variables. Hence, from now φ has only one free set variable X.

Since $\mathscr{C}$ is weakly sparse, there is some $t \in \mathbb{N}$ such that no graph from $\mathscr{C}$ contains $K_{t,t}$ as a subgraph. Consider any $G \in \mathscr{C}$ and $A \subseteq V(G)$ with $\mathrm{rk}(A) \leqslant r$. By Lemma 16, there is a separation (L, R) of G of order at most $p := 2^{r+1} \cdot (t - 1)$ that captures A. Let $S := L \cap R$; then $|S| \leqslant p$. Observe that since (L, R) captures A, A is the union of a subset of S and a collection of connected components of $G - S$. In other words, A is a set of rank 0 in G' – the S-operation of G with respect to the separator logic.

By Lemma 14, there is a formula $\psi(X)$ depending only on φ and p such that for every $B \subseteq V(G)$, we have $G \models \varphi(B)$ iff $G' \models \psi(B)$. In particular, $G' \models \psi(A)$. By Lemma 17, we may find a formula $\psi'(x, \bar{t})$ depending only on ψ and p, an evaluation $\bar{d}$ of variables $\bar{t}$, and a set $A' \subseteq V(G)$ of rank 0 in G', such that $G' \models \psi(A')$ and $A' = \psi'(G', \bar{d})$. By Lemma 15, there is a formula $\varphi'(x, \bar{t}, \bar{t}')$, depending only on ψ' and p, such that for every $u \in V(G)$ and evaluation $\bar{c}$ of $\bar{t}$, we have $G \models \varphi'(u, \bar{c}, \bar{s})$ iff $G' \models \psi'(u, \bar{c})$, where $\bar{s}$ is the tuple of elements from S. All in all, we have $G \models \varphi(A')$ and $A' = \varphi'(G, \bar{d}, \bar{s})$. Therefore, all the conditions required from φ', A' are met. ◀

Now Theorem 1 follows immediately by combining Lemma 20 with Theorem 12.

5 Relation to flip-connectivity logic

In this section we compare the expressive power of low rank MSO with flip-connectivity logic.

5.1 Negative result: Where low-rank MSO is more expressive

As our first result we show that on the class of all graphs low rank MSO is strictly more expressive than flip-connectivity logic. That is, we prove Theorem 2. In the proof we construct a family of graphs G_n and G'_n that can be distinguished by a sentence of low rank MSO, but for every sentence ψ of flip-connectivity logic there is some n such that G_n and G'_n are indistinguishable by ψ.

We begin with an auxiliary graph, which we call F_n. In this graph, there are $n+1$ layers of vertices. There are $10n^2$ vertices in layer 0. The vertices in layer $i \in \{1, \dots, n\}$ are all subsets of vertices in layer $i-1$ that have size $5n$, and the edge relation between layers i and $i-1$ represents membership. Thus, layer 1 has $\binom{10n^2}{5n}$ vertices, layer 2 has $\binom{\binom{10n^2}{5n}}{5n}$ vertices, and so on. Here is a picture:

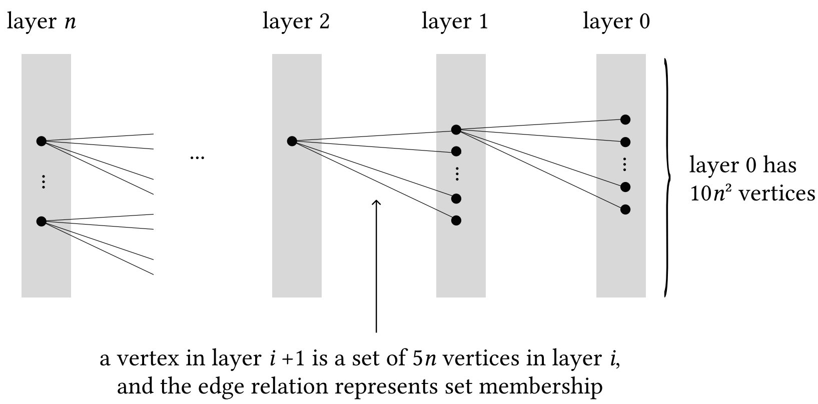

Using the auxiliary graph F_n, we define the two colored graphs G_n and G'_n. Actually, the underlying uncolored graphs of G_n and G'_n will be exactly the same; the difference between them is only in the placement of colors (unary predicates). The uncolored graph underlying G_n and G'_n is constructed as follows:

- take four disjoint copies of F_n;
- merge the first two copies of layer 0;
- merge the last two copies of layer 0;
- create a biclique between the merged copies from the previous two items.

Figure 1 depicts this graph. This completes the description of the uncolored graph underlying G_n and G'_n. To differentiate them, we add four unary predicates that mark vertices of n'th layer in each copy of F_n. These, call them A, B, C, D, are distributed as follows:

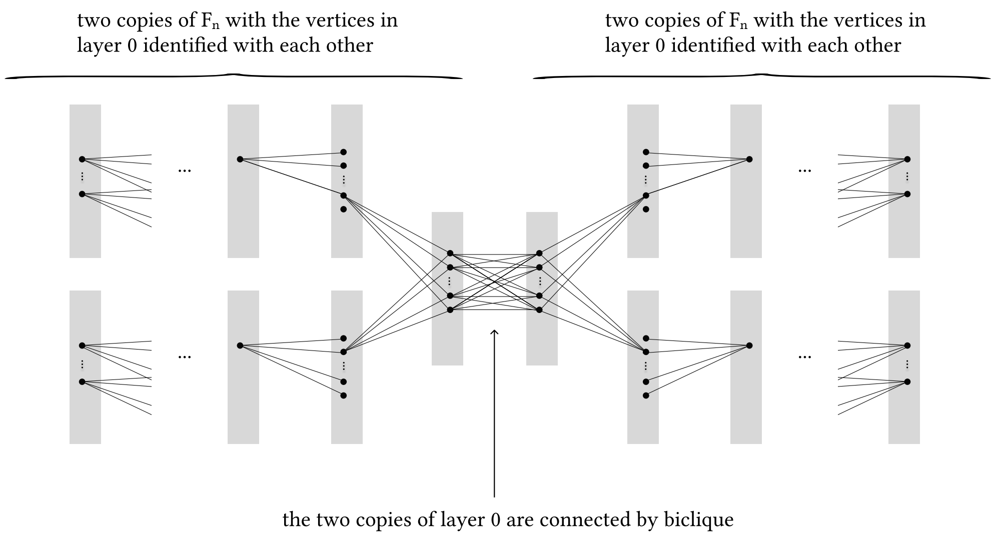

Figure 1 The uncolored graph underlying G_n and G'_n.

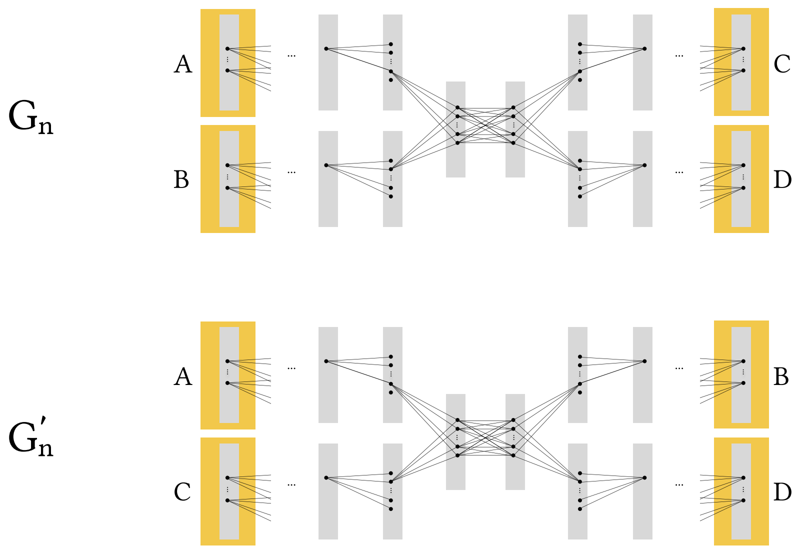

Note that, by construction, all four sets of vertices marked with the predicates A, B, C, and D respectively, are of equal size. It can be easily verified (♠) that G_n and G'_n can be distinguished by a sentence of low rank MSO expressing the following property; more precisely, G_n satisfies the property, while G'_n does not.

There is a set of rank 1 *that contains all the vertices marked with predicate A but no vertices marked with predicate C.*

The hard part of the proof is to prove that for any sentence ψ of flip-connectivity logic, ψ cannot distinguish G_n from G'_n for n large enough. For this, let $q \in \mathbb{N}$ be such that both the quantifier rank and the arity of all flip-connectivity predicates present in ψ is bounded by q. We argue that for n large enough depending on q, any flip of G_n or of G'_n definable from at most q parameters has a very simple structure of connected components: there is one huge connected component containing almost all the vertices, and every other connected component has bounded diameter. This implies that the connectivity relation in this flip can be expressed in FO, which allows us to rewrite flip-connectivity predicates appearing in ψ to formulas of plain FO. We thus obtain an FO sentence ψ' that distinguishes G_n from G'_n for all n large enough; but a standard argument based on Ehrenfeucht–Fraïssé games shows that such a sentence cannot exist. Details of this reasoning can be found in the full version of the paper. We would also like to point that, using standard arguments, the unary predicates in the above construction can be simulated by constant size gadgets, therefore leading to a sentence for Theorem 2 that works in graphs without colors.

5.2 Positive result

We continue to our positive result for flip-connectivity logic, i.e. Theorem 3. Let us start with recalling the notion of VC dimension and stating facts about it that will be of use.

A *set system* over a universe U is just a family $\mathcal{F}$ of subsets of U. For a subset of the universe $X \subseteq U$, we say that X is *shattered* by $\mathcal{F}$ if for every $Y \subseteq X$ there exists $F \in \mathcal{F}$ such that $F \cap X = Y$. The *VC dimension* of $\mathcal{F}$ is the largest cardinality of a set shattered by $\mathcal{F}$. The VC dimension of a graph G is the VC dimension of the set system of neighborhoods $\{N(v) \colon v \in V(G)\}$; and a graph class $\mathscr{C}$ has *bounded VC dimension* if there is $d \in \mathbb{N}$ such that every member of $\mathscr{C}$ has VC dimension at most d.

Next, we will need the following definition of duality of a binary relation, and its connection to the notion of VC dimension. The following definitions and results are taken from [3].

▶ **Definition 21.** *Let $E \subseteq A \times B$ be a binary relation. We say that E has a* duality *of order k if at least one of two cases holds:*

- *there exists $A_0 \subseteq A$ of size at most k such that for every $b \in B$ there is some $a \in A_0$ with $\neg E(a,b)$, or*
- *there exists $B_0 \subseteq B$ of size at most k such that for every $a \in A$ there is some $b \in B_0$ with $E(a,b)$.*

▶ **Theorem 22** (see [3]). *For every $d \in \mathbb{N}$ there is some $k \in \mathbb{N}$ such that the following holds. Let $E \subseteq A \times B$ be a binary relation with both A and B finite such that the set system $\{\{b \in B \mid E(a,b)\} \colon a \in A\}$ has VC dimension at most d. Then E has a duality of order k.*

We will use the following result proven by Bonnet, Dreier, Gajarský, Kreutzer, Mählmann, Simon, and Toruńczyk [3]. Here, a *pseudometric* is a symmetric function $\delta \colon V \times V \to \mathbb{R}_{\geqslant 0} \cup \{+\infty\}$ satisfying the triangle inequality; and for a partition $\mathcal{P}$ of a set V, by $\mathcal{P}(v)$ we denote the unique part of $\mathcal{P}$ containing v.

▶ **Theorem 23** ([3, Theorem 3.5]). *Fix $r, k, t \in \mathbb{N}$. Let V be a finite set equipped with:*

- *a binary relation $E \subseteq V \times V$ such that for all $A \subseteq V$ and $B \subseteq V$, $E \cap (A \times B)$ has a duality of order k,*
- *a pseudometric $\mathrm{dist} \colon V \times V \to \mathbb{R}_{\geqslant 0} \cup \{+\infty\}$, and*
- *a partition $\mathcal{P}$ of V with $|\mathcal{P}| \leqslant t$.*

Suppose further that $E(u,v)$ depends only on $\mathcal{P}(u)$ and $\mathcal{P}(v)$, for all u, v with $\mathrm{dist}(u,v) > r$. (That is, if $\mathrm{dist}(u,v) > r$, $\mathrm{dist}(u',v') > r$, $\mathcal{P}(u) = \mathcal{P}(u')$, and $\mathcal{P}(v) = \mathcal{P}(v')$, we have $E(u,v) \iff E(u',v')$.) Then there is a set $S \subseteq V$ of size $\mathcal{O}(kt^2)$ such that $E(u,v)$ depends only on $E(u,S)$ and $E(S,v)$ for all $u, v \in V$ with $\mathrm{dist}(u,v) > 5r$.

Now we proceed to the proof of Theorem 3. We start with the following combinatorial lemma that is the analogue of Lemma 16 and characterizes low rank sets in graphs of bounded VC dimension.

▶ **Lemma 24.** *Fix a constant $r \in \mathbb{N}$ and a class $\mathscr{C}$ of graphs of bounded VC dimension. Then there exists a constant $\ell \in \mathbb{N}$ such that for every graph $G \in \mathscr{C}$ and every set $A \subseteq V(G)$ of rank at most r, there is a set $S \subseteq V(G)$ of size at most ℓ and a symmetric flip G' of G with parameters S such that A is the union of the vertex sets of a collection of connected components of G'.*

Proof. Since $\mathscr{C}$ has bounded VC dimension, there is some $k \in \mathbb{N}$ such that the edge relation of every graph $G \in \mathscr{C}$ has a duality of order k. Take a graph $G \in \mathscr{C}$ and a subset $A \subseteq V(G)$ of rank at most r. Partition A according to the edge types on the complement of A: two vertices $u, v \in A$ are in the same part if for every $w \in V(G) - A$, we have $E(u, w) \iff E(v, w)$. Since A has rank at most r, the number of parts is at most 2^r. Similarly, partition the complement of A according to the edge types on A; again we get at most 2^r parts. In this way we obtain a partition $\mathcal{P}$ of $V(G)$ into at most 2^{r+1} parts.

Define a pseudometric on $V(G)$ as follows: $\mathrm{dist}(u, v) = 0$ if both $u, v \in A$ or both $u, v \in V(G) - A$, and $\mathrm{dist}(u, v) = +\infty$ otherwise. Note that $E(u, v)$ depends only on $\mathcal{P}(u)$ and $\mathcal{P}(v)$, for all u, v with $\mathrm{dist}(u, v) > 1$. Therefore, by Theorem 23 there is a set $S \subseteq V(G)$ of size at most ℓ for some constant $\ell = \ell(k, r)$ such that $E(u, v)$ depends only on $E(u, S)$ and $E(S, v)$ for all $u, v \in V(G)$ with $\mathrm{dist}(u, v) > 5$. In particular, for every $u \in A$ and $v \in V(G) - A$ we have that $E(u, v)$ depends only on $E(u, S)$ and $E(S, v)$. Therefore, there exists a symmetric flip G' of G with parameters S that has no edges between A and $V(G) - A$. Naturally, then A is the union of the vertex sets of a collection of connected components of G'. ◀

With the combinatorial characterization of low rank sets in graphs of bounded VC dimension provided by Lemma 24, the remainder of the proof of Theorem 3 proceeds in direct analogy to Theorem 1. We first use the same reasoning as in the proof of Lemma 17 to establish the low rank definability property for sets of rank 0, and then we apply Lemmas 14 and 15 together with the characterization of Lemma 24 to establish the low rank definability property in full generality. Details can be found in the full version.

6 Relation to flip-reachability logic

In this section we show that low rank MSO and flip-reachability logic have the same expressive power on all undirected graphs (Theorem 4). The main part of the work is a combinatorial characterization of sets of low rank in undirected graphs, captured in the Low Rank Structure Theorem (Theorem 25). Before presenting it, we give some necessary definitions.

For an undirected graph G and $r \in \mathbb{N}$, we define $\mathsf{LowRank}^r(G)$ to be the family of all vertex sets in G of rank at most r:

$$\mathsf{LowRank}^r(G) \coloneqq \{X \subseteq V(G) \mid \mathrm{rk}(X) \leqslant r\}.$$

Our main goal in this section is to give an effective representation of $\mathsf{LowRank}^r(G)$. More precisely, we will show that $\mathsf{LowRank}^r(G)$ can be written as the union of $|V(G)|^{\mathcal{O}_r(1)}$ *well-structured* families of subsets of $V(G)$, and that each such family can be defined using $\mathcal{O}_r(1)$ vertices of G – here $\mathcal{O}_r(1)$ is a shortcut for $\mathcal{O}(f(r))$, for some function $f : \mathbb{N} \to \mathbb{N}$. Formal definitions follow.

Let G be an undirected graph. A *seed* of G is a tuple $(X_+, X_-, \mathcal{X})$, where

- $X_+, X_- \subseteq V(G)$,
- $\mathcal{X}$ is a collection of nonempty subsets of $V(G)$ (called *parts*), and
- $\mathcal{X} \cup \{X_+, X_-\}$ is a partition of $V(G)$.

The family *spanned* by a seed $(X_+, X_-, \mathcal{X})$ is

$$\mathsf{Span}(X_+, X_-, \mathcal{X}) \coloneqq \{X_+ \cup \bigcup \mathcal{Y} : \mathcal{Y} \subseteq \mathcal{X}\}.$$

In our characterization of $\mathsf{LowRank}^r(G)$, the obtained families covering $\mathsf{LowRank}^r(G)$ can be spanned by seeds with two additional structural restrictions: first, they need to satisfy a specific uniformity condition, which we will introduce in a moment; and second, the seeds themselves need to be defined in terms of at most $\mathcal{O}_r(1)$ vertices of G via fixed formulas of flip-reachability logic.

For any seed $(X_+, X_-, \mathcal{X})$ and $u \in \bigcup \mathcal{X}$, we define $\mathcal{X}(u) \in \mathcal{X}$ as the unique part of $\mathcal{X}$ containing u. Then, given a tuple $\bar{a}$ of vertices of G, we say that a seed $(X_+, X_-, \mathcal{X})$ is *$\bar{a}$-uniform* if for every $u_1, u_2 \in \bigcup \mathcal{X}$ with $\mathrm{atp}(u_1, \bar{a}) = \mathrm{atp}(u_2, \bar{a})$, we have that $N_G(u_1) - (\mathcal{X}(u_1) \cup \mathcal{X}(u_2)) = N_G(u_2) - (\mathcal{X}(u_1) \cup \mathcal{X}(u_2))$. In other words, whenever u_1 and u_2 are twins with respect to $\bar{a}$, they are also twins with respect to $V(G) - (\mathcal{X}(u_1) \cup \mathcal{X}(u_2))$.

Next, we say that a seed $(X_+, X_-, \mathcal{X})$ is *defined* by formulas φ_+, φ_-, $\varphi_\sim$ of flip-reachability logic and a k-tuple of vertices $\bar{a}$ if:

- $X_+ = \{v \in V(G) : G \models \varphi_+(v, \bar{a})\}$,
- $X_- = \{v \in V(G) : G \models \varphi_-(v, \bar{a})\}$, and
- two vertices $u, v \in V(G) - (X_+ \cup X_-)$ belong to the same part of $\mathcal{X}$ if and only if $G \models \varphi_\sim(u, v, \bar{a})$.

Note that, assuming that $\varphi_+(\cdot, \bar{a})$ and $\varphi_-(\cdot, \bar{a})$ describe disjoint subsets of $V(G)$ and $\varphi_\sim(\cdot, \cdot, \bar{a})$ is an equivalence relation on $V(G)$, there exists a unique seed defined by $\varphi_+, \varphi_-, \varphi_\sim$ and $\bar{a}$. We denote this seed by $\mathsf{Seed}(\varphi_+, \varphi_-, \varphi_\sim, \bar{a})$.

We are now ready to state our structural characterization of vertex sets of bounded rank.

▶ **Theorem 25** (Low Rank Structure Theorem). *Let $r \in \mathbb{N}$. Then there exist an integer $k \coloneqq k(r)$ and formulas φ_+, φ_-, $\varphi_\sim$ of flip-reachability logic with the following properties:*

- *For any undirected graph G and k-tuple $\bar{a} \in V(G)^k$, we have that $\varphi_+, \varphi_-, \varphi_\sim, \bar{a}$ define an $\bar{a}$-uniform seed of G.*
- *For any undirected graph G,*

$$\mathsf{LowRank}^r(G) = \bigcup_{\bar{a} \in V(G)^k} \mathsf{Span}(\mathsf{Seed}(\varphi_+, \varphi_-, \varphi_\sim, \bar{a})).$$

The proof of the Low Rank Structure Theorem (Theorem 25) is based on the following two lemmas. The first states that the family of low-rank vertex sets in G can be characterized as *suffixes* of certain digraphs that are flips of G parameterized by a bounded number of vertices.

▶ **Lemma 26** (♠). *Let $r \in \mathbb{N}$. Then there exist an integer $k \coloneqq k(r)$ and a binary relation $A \subseteq \mathrm{atp}^{k+1} \times \mathrm{atp}^{k+1}$ such that, for any undirected graph G, we have*

$$\mathsf{LowRank}^r(G) = \bigcup_{\bar{a} \in V(G)^k} \mathsf{Suffixes}(G \oplus_{\bar{a}} A),$$

where $\mathsf{Suffixes}(H)$ denotes the set of suffixes of a directed graph H:

$$\mathsf{Suffixes}(H) \coloneqq \{S \subseteq V(H) : \forall_{\vec{uv} \in E(H)}\, u \in S \Rightarrow v \in S\}.$$

The proof of Lemma 26 goes roughly as follows. For every set X of low rank there is a small number of representative vertices in X and in $\widehat{X} := V(G) - X$ representing the adjacencies between these two sets [16]. Our aim is to recover X, given a guess of representatives. For vertices that have a twin representative only on one side of the cut $(X, \widehat{X})$, we know whether we should include them in X or in $\widehat{X}$. For the rest of the vertices, we show that there is a flip $H := G \oplus_{\bar{a}} A$, with $\bar{a}$ consisting precisely of the representative vertices, with the property that whenever $\vec{uv}$ is an arc of H, then $u \in X$ implies $v \in X$. This is the digraph whose suffix is X. Details can be found in the full version of the paper.

The next lemma essentially shows that the family of all suffixes of a directed graph can be described in a concise way, provided that the graph is a definable flip of an undirected graph.

▶ **Lemma 27 (♠).** *Let $k \in \mathbb{N}$ and $A \subseteq \mathrm{atp}^{k+1} \times \mathrm{atp}^{k+1}$. Then there exist $\ell \in \mathbb{N}$ and formulas φ_+, φ_-, $\varphi_\sim$ of flip-reachability logic with the following properties:*
- *For any undirected graph G and any $(k+\ell)$-tuple $\bar{a}\bar{b} \in V(G)^{k+\ell}$ with $|\bar{a}| = k$, $|\bar{b}| = \ell$, the quadruple $\varphi_+, \varphi_-, \varphi_\sim, \bar{a}\bar{b}$ defines an $\bar{a}$-uniform seed of G.*
- *For any undirected graph G and any k-tuple $\bar{a} \in V(G)^k$,*

$$\mathsf{Suffixes}(G \oplus_{\bar{a}} A) = \bigcup_{\bar{b} \in V(G)^\ell} \mathsf{Span}(\mathsf{Seed}(\varphi_+, \varphi_-, \varphi_\sim, \bar{a}\bar{b})).$$

The intuition behind the proof of Lemma 27 is as follows. Given a digraph H, the suffixes of H are unions of collections of strongly connected components of H. Therefore, in each seed $(X_+, X_-, \mathcal{X})$ that we shall construct, both sets X_+ and X_- will be unions of strongly connected components of H (with X_+ being a suffix and X_- a prefix of H), while the remaining components – the parts of $\mathcal{X}$ – form an *antichain* residing "between X_+ and X_-": they are pairwise not reachable from each other. The key step in the proof is to show that although the antichain forming $\mathcal{X}$ can a priori consist of an unbounded number of components, it can be described concisely using a bounded-size tuple of vertices. This argument crucially exploits the assumption that H is a flip of an undirected graph G defined by some k parameters.

With the Low Rank Structure Theorem (Theorem 25) in hand, the only missing ingredient to complete the proof of Theorem 4 is the low rank definability property for flip-reachability logic. This is done using the formulas provided by the Low Rank Structure Theorem and by additionally arguing that only a small number of parts from $\mathcal{X}$ is essential for the satisfaction of a given formula (analogously to Lemma 17). The uniformity condition of the seed is crucial for the correctness of this argument. Details in the full version.

7 Conclusions

Beyond graphs. In this work we considered low rank MSO on undirected graphs. However, the definition can be naturally extended to any kind of structures in which a meaningful notion of the rank of a set can be considered. Take for example binary relational structures. For a partition $(X, \widehat{X})$ of the universe of a structure $\mathbb{A}$, we can define the bipartite adjacency matrix $\mathsf{Adj}_{\mathbb{A}}[X, \widehat{X}]$ by placing at the intersection of the row of $u \in X$ and the column of $v \in X$, the atomic type of the pair (u, v). This way, the adjacency matrix is no longer binary, but we can still measure its *diversity*: the number of distinct rows plus the number of distinct columns. (See the discussion in Section 2 of the full version.) Now, in low rank MSO over binary structures we would stipulate the set quantification to range only over sets that induce cuts of bounded diversity. For structures of higher arity, say bounded by k, one natural definition would be to index rows of $\mathsf{Adj}_{\mathbb{A}}[X, \widehat{X}]$ by $(< k)$-tuples of elements of X, the columns by $(< k)$-tuples of elements of Y, and at the intersection of row $\bar{u} \in X^{<k}$ and column $\bar{v} \in \widehat{X}^{<k}$ put the atomic type of $\bar{u}\bar{v}$.

To see an example of this definition in practice, consider the setting of finite words over a finite alphabet Σ. There are two ways of encoding such words as relational structures: either by equipping the set of positions by the total order, or only by the successor relation. Deploying first-order logic FO on these two encodings yields two logics with different expressive power. It is customary to consider the ordered one as the right notion of FO on words, mainly due to classic connections with star-free expressions [14] and aperiodic monoids [21]. In contrast, it is not difficult to see that in both encodings, sets of bounded rank can be characterized as unions of a bounded number of intervals of positions. Consequently, regardless of the encoding, low rank MSO on finite words is equivalent to FO with access to the total order on positions.

Another case is that of directed graphs, which can be seen as binary structures consisting of the vertex set equipped with a (not necessarily symmetric) arc relation. Recall that we proved that in undirected graphs, properties expressible in low rank MSO can be decided in polynomial time (Corollary 5), and this was a consequence of the equivalence of low rank MSO and flip-reachability logic (Theorem 4). We do not know whether this equivalence holds in general directed graphs, hence the following question is open.

▶ **Question 28.** *Can every property of directed graphs definable in low rank MSO be decided in polynomial time?*

Finally, another interesting class of structures on which low rank MSO can be deployed are matroids. This setting is slightly tricky, as it is natural to take the connectivity function as the concept of the rank of a set, instead of the standard matroid rank function (i.e., maximum size of an independent subset). That is, if $\mathcal{M}$ is a matroid with ground set E and $(X, \widehat{X})$ is a partition of E, then the connectivity function of X is $\mathrm{rk}_{\mathcal{M}}(X) + \mathrm{rk}_{\mathcal{M}}(\widehat{X}) - \mathrm{rk}_{\mathcal{M}}(E)$, where $\mathrm{rk}_{\mathcal{M}}(\cdot)$ is the standard rank function of $\mathcal{M}$; this way the connectivity function of X and of $\widehat{X}$ are equal. Given the tight links between the structural theory of matroids and the theory of vertex-minors, where sets of small cutrank play a key role, one might suspect that low rank MSO has interesting properties on matroids as well.

Monadic dependence. The recent developments on computational aspects of FO have highlighted the notion of *monadically dependent* graph classes as a key structural dividing line. In a nutshell, a class of graphs $\mathscr{C}$ is monadically dependent if one cannot interpret the class of all graphs in vertex-colored graphs from $\mathscr{C}$ using a fixed FO interpretation; see [17, Section 4.1] for a formal definition and a discussion. The notion of monadic dependence can be naturally defined also for logics other than FO. What would be then monadic dependence with respect to low rank MSO? It is not hard to prove using the results of Pilipczuk et al. [18] that on subgraph-closed classes, monadic dependence with respect to separator logic coincides with the property of excluding a fixed graph as a topological minor. So by analogy, monadic dependence with respect to low rank MSO should be a dense, logically-motivated analogue of excluding topological minors. It must be, however, different from the property of excluding a fixed vertex-minor, for the latter is known not to imply monadic dependence with respect to FO [10].

We pose the following question as an excuse to explore combinatorial properties of graph classes that are monadically dependent with respect to low rank MSO. It is a low rank MSO counterpart of the analogous question posed for FO (see e.g. [17, Conjecture 2]), which is currently under intensive investigation.

▶ **Question 29.** *Is the model-checking problem for low rank* MSO *fixed-parameter tractable on every class of graphs that is monadically dependent with respect to low rank* MSO*?*

Space complexity. Note that every graph property definable in FO can be decided in L (deterministic logarithmic space): a brute-force model-checking algorithm needs only to remember constantly many variables at a time. Due to Reingold's result that undirected reachability is in L [19], the same can be said about separator logic and flip-connectivity logic; hence also about low rank MSO on any graph class of bounded VC dimension, by Theorem 3. However, the argument breaks for flip-reachability logic, for testing flip-reachability predicates requires solving directed reachability, which is NL-complete. Consequently, thanks to Immerman–Szelepcsényi Theorem [11, 22] and Theorem 4, we can place the model-checking problem for low rank MSO in slicewise NL, but membership in slicewise L is unclear. We ask the following.

▶ **Question 30.** *Can every property of undirected graphs definable in low rank* MSO *be decided in deterministic logarithmic space?*

References

1 Mikołaj Bojańczyk. Separator logic and star-free expressions for graphs. *ArXiv preprint*, abs/2107.13953, 2021. `arXiv:2107.13953`.

2 Mikołaj Bojańczyk, Michał Pilipczuk, Wojciech Przybyszewski, Marek Sokołowski, and Giannos Stamoulis. Low rank MSO. *ArXiv preprint*, abs/2502.08476, 2025. `doi:10.48550/arXiv.2502.08476`.

3 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 *37th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2022*, pages 54:1–54:13. ACM, 2022. `doi:10.1145/3531130.3533367`.

4 Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, March 1990. `doi:10.1016/0890-5401(90)90043-H`.

5 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. `doi:10.1007/s002249910009`.

6 Jan Dreier, Ioannis Eleftheriadis, Nikolas Mählmann, Rose McCarty, Michał Pilipczuk, and Szymon Toruńczyk. First-Order model checking on monadically stable graph classes. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS2024*, pages 21–30. IEEE, 2024. `doi:10.1109/FOCS61266.2024.00012`.

7 Jan Dreier, Nikolas Mählmann, Sebastian Siebertz, and Szymon Toruńczyk. Indiscernibles and flatness in monadically stable and monadically NIP classes. In *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023*, volume 261 of *LIPIcs*, pages 125:1–125:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. `doi:10.4230/LIPIcs.ICALP.2023.125`.

8 Jan Dreier, Nikolas Mählmann, and Szymon Toruńczyk. Flip-breakability: A combinatorial dichotomy for monadically dependent graph classes. In *56th Annual ACM Symposium on Theory of Computing, STOC 2024*, pages 1550–1560. ACM, 2024. `doi:10.1145/3618260.3649739`.

9 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 *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023*, volume 261 of *LIPIcs*, pages 128:1–128:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. `doi:10.4230/LIPIcs.ICALP.2023.128`.

10 Petr Hliněný and Filip Pokrývka. Twin-width and limits of tractability of FO model checking on geometric graphs. *ArXiv preprint*, abs/2204.13742, 2022. doi:10.48550/arXiv.2204.13742.

11 Neil Immerman. Nondeterministic space is closed under complementation. *SIAM Journal on Computing*, 17(5):935–938, 1988. doi:10.1137/0217058.

12 Donggyu Kim and Sang-il Oum. Vertex-minors of graphs: A survey. *Discrete Applied Mathematics*, 351:54–73, 2024. doi:10.1016/J.DAM.2024.03.011.

13 Rose McCarty. *Local structure for vertex-minors.* PhD thesis, University of Waterloo, 2021. Available at http://hdl.handle.net/10012/17633.

14 Robert McNaughton and Seymour Papert. *Counter-free automata.* The M.I.T. Press, Cambridge, Mass.-London, 1971.

15 Sang-il Oum. Rank-width and vertex-minors. *Journal of Combinatorial Theory, Series B*, 95(1):79–100, 2005. doi:10.1016/j.jctb.2005.03.003.

16 Sang-il Oum and Paul D. Seymour. Approximating clique-width and branch-width. *Journal of Combinatorial Theory, Series B*, 96(4):514–528, 2006. doi:10.1016/J.JCTB.2005.10.006.

17 Michał Pilipczuk. Graph classes through the lens of logic. *ArXiv preprint*, abs/2501.04166, 2025. doi:10.48550/arXiv.2501.04166.

18 Michał Pilipczuk, Nicole Schirrmacher, Sebastian Siebertz, Szymon Toruńczyk, and Alexandre Vigny. Algorithms and data structures for first-order logic with connectivity under vertex failures. In *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022*, volume 229 of *LIPIcs*, pages 102:1–102:18, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2022.102.

19 Omer Reingold. Undirected connectivity in log-space. *Journal of the ACM*, 55(4):17:1–17:24, 2008. doi:10.1145/1391289.1391291.

20 Nicole Schirrmacher, Sebastian Siebertz, and Alexandre Vigny. First-order logic with connectivity operators. *ACM Transactions on Computational Logic*, 24(4):1–23, 2023. doi:10.1145/3595922.

21 Marcel-Paul Schützenberger. On finite monoids having only trivial subgroups. *Information and Control*, 8:190–194, 1965. doi:10.1016/S0019-9958(65)90108-7.

22 Róbert Szelepcsényi. The method of forced enumeration for nondeterministic automata. *Acta Informatica*, 26(3):279–284, 1988. doi:10.1007/BF00299636.

23 Szymon Toruńczyk. Flip-width: Cops and Robber on dense graphs. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*, pages 663–700. IEEE, 2023. doi:10.1109/FOCS57990.2023.00045.

24 Egon Wanke. k-NLC graphs and polynomial algorithms. *Discrete Applied Mathematics*, 54(2-3):251–266, 1994. doi:10.1016/0166-218X(94)90026-4.