

Approximation Schemes for First-Order Definable Optimisation Problems

Anuj Dawar* Martin Grohe** Stephan Kreutzer** Nicole Schweikardt**

* University of Cambridge, U.K., Email: anuj.dawar@cl.cam.ac.uk

** Humboldt-Universität zu Berlin, Germany, Email: {grohe | kreutzer | schweika}@informatik.hu-berlin.de

Abstract

Let $\varphi(X)$ be a first-order formula in the language of graphs that has a free set variable X , and assume that X only occurs positively in $\varphi(X)$. Then a natural minimisation problem associated with $\varphi(X)$ is to find, in a given graph G , a vertex set S of minimum size such that G satisfies $\varphi(S)$. Similarly, if X only occurs negatively in $\varphi(X)$, then $\varphi(X)$ defines a maximisation problem. Many well-known optimisation problems are first-order definable in this sense, for example, MINIMUM DOMINATING SET or MAXIMUM INDEPENDENT SET.

We prove that for each class \mathcal{C} of graphs with excluded minors, in particular for each class of planar graphs, the restriction of a first-order definable optimisation problem to the class \mathcal{C} has a polynomial time approximation scheme.

A crucial building block of the proof of this approximability result is a version of Gaifman’s locality theorem for formulas positive in a set variable. This result may be of independent interest.

1. Introduction

It has long been known that many combinatorial optimisation problems that are hard to approximate in general have polynomial time approximation schemes (PTASs) on planar graphs, that is, they can be approximated to any given ratio in polynomial time. Among these problems are MINIMUM DOMINATING SET, MINIMUM VERTEX COVER, and MAXIMUM INDEPENDENT SET. The two main techniques for proving such approximation results on planar graphs are Lipton and Tarjan’s planar separator theorem [20] and Baker’s layerwise decomposition technique [4]. Both techniques have been generalised from planar graphs to more general graph classes such as graphs of bounded genus and ultimately to arbitrary classes of graphs with excluded minors [1, 10, 14, 7]. Recall that a *minor* of a graph G is a graph that can be obtained from a subgraph of G by contracting edges. We say that a class \mathcal{C} of graphs has an *excluded minor* if there is some graph H such that H is not a

minor of any graph in \mathcal{C} . For example, the complete graph on five vertices, K_5 , is an excluded minor of the class of all planar graphs. Most approximation results on general classes of graphs with excluded minors make heavy use of Robertson and Seymour’s structure theory for graphs with excluded minors [23]. In a recent paper, Demaine, Hajiaghayi, and Kawarabayashi [7] have proved algorithmic versions of some of the central parts of this theory and use these to obtain several new approximability results.

What kind of problems are approximable on graphs with excluded minors? Demaine et al. [7] gave a general criterion that is met by most problems known to be approximable, but is somewhat unsatisfactory because it describes when a certain proof technique works rather than describing a “natural” class of problems. On planar graphs, Khanna and Motwani [17] tried a more systematic “syntactic” approach: They defined three “generic” problems based on propositional logic and showed that the planar versions of these problems have PTASs. Then they showed that most problems which at that time were known to have PTASs can easily be reduced to one of these three problems. In this paper, we carry out a different logic based approach towards identifying a large class of problems that have PTASs on classes of graphs with excluded minors. Our approach, in contrast to that of Khanna and Motwani, is based on first-order logic: Let $\varphi(X)$ be a first-order formula in the language of graphs that has a free set variable X , and assume that X only occurs positively in $\varphi(X)$. Then a natural minimisation problem $\text{MIN}_{\varphi(X)}$ associated with $\varphi(X)$ is to find, in a given graph G , a vertex set S of minimum size such that G satisfies $\varphi(S)$. Many natural minimisation problems can be described as problems $\text{MIN}_{\varphi(X)}$ for a suitable formula $\varphi(X)$. For example, the MINIMUM DOMINATING SET problem is $\text{MIN}_{\varphi(X)}$ for the formula $\varphi(X) = \forall x(Xx \vee \exists y(Xy \wedge E_{xy}))$. The condition that $\varphi(X)$ be positive in X is imposed to guarantee monotonicity, which is necessary to exclude pathological examples (see Example 11). Similarly, if X only occurs negatively in a formula $\psi(X)$ then this formula defines a natural maximisation problem $\text{MAX}_{\psi(X)}$. For example, MAX-

MINIMUM INDEPENDENT SET is $\text{MAX}_{\psi(X)}$ for the formula $\psi(X) = \forall x \forall y (\neg Xx \vee \neg Xy \vee \neg Exy)$. Papadimitriou and Yannakakis [22] considered similar syntactically defined optimisation problems to introduce their well-known class MAXSNP; a more detailed analysis of such problems was later carried out by Kolaitis and Thakur [18, 19]. The syntactical problems studied by these authors can be directly translated to problems of the form $\text{MAX}_{\varphi(X)}$ or $\text{MIN}_{\varphi(X)}$, where φ is a formula of first-order or existential second-order logic and X a (not necessarily unary) relation variable.

We call problems of the form $\text{MIN}_{\varphi(X)}$ or $\text{MAX}_{\varphi(X)}$ *first-order definable optimisation problems*. For a class \mathcal{C} of graphs, we write $\text{MIN}_{\varphi(X)}(\mathcal{C})$ and $\text{MAX}_{\varphi(X)}(\mathcal{C})$ to denote the restriction of $\text{MIN}_{\varphi(X)}$ ($\text{MAX}_{\varphi(X)}$, respectively) to instances from \mathcal{C} . Our main result states that the restrictions of first-order definable optimisation problems to graphs with excluded minors have PTASs. More precisely:

Theorem 1. *Let $\varphi(X)$ be a first-order formula in the language of graphs that is positive in a set variable X , and let \mathcal{C} be a class of graphs with an excluded minor. Then $\text{MIN}_{\varphi(X)}(\mathcal{C})$ has a PTAS.*

Similarly, if $\varphi(X)$ is negative in X , then $\text{MAX}_{\varphi(X)}(\mathcal{C})$ has a PTAS.

In fact, we prove a stronger result; we actually get an *EPTAS* (efficient polynomial time approximation scheme, cf. [11, 21]). This means that the running time of the ε -approximation algorithm is $f(\varepsilon) \cdot n^{O(1)}$ for some function f . Let us remark that the statement of Theorem 1 does not extend, for example, to monadic second-order logic (cf. Example 11).

We see our result as an algorithmic meta theorem in the style of Courcelle’s theorem [6]. It gives a uniform and natural explanation for a large family of algorithmic results, and for many problems it gives a quick way of proving that they have a PTAS on graphs with excluded minors. Consider, for example, the MINIMUM DISTANCE d DOMINATING SET problem (for a constant $d \geq 1$), which asks for a minimum set S of vertices of a graph such that each vertex is within distance d of the set S . It is easy to find a first-order formula $\varphi(X)$ such that this problem is $\text{MIN}_{\varphi(X)}$. Thus it has a PTAS on all classes with excluded minors. This was not known before, and the problem also does not meet Demaine et al.’s general approximability criterion mentioned above. Our theorem can easily be extended from graphs to arbitrary relational structures and also to weighted versions of the problems, but we defer the details of these generalisations to the full version of this paper. Even with these generalisations, we do not claim that our theorem captures all problems that have PTASs on classes of graphs with excluded minors. Indeed, it is easy to find problems that meet Demaine et al.’s approximability criterion, but are not first-order definable in our sense. Even on planar graphs, our

approach seems incomparable with Khanna and Motwani’s in that there is no obvious translation of our syntactically defined problems into theirs or vice versa. An important difference between our result and those of Demaine et al. and Khanna and Motwani is that we obtain an EPTAS. For the problem PLANAR TMIN, for which Khanna and Motwani obtained a PTAS, it can actually be proved that, under reasonable complexity theoretic assumptions, it does not have an EPTAS [5, 21].

The proof of Theorem 1 has two parts: the second, algorithmic, part builds on techniques that were first applied in [14] to classes of graphs with excluded minors and generalise Baker’s layerwise decomposition technique [4]. However, the techniques have to be generalised considerably to handle the very general class of problems we consider here. The crucial property of first-order definable optimisation problems that our algorithms exploit is the locality of first-order logic. In the first part of the proof of Theorem 1, we prove a “positive version” of Gaifman’s locality theorem, a result which may be of independent interest:

Theorem 2. *Let $\varphi(X)$ be a first-order sentence that is positive in the set variable X . Then there is a Boolean combination $\psi(X)$ of basic local sentences so that $\psi(X)$ is positive in X and equivalent to $\varphi(X)$.*

The necessary definitions will be given later. Rather unexpectedly, the proof of this theorem proved to be fairly difficult, as we were unable to adapt the known proofs of Gaifman’s theorem [12] (see [9, 16] for alternative proofs) or of its existential version [15]. Our proof of the positive version uses ideas from [3] to analyse the spatial distribution of the types occurring in a structure, and it uses a lemma from [15] to get from a nonuniform to a uniform version of the theorem, but the core combinatorial argument is new.

2. Preliminaries

A *vocabulary* is a finite set of relation symbols and constant symbols. Associated with every relation symbol R is a positive integer called the *arity* of R . In the following, τ always denotes a vocabulary. τ is called *relational* if it does not contain any constant symbol.

A τ -structure \mathcal{A} consists of a non-empty set A , called the *universe* of \mathcal{A} , an element $c^{\mathcal{A}} \in A$ for each constant symbol $c \in \tau$, and a relation $R^{\mathcal{A}} \subseteq A^r$ for each r -ary relation symbol $R \in \tau$.

The *Gaifman graph* of a τ -structure \mathcal{A} is the (undirected, loop-free) graph $\mathcal{G}_{\mathcal{A}}$ with vertex set A and an edge between two vertices $a, b \in A$ iff there exists an $R \in \tau$ and a tuple $(a_1, \dots, a_r) \in R^{\mathcal{A}}$ such that $a, b \in \{a_1, \dots, a_r\}$.

The *distance* between two elements $a, b \in A$ in \mathcal{A} , denoted by $\text{dist}^{\mathcal{A}}(a, b)$, is defined to be the length (that is, number of edges) of the shortest path from a to b in the Gaif-

man graph of \mathcal{A} . For $r \geq 0$ and $a \in A$, the r -neighbourhood of a in \mathcal{A} is the set $N_r^{\mathcal{A}}(a) = \{b \in A : \text{dist}^{\mathcal{A}}(a,b) \leq r\}$.

The induced substructure of \mathcal{A} with universe $N_r^{\mathcal{A}}(a)$ is denoted by $\mathcal{N}_r^{\mathcal{A}}(a)$. We omit superscripts \mathcal{A} if \mathcal{A} is clear from the context.

We write $\text{FO}(\tau)$ to denote the class of all formulae in first-order logic over the vocabulary τ , and we write $\text{qr}(\varphi)$ to denote the *quantifier rank* of an $\text{FO}(\tau)$ -formula φ . If X is a unary relation symbol not in τ , then an occurrence of X in an $\text{FO}(\tau \cup \{X\})$ -formula φ is said to be *positive* if it is within the scope of an *even* number of negations and it is said to be *negative* otherwise. We say that the formula φ is *positive in X* (or *X -positive*) if all occurrences of X in φ are positive. Similarly, we say that φ is *negative in X* (or *X -negative*) if all occurrences of X in φ are negative.

For every $r \geq 0$, we let $\text{dist}_{\leq r}(x,y)$ be an $\text{FO}(\tau)$ -formula expressing that the distance between x and y is at most r . We often write $\text{dist}(x,y) \leq r$ instead of $\text{dist}_{\leq r}(x,y)$ and $\text{dist}(x,y) > r$ or $\text{dist}_{> r}(x,y)$ instead of $\neg \text{dist}_{\leq r}(x,y)$.

The r -relativisation of a formula $\varphi(x)$ is the formula $\varphi^r(x)$ obtained from φ by first renaming all bound variables so that they are different from x and then replacing each subformula of the form $\exists y \psi$ by $\exists y (\text{dist}(x,y) \leq r \wedge \psi)$ and each subformula of the form $\forall y \psi$ by $\forall y (\text{dist}(x,y) \leq r \rightarrow \psi)$. Clearly, the r -relativisation of every formula $\varphi(x)$ is r -local, that is, for every τ -structure \mathcal{A} and every $a \in A$ we have $\mathcal{A} \models \varphi^r(a) \iff \mathcal{N}_r^{\mathcal{A}}(a) \models \varphi^r(a)$. Note that we also have $\mathcal{A} \models \varphi^r(a) \iff \mathcal{N}_r^{\mathcal{A}}(a) \models \varphi(a)$.

A (symmetric) *basic local sentence* (with parameters k, r, q) is a sentence of the form

$$\exists x_1 \dots \exists x_k \left(\bigwedge_{1 \leq i < j \leq k} \text{dist}(x_i, x_j) > 2r \wedge \bigwedge_{1 \leq i \leq k} \psi(x_i) \right),$$

where $k, r \geq 1$ and $\psi(x)$ is r -local and of quantifier rank q (here, the adjective ‘‘symmetric’’ emphasises that the same formula ψ is used for each of the variables x_i).

Theorem 3 (Gaifman [12]). *Every first-order sentence over a relational vocabulary is equivalent to a Boolean combination of basic local sentences.*

3. A positive locality theorem

In this section we present a proof of the version of Gaifman’s theorem for formulae *positive* in a unary relation symbol, stated in Theorem 2. From now on, fix a relational vocabulary τ and a unary relation symbol $X \notin \tau$. For proving Theorem 2 we adopt the approach of [15] of using *asymmetric* basic local formulae in an intermediate step.

An *asymmetric basic local sentence with parameters k, κ, r, q* is a sentence of the form

$$\exists x_1 \dots \exists x_k \left(\bigwedge_{1 \leq i < j \leq k} \text{dist}(x_i, x_j) > \kappa \cdot 2r \wedge \bigwedge_{i=1}^k \psi_i(x_i) \right),$$

where $\psi_i(x_i)$ is r -local and of quantifier rank at most q . We denote the set of all asymmetric basic local sentences with parameters $k' \leq k$, $\kappa, r' \leq r$, and $q' \leq q$ by $\text{ABL}(k, \kappa, r, q)$. By $\text{ABL}^+(k, \kappa, r, q)$ (respectively, $\text{ABL}^-(k, \kappa, r, q)$) we denote the set of all sentences in $\text{ABL}(k, \kappa, r, q)$ that are *positive* (respectively, *negative*) in X .

Similarly, we write $\text{BL}(k, r, q)$, $\text{BL}^+(k, r, q)$, and $\text{BL}^-(k, r, q)$ for, respectively, the set of all, all X -positive, and all X -negative symmetric basic local sentences with parameters $k' \leq k$, $r' \leq r$, and $q' \leq q$.

For a sentence $\varphi \in \text{ABL}(k, \kappa, r, q)$ of the form

$$\exists x_1 \dots \exists x_k \left(\bigwedge_{1 \leq i < j \leq k} \text{dist}(x_i, x_j) > \kappa \cdot 2r \wedge \bigwedge_{i=1}^k \psi_i(x_i) \right)$$

we write $\varphi_{[1/\kappa]}$ to denote the sentence

$$\exists x_1 \dots \exists x_k \left(\bigwedge_{1 \leq i < j \leq k} \text{dist}(x_i, x_j) > 2r \wedge \bigwedge_{i=1}^k \psi_i(x_i) \right)$$

(in particular, $\varphi_{[1/\kappa]} \in \text{ABL}(k, 1, r, q)$).

The two major steps in proving Theorem 2 consist of showing the following two technical lemmas:

Lemma 4. *Let $K, Q, R \geq 2$ and let $\kappa := 2^{K^2-1}$. Suppose \mathcal{A} and \mathcal{B} are $\tau \cup \{X\}$ -structures such that every X -positive (resp. X -negative) basic local sentence in $\text{BL}(K, \kappa \cdot 2R, Q)$ that holds in \mathcal{A} also holds in \mathcal{B} . Then we have for every X -positive (resp. X -negative) sentence $\varphi \in \text{ABL}(K, \kappa, R, Q)$ that $\mathcal{A} \models \varphi$ implies $\mathcal{B} \models \varphi_{[1/\kappa]}$.*

We omit the proof of Lemma 4 since it is virtually identical to the proof of Lemma 4 in [15]. We will use Lemma 4 as an intermediate step in proving the following:

Lemma 5. *For every $q \geq 0$ there exist $K, R, Q \geq 2$ such that for all $\tau \cup \{X\}$ -structures \mathcal{A}, \mathcal{B} the following holds: If for every $\varphi \in \text{BL}^+(K, R, Q)$, $\mathcal{A} \models \varphi$ implies $\mathcal{B} \models \varphi$, and for every $\varphi \in \text{BL}^-(K, R, Q)$, $\mathcal{B} \models \varphi$ implies $\mathcal{A} \models \varphi$, then we have for every X -positive $\text{FO}(\tau \cup \{X\})$ -sentence ζ of quantifier rank at most q that $\mathcal{A} \models \zeta$ implies $\mathcal{B} \models \zeta$.*

Note that by using Lemma 5 one easily obtains a proof of Theorem 2 (details of this will be given in the full version of the paper).

The remainder of Section 3 is devoted to the proof of Lemma 5. To prove Lemma 5, we use the following ‘‘ X -positive’’ variant of the classical *Ehrenfeucht-Fraïssé game* (EF-game, for short) for first-order logic.

3.1. The X -positive EF-game.

The rules of this game are the same as for the ‘‘classical’’ EF-game for first-order logic (cf., e.g. [9]), the winning condition, however, is slightly different. To be precise, the ‘‘ X -positive’’ EF-game is defined as follows:

Let q be a positive integer. The q -round X -positive EF-game is played by two players, *the spoiler* and *the duplicator*, on two $\tau \cup \{X\}$ -structures \mathcal{A} and \mathcal{B} . The spoiler's intention is to show a difference between the two structures, while the duplicator tries to make them look alike. There is a fixed number q of rounds. Each round $i \in \{1, \dots, q\}$ is played as follows: First, the spoiler chooses either an element a_i in A or an element b_i in B . Next, the duplicator chooses an element in the other structure. I.e., she chooses an element b_i in B if the spoiler's move was in A , or an element a_i in A if the spoiler's move was in B . After q rounds the game ends with elements a_1, \dots, a_q chosen in A and b_1, \dots, b_q chosen in B . The *duplicator has won the game* iff the mapping f defined via $(a_j \mapsto b_j)_{j=1, \dots, q}$ is an X -positive partial isomorphism from \mathcal{A} to \mathcal{B} , i.e.,

- (i) for any tuple of elements (v_1, \dots, v_s) within the domain of f and any relation symbol $R \in \tau$ of arity s , we have $(v_1, \dots, v_s) \in R^{\mathcal{A}} \iff (f(v_1), \dots, f(v_s)) \in R^{\mathcal{B}}$, and
- (ii) for any element v within the domain of f and for the particular unary relation symbol X , we have $v \in X^{\mathcal{A}} \implies f(v) \in X^{\mathcal{B}}$.

Otherwise, the spoiler has won the game. Since the game is finite, one of the two players must have a *winning strategy*, i.e., he or she can always win the game, no matter how the other player plays. We write $\mathcal{A} \xrightarrow{X\text{-pos}} \mathcal{B}$ to denote that the duplicator has a winning strategy in the q -round X -positive EF-game on \mathcal{A} and \mathcal{B} . Note that the relation defined by $\xrightarrow{X\text{-pos}}$ on the class of all $\tau \cup \{X\}$ -structures is reflexive and transitive, but not symmetric.

The fundamental use of the q -round X -positive EF-game comes from the fact that it characterises definability by X -positive first-order sentences in the following sense:

Proposition 6. *Let \mathcal{A} and \mathcal{B} be $\tau \cup \{X\}$ -structures and let q be a positive integer. If $\mathcal{A} \xrightarrow{X\text{-pos}} \mathcal{B}$, then we have for every X -positive FO($\tau \cup \{X\}$)-sentence ζ of quantifier rank at most q , that $\mathcal{A} \models \zeta$ implies $\mathcal{B} \models \zeta$.*

The proof is straightforward.

Now, note that Lemma 5 is an immediate consequence of Proposition 6 and the following lemma.

Lemma 7. *For every $q \geq 0$ there exist $K, R, Q \geq 2$ such that for all $\tau \cup \{X\}$ -structures \mathcal{A}, \mathcal{B} : If for every $\varphi \in \text{BL}^+(K, R, Q)$, $\mathcal{A} \models \varphi$ implies $\mathcal{B} \models \varphi$, and for every $\varphi \in \text{BL}^-(K, R, Q)$, $\mathcal{B} \models \varphi$ implies $\mathcal{A} \models \varphi$, then $\mathcal{A} \xrightarrow{X\text{-pos}} \mathcal{B}$.*

Subsection 3.2 below is devoted to the proof of Lemma 7.

3.2. Proof of Lemma 7.

Before describing the duplicator's winning strategy we need some preparation.

Let $r, q \geq 0$, \mathcal{A} a $\tau \cup \{X\}$ -structure and $a \in A$. The *full* (r, q) -type of a in \mathcal{A} is the set $\text{full-}(r, q)\text{-type}^{\mathcal{A}}(a) :=$

$$\{\varphi^r(x) \mid \varphi \in \text{FO}(\tau \cup \{X\}), \text{qr}(\varphi) \leq q, \mathcal{A} \models \varphi^r(a)\}.$$

Note that there is a formula

$$\theta_{(r, q, \mathcal{A}, a)}(x) := \bigwedge_{\varphi \in \text{full-}(r, q)\text{-type}^{\mathcal{A}}(a)} \varphi(x),$$

defining an element's full (r, q) -type and for all $\tau \cup \{X\}$ -structures \mathcal{B} and all $b \in B$ we have $\mathcal{B} \models \theta_{(r, q, \mathcal{A}, a)}(b) \iff \text{full-}(r, q)\text{-type}^{\mathcal{B}}(b) = \text{full-}(r, q)\text{-type}^{\mathcal{A}}(a)$. Also, all the formulae $\theta_{(r, q, \mathcal{A}, a)}$ are r -local and of quantifier rank at most \tilde{q} (where $\tilde{q} \geq q$ only depends on r, q , and $\tau \cup \{X\}$). In the following, we often identify the types with these formulae. We denote the set of all full (r, q) -types by $\Theta_{(r, q)}$.

A type $\theta(x)$ is *realised* in a structure \mathcal{A} if there is an element $a \in A$ such that $\mathcal{A} \models \theta(a)$. The element a is called a *realisation* of θ in \mathcal{A} .

Let \mathcal{A} be a structure, let $C \subseteq A$, let $\theta \in \Theta_{(r, q)}$ be a full (r, q) -type, and let $R, K \geq 0$. We say that

θ is R -covered by C

if for all realisations a of θ in \mathcal{A} we have $a \in N_R(C)$. We say that

θ is (R, K) -free over C

if there are realisations a_1, \dots, a_K of θ in \mathcal{A} such that $a_i \notin N_R(C)$ for all $i \in \{1, \dots, K\}$ and $\text{dist}(a_i, a_j) > R$ for all $i, j \in \{1, \dots, K\}$ with $i \neq j$.

The next lemma analyses the spatial distribution of the types occurring in a structure.

Lemma 8. *For all $k, r, q \geq 0$ there are $\hat{K} \geq k$ and $\hat{R} \geq r$, such that for all $\tau \cup \{X\}$ -structures \mathcal{A} and \mathcal{B} there are K, R with $k \leq K \leq \hat{K}$ and $r \leq R \leq \hat{R}$ and sets $C^{\mathcal{A}} \subseteq A$ and $C^{\mathcal{B}} \subseteq B$ such that for $\kappa := 2^{K^2-1}$ and each $\mathcal{D} \in \{\mathcal{A}, \mathcal{B}\}$, the following properties are satisfied:*

- (1) $|C^{\mathcal{D}}| \leq K$;
- (2) $\text{dist}(c, c') > \kappa \cdot 10R$, for all $c, c' \in C^{\mathcal{D}}$ with $c \neq c'$;
- (3) each $\theta \in \Theta_{(r, q)}$ realised in \mathcal{D} is either R -covered by $C^{\mathcal{D}}$ or $(\kappa \cdot 10R, 10K)$ -free over $C^{\mathcal{D}}$.

Due to lack of space, we defer the proof of Lemma 8 to the full version of the paper. For the proof of Lemma 7 we also need the notions of *positive* types and *negative* types of an element a in a $\tau \cup \{X\}$ -structure \mathcal{A} . The *positive* (r, q) -type of a is the set

$$\text{pos-}(r, q)\text{-type}^{\mathcal{A}}(a) := \{\varphi^r(x) \mid \varphi \in \text{FO}(\tau \cup \{X\}) \text{ positive in } X, \text{qr}(\varphi) \leq q, \mathcal{A} \models \varphi^r(a)\}.$$

Similarly, the *negative* (r, q) -type of a is the set

$$\text{neg-}(r, q)\text{-type}^{\mathcal{A}}(a) := \{\varphi^r(x) : \varphi \in \text{FO}(\tau \cup \{X\}) \text{ negative in } X, \text{qr}(\varphi) \leq q, \mathcal{A} \models \varphi^r(a)\}.$$

Note that $\text{pos-}(r, q)\text{-type}^{\mathcal{A}}(a) \subseteq \text{full-}(r, q)\text{-type}^{\mathcal{A}}(a)$ and $\text{neg-}(r, q)\text{-type}^{\mathcal{A}}(a) \subseteq \text{full-}(r, q)\text{-type}^{\mathcal{A}}(a)$. The formula

$$\theta_{(r, q, \mathcal{A}, a)}^+(x) := \bigwedge_{\varphi \in \text{pos-}(r, q)\text{-type}^{\mathcal{A}}(a)} \varphi(x),$$

defines the positive (q, r) -type, in the sense that for all $(\tau \cup \{X\})$ -structures \mathcal{B} and all $b \in B$ with $\mathcal{B} \models \theta_{(r, q, \mathcal{A}, a)}^+(b)$ we have $\text{pos-}(r, q)\text{-type}^{\mathcal{B}}(b) \supseteq \text{pos-}(r, q)\text{-type}^{\mathcal{A}}(a)$. Analogously, one obtains a formula $\theta_{(r, q, \mathcal{A}, a)}^-(x)$ that defines the negative (q, r) -type of a in \mathcal{A} . Note that the formulae $\theta_{(r, q, \mathcal{A}, a)}^+$ and $\theta_{(r, q, \mathcal{A}, a)}^-$ are r -local and of quantifier rank at most \tilde{q} (where $\tilde{q} \geq q$ only depends on q, r , and $\tau \cup \{X\}$). Furthermore, $\theta_{(r, q, \mathcal{A}, a)}^+(x)$ is positive in X , whereas $\theta_{(r, q, \mathcal{A}, a)}^-(x)$ is negative in X . In the following, we often identify the types with these formulae.

We denote the set of all positive and negative (r, q) -types by $\Theta_{(r, q)}^+$ and $\Theta_{(r, q)}^-$, respectively. A (positive or negative) type $\theta(x)$ is *realised* in a structure \mathcal{A} if there is an $a \in A$ such that $\mathcal{A} \models \theta(a)$. We call a a *realisation* of θ in \mathcal{A} .

Proof of Lemma 7:

We fix $q \geq 0$ and let $k := q$, $r := 3^q$, and $Q := \widetilde{q+1}$. Let \hat{K}, \hat{R} be chosen according to Lemma 8. Now let \mathcal{A} and \mathcal{B} be $\tau \cup \{X\}$ -structures such that

- (*) for every $\varphi \in \text{BL}^+(\hat{K}, 2^{\hat{K}^2-1} \cdot 10\hat{R}, Q)$, $\mathcal{A} \models \varphi$ implies $\mathcal{B} \models \varphi$, and for every $\varphi \in \text{BL}^-(\hat{K}, 2^{\hat{K}^2-1} \cdot 10\hat{R}, Q)$, $\mathcal{B} \models \varphi$ implies $\mathcal{A} \models \varphi$.

Before we can describe the duplicator's winning strategy in the q -round X -positive EF-game on \mathcal{A} and \mathcal{B} , we first need to explore the "playing fields" \mathcal{A} and \mathcal{B} . To this end, we first apply Lemma 8 to \mathcal{A} and \mathcal{B} (with k, r, q) to obtain numbers $K \leq \hat{K}, R \leq \hat{R}$ and sets $C := C^{\mathcal{A}} \subseteq A$ and $D := C^{\mathcal{B}} \subseteq B$ such that for $\kappa := 2^{K^2-1}$ and for each $\mathcal{D} \in \{\mathcal{A}, \mathcal{B}\}$ the properties (1)–(3) of Lemma 8 are satisfied.

Note that, since $K \leq \hat{K}$ and $R \leq \hat{R}$, (*) in particular holds when replacing \hat{K} with K and \hat{R} with R . Thus, by applying Lemma 4 (both the X -positive and the X -negative version, while interchanging the roles of \mathcal{A} and \mathcal{B} when applying the X -negative version), we obtain

- (**) for every $\varphi \in \text{ABL}^+(K, \kappa, 5R, Q)$, $\mathcal{A} \models \varphi$ implies $\mathcal{B} \models \varphi_{[1/\kappa]}$ and for every $\varphi \in \text{ABL}^-(K, \kappa, 5R, Q)$, $\mathcal{B} \models \varphi$ implies $\mathcal{A} \models \varphi_{[1/\kappa]}$.

Let us proceed with fixing some more notation. We call a *full* (r, q) -type $\theta \in \Theta_{(r, q)}$ *rare* in \mathcal{A} (in \mathcal{B}) if it is R -covered by C (by D , respectively); otherwise, we call θ *frequent*.

A *positive* (r, q) -type $\pi \in \Theta_{(r, q)}^+$ is called *saturated* if there is a *full* type $\theta \in \Theta_{(r, q)}$ that is frequent in \mathcal{B} such that $\pi \subseteq \theta$. Similarly, a *negative* (r, q) -type $\nu \in \Theta_{(r, q)}^-$ is called *saturated* if there is a full type $\theta \in \Theta_{(r, q)}$ that is frequent in \mathcal{A} such that $\nu \subseteq \theta$.

We define a bipartite Graph \mathcal{G} on $C \cup D$ by drawing an edge from $c \in C$ to $d \in D$ if there are a $c' \in N_{2R}(c)$ and a $d' \in N_{2R}(d)$ such that the positive $(4R, q+1)$ -type of c' is contained in the positive $(4R, q+1)$ -type of d' and hence the negative $(4R, q+1)$ -type of d' is contained in the negative $(4R, q+1)$ -type of c' .

We call an element $c \in C$ *special* if there is an $a \in N_R(c)$ such that the positive (r, q) -type of a is *not saturated* (i.e., every *full* (r, q) -type $\theta \supseteq \text{pos-}(r, q)\text{-type}^{\mathcal{A}}(a)$ is *rare* in \mathcal{B} , i.e., realised only by elements in $N_R(D)$). Similarly, an element $d \in D$ is called *special* if there is a $b \in N_R(d)$ such that the negative (r, q) -type of b is *not saturated*.

Let $C_S \subseteq C$ and $D_S \subseteq D$ be the sets of all special vertices.

Claim 9. *The bipartite graph \mathcal{G} has a matching M such that each special element of C and D is an endpoint of an edge in M .*

Proof of Claim 9: Let $\ell = |C_S|$ and $C_S = \{c_1, \dots, c_\ell\}$. For every $i \in \{1, \dots, \ell\}$ let $\pi_i(x)$ be the positive $(4R, q+1)$ -type of c_i . Then \mathcal{A} satisfies the X -positive asymmetric basic local sentence $\varphi := \exists x_1 \dots \exists x_\ell \left(\bigwedge_{1 \leq i < j \leq \ell} d(x_i, x_j) > \kappa \cdot 10R \wedge \bigwedge_{i=1}^{\ell} \pi_i(x_i) \right)$. Thus \mathcal{B} satisfies the sentence $\varphi_{[1/\kappa]}$ (to see this, recall the assumption (**)) on \mathcal{A} and \mathcal{B} on the sentences in $\text{ABL}^+(K, \kappa, 5R, Q)$, note that $\ell \leq K$ since $C_S \subseteq C$ and $|C| \leq K$, and recall from Lemma 8 that elements in C have pairwise distance $> \kappa \cdot 10R$). Hence we can find $b_1, \dots, b_\ell \in B$ of pairwise distance greater than $10R$ such that $\mathcal{B} \models \pi_i(b_i)$ for every $i \in \{1, \dots, \ell\}$.

Let us next note that each of the b_i belongs to $N_{2R}(D)$. This can be seen as follows. Since c_i is special, there exists an $a_i \in N_R(c_i)$ whose positive (r, q) -type π' is not saturated, i.e., each full (r, q) -type that contains π' is realised in \mathcal{B} only by elements in $N_R(D)$. Since the positive $(4R, q+1)$ -type of b_i contains the positive $(4R, q+1)$ -type π_i of c_i and $\mathcal{N}_{4R}(c_i)$ satisfies the X -positive formula $\exists y (\text{dist}(c_i, y) \leq R \wedge \pi'(y))$, we know that also $\mathcal{N}_{4R}(b_i)$ satisfies this formula, and thus there exists an element b'_i with $\text{dist}(b_i, b'_i) \leq R$ whose full (r, q) -type contains π' . Since π' is not saturated, we conclude that $b'_i \in N_R(D)$ and hence $b_i \in N_{2R}(D)$.

Since each b_i belongs to $N_{2R}(D)$, there are $d_1, \dots, d_\ell \in D$ such that $\text{dist}(d_i, b_i) \leq 2R$, for every $i \in \{1, \dots, \ell\}$. The vertices d_1, \dots, d_ℓ are pairwise distinct, since $\text{dist}(b_i, b_j) > 10R$ and thus $\text{dist}(d_i, d_j) > 6R$, for all $i, j \in \{1, \dots, \ell\}$ with $i \neq j$. Furthermore, by the definition of the graph \mathcal{G} , there is an edge between c_i and d_i , for every $i \in \{1, \dots, \ell\}$ (to see

this, note that $c'_i := c_i \in N_{2R}(c_i)$ has the positive $(4R, q+1)$ -type π_i , and $d'_i := b_i \in N_{2R}(d_i)$ has a positive $(4R, q+1)$ -type that contains π_i . It follows that each set $C' \subseteq C_S$ of special vertices has at least $|C'|$ neighbours in D .

Analogously, we can show that each set $D' \subseteq D_S$ of special vertices has at least $|D'|$ neighbours in C .

Now Claim 9 immediately follows from the following purely combinatorial lemma, which may be viewed as an extension of Hall's well known marriage theorem. Let us say that a vertex is *covered* by a matching if it is an endpoint of an edge in the matching.

Lemma 10. *Let \mathcal{G} be a bipartite graph with bipartition C, D of the vertex set. Let $C_S \subseteq C$ and $D_S \subseteq D$, and suppose that each $C' \subseteq C_S$ has at least $|C'|$ neighbours in D and each $D' \subseteq D_S$ has at least $|D'|$ neighbours in C . Then there is a matching M of \mathcal{G} that covers each vertex in $C_S \cup D_S$.*

The proof of Lemma 10 can be found in the full version of this paper. To proceed with the proof of Lemma 7 let us now fix a matching M that covers all special vertices (such a matching exists by Claim 9). Let $c_1, \dots, c_m \in C$ and $d_1, \dots, d_m \in D$ be the vertices covered by M via an edge between c_i and d_i , for each $i \in \{1, \dots, m\}$. By the definition of the graph \mathcal{G} , for $i \in \{1, \dots, m\}$, there are vertices $c'_i \in N_{2R}(c_i)$ and $d'_i \in N_{2R}(d_i)$ such that the positive $(4R, q+1)$ -type of d'_i contains the positive $(4R, q+1)$ -type of c'_i . In particular, the duplicator has a winning strategy for the q -round X -positive EF-game on $\mathcal{N}_{4R}(c'_i)$ and $\mathcal{N}_{4R}(d'_i)$.

Recall that, by the definition of *special* vertices, every $a \in A$ whose positive (r, q) -type is not saturated is in the R -neighbourhood of some special vertex of C and hence, in particular, in $N_R(c_i) \subseteq N_{2R}(c'_i)$, for some $i \in \{1, \dots, m\}$ (to see this, note that (1) a has to belong to $N_R(C)$ due to Lemma 8, and (2) the vertex from C in whose R -neighbourhood a lies has to be special). Similarly, every $b \in B$ whose negative (r, q) -type is not saturated is in $N_R(d_i)$ and hence in $N_{2R}(d'_i)$ for some $i \in \{1, \dots, m\}$.

Now it is easy to define a winning strategy for the duplicator in the q -round X -positive EF-game on \mathcal{A} and \mathcal{B} : If the spoiler plays near a vertex c'_i or d'_i , the duplicator answers according to the local strategy there. If the spoiler plays near a vertex played before, the duplicator answers according to the local strategy there. Otherwise, the spoiler plays a saturated vertex far away from everything, and the duplicator can always find an answer. The meaning of "near" varies with the number j of moves remaining in the game. The duplicator seeks to preserve neighbourhoods of radius 3^j around previously played elements or $2R + 3^j$ around c'_i or d'_i .

This finally completes the proof of Lemma 7 and thus, altogether, the proof of Theorem 2. \square

4. Graph Decompositions

In this section we fix some notation and briefly present the basic notions from graph minor theory used later on. See the last chapter of [8] or the survey [24].

The vertex set of a graph \mathcal{G} is denoted by $V^{\mathcal{G}}$ and its edge set is denoted by $E^{\mathcal{G}}$. For $U \subseteq V^{\mathcal{G}}$ we write $\langle U \rangle$ for the subgraph of \mathcal{G} induced by U . A tree is an acyclic, connected graph. We usually use rooted directed trees where edges are directed from the root towards the leaves.

A *minor* of a graph \mathcal{G} is a graph \mathcal{H} that can be obtained from a subgraph of \mathcal{G} by contracting edges. We write $\mathcal{H} \preceq \mathcal{G}$ to denote that \mathcal{H} is a minor of \mathcal{G} . A class \mathcal{C} of graphs is *minor closed* if, and only if, for all $\mathcal{G} \in \mathcal{C}$ and $\mathcal{H} \preceq \mathcal{G}$ also $\mathcal{H} \in \mathcal{C}$. A class \mathcal{C} of graphs is *\mathcal{H} -free* if $\mathcal{H} \not\preceq \mathcal{G}$ for all $\mathcal{G} \in \mathcal{C}$. We then call \mathcal{H} an *excluded minor* of \mathcal{C} .

A *tree-decomposition* of a graph \mathcal{G} is a pair $(T, (B_t)_{t \in V^T})$, where T is a directed tree and $B_t \subseteq V^{\mathcal{G}}$ for all $t \in V^T$ such that $\bigcup_{t \in V^T} \langle B_t \rangle = \mathcal{G}$ and for every $v \in V^{\mathcal{G}}$ the set $\{t : v \in B_t\}$ is connected. The sets B_t are called *blocks* of the decomposition. The *width* of $(T, (B_t)_{t \in V^T})$ is $\max\{|B_t| : t \in V^T\} - 1$ and the *tree-width* $\text{tw}(\mathcal{G})$ of a graph \mathcal{G} is the minimal width of any of its tree-decompositions. A class \mathcal{C} of graphs has *bounded tree-width*, if there is a constant k bounding the tree-width of all members of \mathcal{C} .

For a tree-decomposition $(T, (B_t)_{t \in V^T})$ and $t \in V^T$ with parent $s \in V^T$ we let $A_t := B_t \cap B_s$. For the root r of T we let $A_r := \emptyset$. The *adhesion* of $(T, (B_t)_{t \in V^T})$ is the number $\text{ad}(T, (B_t)_{t \in V^T}) := \max\{|A_t| : t \in V^T\}$. The *torso* $[B_t]$ of $(T, (B_t)_{t \in V^T})$ at $t \in V^T$ is the graph with vertex set B_t and with an edge between $u, v \in B_t$ if $(u, v) \in E^{\mathcal{G}}$ or u, v both belong to a block B_s with $s \neq t$.

A tree-decomposition of a graph \mathcal{G} over a class \mathcal{B} of graphs is a tree-decomposition $(T, (B_t)_{t \in V^T})$ whose torsos $[B_t]$ are contained in \mathcal{B} .

We also need the following notion. The *local tree-width* of a graph \mathcal{G} is the function $\text{ltw}^{\mathcal{G}} : \mathbb{N} \rightarrow \mathbb{N}$ defined as $\text{ltw}^{\mathcal{G}}(r) := \max\{\text{tw}(\langle N_r^{\mathcal{G}}(v) \rangle) : v \in V^{\mathcal{G}}\}$. A class \mathcal{C} of graphs has *bounded local tree-width* if there is an $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $\text{ltw}^{\mathcal{G}}(r) \leq f(r)$ for all $\mathcal{G} \in \mathcal{C}$ and $r \in \mathbb{N}$.

5. First-order definable optimisation problems

In this section we present a proof of Theorem 1. Here, we only prove the minimisation version of the theorem. The maximisation version is proved similarly using techniques from [14]. We defer the details to the full version of the paper. We begin with a formal definition. Let X_{\min} be an optimal solution for $\text{MIN}_{\varphi(X)}$ on input \mathcal{G} . For $\varepsilon > 0$ we call a solution X , i.e. a set X with $(\mathcal{G}, X) \models \varphi$, ε -close if $|X| \leq (1 + \varepsilon)|X_{\min}|$. A *polynomial-time approximation scheme* (PTAS) for $\text{MIN}_{\varphi(X)}$ is a uniform family $(A_{\varepsilon})_{\varepsilon > 0}$ of algorithms, where A_{ε} , given an instance \mathcal{G} , computes an ε -close solution for \mathcal{G} in polynomial time. Uniform here

means that there is an algorithm that, given ε , generates A_ε . A PTAS is called *efficient*, (or, it is an EPTAS), if the degree of the polynomial bounding the running time of A_ε does not depend on ε . Our proof of Theorem 1 establishes an EPTAS for first-order definable optimisation problems.

Example 11. It is well known that the class of planar graphs excludes a minor. Thus, by Theorem 1, every optimisation problem definable by an X -positive or X -negative first-order formula has a PTAS on the class of planar graphs.

However, the result neither extends to monadic second-order logic (MSO) nor to first-order formulae which are not monotone in X . For this, note that 3-colourability is NP-complete even on the class of planar graphs (see [13]). As 3-colourability can easily be defined by a formula $\psi \in \text{MSO}$, the minimisation problem defined by $\varphi(X) := \psi \rightarrow \forall x Xx$ cannot have a PTAS (unless $P = \text{NP}$). Similarly, a simple reduction shows that the 3-colourability problem on planar graphs can be reduced to a minimisation problem on planar graphs defined by a non-monotone first-order formula. \square

To prove Theorem 1 we first need some preparation. Let φ be a first-order formula positive in X . By Theorem 2 we can assume that $\varphi := \bigvee_i \bigwedge_j \psi_{i,j}$, where each $\psi_{i,j}$ is X -positive and either basic local or the negation of a basic local formula. To compute a minimal set X satisfying the formula, it suffices to consider the disjuncts $\bigwedge_j \psi_{i,j}$ separately. Thus we can assume that $\varphi := \bigwedge_j \psi_j$, where each ψ_j is either an existential basic local formula of the form $\exists x_1 \dots \exists x_r \bigwedge_{s \neq t} \text{dist}(a_s, a_t) > 2r \wedge \bigwedge_{s=1}^r \vartheta_j(x_s)$ or a *universal basic local formula* of the form $\forall x_1 \dots \forall x_r (\bigwedge_{s \neq t} \text{dist}(a_s, a_t) > 2r \rightarrow \bigvee_{s=1}^r \vartheta_j(x_s))$. Here, the formulae ϑ_j are X -positive and restricted to the r -neighbourhoods of x_l . We will transform φ even further. For this, observe that for any fixed graph $\mathcal{G} := (V, E)$, the formula $\exists x_1 \dots \exists x_k \bigwedge_{s \neq t} \text{dist}(a_s, a_t) > 2r \wedge \bigwedge_{s=1}^k \vartheta(x_s)$ is equivalent to $\bigvee_{(a_1, \dots, a_k) \in (V^\mathcal{G})^k} (\bigwedge_{s \neq t} \text{dist}(a_s, a_t) > 2r \wedge \bigwedge_{s=1}^k \vartheta(a_s))$ which can be simplified even further to

$$\bigvee \left\{ \bigwedge_{s=1}^k \vartheta(a_s) : (a_1, \dots, a_k) \in (V^\mathcal{G})^k \text{ and } \bigwedge_{s \neq t} \text{dist}(a_s, a_t) > 2r \right\}. \quad (*)$$

Here, the elements a_i are used as constants in the formulae. A similar, although more complicated, translation can be given for the universal basic local sentences. Let \mathcal{G} be a graph and $\psi(X) := \forall x_1 \dots \forall x_k (\bigwedge_{s \neq t} \text{dist}(x_s, x_t) > 2r \rightarrow \bigvee_{s=1}^k \vartheta(x_s))$ be a universal basic local sentence. For every $0 \leq j \leq k-1$ let π_j be the set of functions $f : \{1, \dots, j\} \rightarrow \mathbb{N}$ such that $\sum_{i=1}^j f(i) < k$. Let $\psi^*(X)$ be the formula

$$\bigvee_{j < k} \left(\bigvee_{(a_1, \dots, a_j) \in (V^\mathcal{G})^j} \left(\bigvee_{f \in \pi_j} \bigvee_{n=0}^{k-j} \chi_{j, \bar{a}, f, n} \right) \right) \quad (**)$$

where $\chi_{j, \bar{a}, f, n} :=$

$$\begin{aligned} & \bigwedge_{1 \leq l < l' \leq j} \text{dist}(a_l, a_{l'}) > 5^n \cdot 8r \wedge \\ & \forall x \left(\bigwedge_{i=1}^j \text{dist}(x, a_i) \geq 5^n \cdot 4r \rightarrow \vartheta(x) \right) \wedge \\ & \bigwedge_{s=1}^j \forall x_1 \dots \forall x_{f(s)+1} \left(\bigwedge_{l=1}^{f(s)+1} \text{dist}(x_l, a_s) < 5^n 4r \wedge \right. \\ & \quad \left. \bigwedge_{1 \leq l < l' \leq f(s)+1} \text{dist}(x_l, x_{l'}) > 2r \right) \rightarrow \bigvee_{l=1}^{f(s)+1} \vartheta(x_l). \end{aligned}$$

We claim that ψ and ψ^* are equivalent on \mathcal{G} in the following sense. Due to lack of space, we defer the proof to the full version of the paper.

Lemma 12. *For every set $X \subseteq V^\mathcal{G}$, $(\mathcal{G}, X) \models \psi$ if, and only if, $(\mathcal{G}, X) \models \psi^*$.*

Lemma 12 implies that for a given graph \mathcal{G} we can translate the formula $\varphi(X)$ into a conjunction of formulae of the form $(*)$ and $(**)$. By distributivity, we can translate this into a disjunction of conjunctions of formulae $\chi_{j, \bar{a}, f, n}$ and formulae $\xi := \bigwedge_{s=1}^k \vartheta(b_l)$ for tuples of constants \bar{a}_i, \bar{b}_i , functions f and numbers j and n . As the arity of the tuples \bar{a}, \bar{b} is bounded by a function of φ and the ranges of the variables j, n , and f also only depend on φ , the translation can be done in polynomial time in the size of \mathcal{G} .

Further, the first line of the formula $\chi_{j, \bar{a}, f, n}$ only imposes conditions on the choice of the tuple \bar{a} . It follows that for computing an approximation of a set X satisfying φ in \mathcal{G} it suffices to compute an approximation of a set X satisfying the conjunction of formulae

- $\bigwedge_{s=1}^k \vartheta(a_l)$ for an r -local formula ϑ and a tuple of constants (a_1, \dots, a_k) with $\text{dist}(a_s, a_t) > 2r$ for all $i \neq j$,
- $\forall x (\bigwedge_{i=1}^j \text{dist}(x, a_i) \geq q \rightarrow \vartheta(x))$ for an r -local formula ϑ , some $q > 2r$ and a tuple of elements (a_1, \dots, a_j) of distance $\text{dist}(a_s, a_t) > 2q$ for all $i \neq j$, and
- $\forall x_1 \dots \forall x_{f(s)+1} \left(\bigwedge_{l=1}^{f(s)+1} \text{dist}(x_l, a_s) < 5^n 4r \wedge \bigwedge_{1 \leq l < l' \leq f(s)+1} \text{dist}(x_l, x_{l'}) > 2r \right) \rightarrow \bigvee_{l=1}^{f(s)+1} \vartheta(x_l)$ which are q -local around a_s for some $q > 2r$ only depending on φ .

Note that the formulas in $c)$ are $5^n \cdot 4r$ -local around a_s . Hence, Theorem 1 follows from the following lemma.

Lemma 13. *Let \mathcal{C}' be a class of graphs with an excluded minor and let $\sigma := \{a_1, \dots, a_k\}$ be a set of constant symbols. Let $q > 0$ and let $\varphi(X) \in \text{FO}$ be an X -positive conjunction of q -local formulae $\varphi_i(a)$ using only one constant symbol $a \in \sigma$ and formulae $\psi_t := \forall x (\bigwedge_{s=1}^{k_t} \text{dist}(x, a_{s,t}) > q_t \rightarrow \vartheta(x))$, for an r -local formula ϑ with $q_t > 2r$, using constant symbols $\bar{a}_t \subseteq \sigma$.*

Let $\mathcal{C} := \{(\mathcal{G}, \bar{a}) : \mathcal{G} \in \mathcal{C}' \text{ and } \bar{a} \subseteq V^\mathcal{G}\}$. Then $\text{MIN}_{\varphi(X)}(\mathcal{C})$ has a polynomial time approximation scheme.

To prove the lemma we use a decomposition theorem for classes of graphs with an excluded minor that is due to [14]. We first introduce some notation.

For $\lambda, \mu \geq 0$ we let

$$\begin{aligned} \mathcal{L}(\lambda) &:= \left\{ \mathcal{G} : \begin{array}{l} \text{for all } \mathcal{H} \preceq \mathcal{G}, \text{ for all } r \geq 0 \\ \text{ltw}^{\mathcal{H}}(r) \leq \lambda \cdot r \end{array} \right\} \\ \mathcal{L}(\lambda, \mu) &:= \left\{ \mathcal{G} : \begin{array}{l} \text{there is } X \subseteq V^{\mathcal{G}} \\ \text{s.th. } (|X| \leq \mu \wedge \mathcal{G} \setminus X \in \mathcal{L}(\lambda)) \end{array} \right\} \end{aligned}$$

Note that $\mathcal{L}(\lambda, \mu)$ is minor closed. The proof of Lemma 13 is based on the following decomposition theorem for classes of graphs with an excluded minor.

Theorem 14 ([14]). *Let \mathcal{C} be a class of graphs with an excluded minor. Then there exist $\lambda, \mu \in \mathbb{N}$ such that all $\mathcal{G} \in \mathcal{C}$ have a tree-decomposition over $\mathcal{L}(\lambda, \mu)$.*

For the approximation algorithms we want to show, we need an algorithmic version of this theorem.

Theorem 15 ([14]). *Let \mathcal{C} be a minor closed class of graphs. Then there is a polynomial-time algorithm that computes for a given graph \mathcal{G} a tree-decomposition of \mathcal{G} over \mathcal{C} or rejects \mathcal{G} if no such decomposition exists.*

We apply this result to the minor closed classes $\mathcal{L}(\lambda, \mu)$. Let \mathcal{G} be a graph. For every vertex $v \in V^{\mathcal{G}}$ and integers $j \geq i \geq 0$ we define

$$L_v^{\mathcal{G}}[i, j] := \{w \in V^{\mathcal{G}} : i \leq \text{dist}^{\mathcal{G}}(v, w) \leq j\},$$

where $\text{dist}^{\mathcal{G}}(v, w)$ denotes the distance between v and w in \mathcal{G} . To simplify notation, we will use $L_v^{\mathcal{G}}[i, j]$ for arbitrary integers i, j and set $L_v^{\mathcal{G}}[i, j] := \emptyset$ if $i > j$ and $L_v^{\mathcal{G}}[i, j] := L_v^{\mathcal{G}}[0, j]$ for $i \leq 0$. The following lemma follows easily.

Lemma 16. *Let $\lambda \in \mathbb{N}$. Then $\text{tw}(\langle L_v^{\mathcal{G}}[i, j] \rangle) \leq \lambda \cdot (j - i + 1)$ for all $\mathcal{G} \in \mathcal{L}(\lambda)$, $v \in V^{\mathcal{G}}$ and $i, j \in \mathbb{Z}$ with $i \leq j$.*

Now, let φ be as in the hypothesis of Lemma 13 and let \mathcal{C} be a class of graphs with an excluded minor. By Theorem 14, we can choose $\lambda, \mu \in \mathbb{N}$ such that every graph $\mathcal{G} \in \mathcal{C}$ has a tree-decomposition over $\mathcal{L}(\lambda, \mu)$. Let $\varepsilon > 0$. We describe a polynomial time algorithm that, on input $\mathcal{G} \in \mathcal{C}$ and $\bar{a} \in V^{\mathcal{G}}$, computes an ε -close solution for $\text{MIN}_{\varphi(X)}(\mathcal{C})$ on (\mathcal{G}, \bar{a}) . To ease notation we will consider the tuple \bar{a} as part of the graph and use notation such as $\mathcal{G} \models \varphi$ for $(\mathcal{G}, \bar{a}) \models \varphi$.

The proof of Lemma 13 is split into two steps. In the first step, which we present in the next subsection, we prove the lemma for the classes $\mathcal{L}(\lambda)$ and $\mathcal{L}(\lambda, \mu)$ of graphs. Here, we use the corresponding result for graphs of bounded tree-width which essentially follows from [2].

Theorem 17 ([2]). *Let $\varphi(X)$ be an X -positive formula of MSO. Then $\text{MIN}_{\varphi(X)}(\mathcal{C})$ can be solved in linear time on any class \mathcal{C} of graphs of bounded tree-width.*

In Section 5.3, we extend the proof to graphs which have a tree-decomposition over $\mathcal{L}(\lambda, \mu)$, i.e. to all graphs in \mathcal{C} .

5.1. The levels of graphs of bounded local tree-width.

In the first step of the proof of Lemma 13 we show that the restriction of $\text{MIN}_{\varphi(X)}(\mathcal{C})$ to instances in $\mathcal{L}(\lambda)$ has a PTAS. Let $\varphi := \bigwedge_{i \in I_e} \varphi_i(a_i) \wedge \bigwedge_{t \in I_u} \psi_t$, where the $\varphi_i(a_i)$ are r -local formulae and the $\psi_t := \forall x (\bigwedge_{s=1}^{k_t} \text{dist}(x, a_{s,t}) > r_t \rightarrow \vartheta_t(x))$, for a q -local formula ϑ_t with $r_t > 2q$, using constant symbols $\bar{a}_t \subseteq \sigma$. For simplicity we assume w.l.o.g. that $r_t = r_{t'}$ for all t, t' . Let $r := r_t$.

Let $k := \lceil \frac{2r}{\varepsilon} \rceil$. Note that $\frac{k+2r}{k} \leq (1 + \varepsilon)$. Let $\mathcal{G} \in \mathcal{L}(\lambda)$. Choose a node $v \in V^{\mathcal{G}}$ arbitrarily. For $1 \leq i \leq k$ and $j \geq 0$ let $L_{ij} := L_v^{\mathcal{G}}[(j-1) \cdot k - r + i, j \cdot k + r + i]$. By Lemma 16, $\text{tw}(\langle L_{ij} \rangle) \leq \lambda(k + 2r + 1)$.

For all $1 \leq i \leq k$, $j \geq 0$ let X_{ij} be a set of minimal cardinality such that

- (1) $(\langle L_{ij} \rangle, X_{ij}) \models \varphi_l(a_l)$ for all $l \in I_e$ such that a_l and its r -neighbourhood is contained in L_{ij} and
- (2) $(\langle L_{ij} \rangle, X_{ij}) \models \psi_t$ for all $i \in I_u$.

Note that as $\psi_t := \forall x (\bigwedge_{s=1}^{k_t} \text{dist}(x, a_{s,t}) > r_t \rightarrow \vartheta_t(x))$ also mentions constants interpreted by vertices outside of L_{ij} , this is, strictly speaking, not well defined. However, as the a_t are constants, we can easily check whether x is close to any constant interpreted by an element outside of L_{ij} . For instance, we could colour the r -neighbourhoods of a_1, \dots, a_{k_t} and then check in ψ_t that x is outside a coloured area. For ease of presentation we will therefore simply write $(\langle L_{ij} \rangle, X_{ij}) \models \psi_t$ even in cases where some or all of the constants are interpreted by elements outside of L_{ij} .

By Theorem 17 the sets X_{ij} can be computed in linear time. For $1 \leq i \leq k$ let $X_i := \bigcup_{j \geq 0} X_{ij}$. As φ is monotone in X , every X_i is a solution of $\text{MIN}_{\varphi(X)}(\mathcal{C})$ on \mathcal{G} .

Let X_{\min} be an optimal solution of $\text{MIN}_{\varphi(X)}(\mathcal{C})$ for \mathcal{G} , i.e. a set of minimal cardinality such that $(\mathcal{G}, X_{\min}) \models \varphi$. Clearly, $X_{\min} \cap L_{ij}$ satisfies the conditions (1) and (2) above for all levels L_{ij} . Hence,

$$\sum_{i=1}^k |X_i| \leq \sum_{i=1}^k \sum_{j \geq 0} |X_{ij}| \leq \sum_{i=1}^k \sum_{j \geq 0} |L_{ij} \cap X_{\min}| \leq (k + 2r) |X_{\min}|.$$

The last inequality follows as every node $v \in V^{\mathcal{G}}$ can be contained in at most $k + 2r$ levels L_{ij} . Choose m , $1 \leq m \leq k$ such that $|X_m| := \min\{|X_1|, \dots, |X_k|\}$. Then

$$|X_m| \leq \frac{(k + 2r)}{k} |X_{\min}| \leq (1 + \varepsilon) |X_{\min}|.$$

Hence, X_m is an ε -close solution of $\text{MIN}_{\varphi(X)}(\mathcal{C})$ on \mathcal{G} . As every set X_i can be computed in polynomial time, the set X_m can also be computed in polynomial time.

5.2. Extension to the classes $L(\lambda, \mu)$.

In a second step we show how this approximation algorithm can be extended to the classes $\mathcal{L}(\lambda, \mu)$ for constants $\lambda, \mu \geq 0$. Let $\mathcal{G} \in \mathcal{L}(\lambda, \mu)$ and let $U \subseteq V^{\mathcal{G}}$ be such that $|U| \leq \mu$ and $\mathcal{G} \setminus U \in L(\lambda)$. The following extension of Theorem 17 can easily be proved by dynamic programming.

Theorem 18. *For every $k \geq 0$ and every first-order formula $\varphi(X)$ which is positive in the set-variable X , the following problem can be solved in linear time. Given a graph \mathcal{G} , a set $U \subseteq V^{\mathcal{G}}$ so that $\text{tw}(\mathcal{G} \setminus U) \leq k$ and a subset $Y \subseteq U$, find a set $X \subseteq V^{\mathcal{G}} \setminus U$ of minimal cardinality such that $(\mathcal{G}, X \cup Y) \models \varphi$ or determine that no such set exists.*

Let again $\varphi(X)$ be an X -positive first-order formula. For every $Y \subseteq U$ let $X(Y)$ be a subset of $V^{\mathcal{G}} \setminus U$ such that $(\mathcal{G}, X(Y) \cup Y) \models \varphi$ and

$$|X(Y)| \leq (1 + \varepsilon) \min\{|X| : X \subseteq V^{\mathcal{G}} \setminus U \text{ and } (\mathcal{G}, X \cup Y) \models \varphi\}$$

or $X(Y) := \perp$ if no such set exists. If $X(Y) = \perp$ for all $Y \subseteq U$ then $\text{MIN}_{\varphi(X)}(\mathcal{C})$ has no solution on \mathcal{G} and we are done. Otherwise let $Y_0 \subseteq U$ be such that $|X(Y_0) \cup Y_0|$ is minimal among $\{|X(Y) \cup Y| : Y \subseteq U \text{ and } X(Y) \neq \perp\}$. Then clearly, $X(Y_0) \cup Y_0$ is an ε -close solution for $\text{MIN}_{\varphi(X)}(\mathcal{C})$.

Using Theorem 18 instead of Theorem 17, the sets $X(Y)$ can be computed in polynomial time analogously to the first step. As there are only 2^μ possible subsets of U – recall that μ is a constant only depending on the class \mathcal{C} – and for each $Y \subseteq U$, $X(Y)$ can be computed in polynomial time, the solution $X(Y_0) \cup Y_0$ can be computed in polynomial time.

5.3. Excluded Minors.

In the last step, we show how the approximation algorithm can be extended to graphs that have a tree-decomposition over $\mathcal{L}(\lambda, \mu)$, i.e. to all graphs in \mathcal{C} .

Let $\mathcal{G} \in \mathcal{C}$. We first compute a tree-decomposition $(T, (B_t)_{t \in V^T})$ over $\mathcal{L}(\lambda, \mu)$. By Theorem 15, this can be done in polynomial time. Let r be the root of T and for every $t \in V^T$ with parent s let $A_t := B_t \cap B_s$. We set $A_r := \emptyset$. Further, for every node $t \in V^T$ let T_t be the subtree of T rooted at t and let $\mathcal{B}_t := \bigcup_{s \in T_t} B_s$.

In what follows we will construct for subgraphs \mathcal{B} of \mathcal{G} sets X such that $(\langle \mathcal{B} \rangle, X)$ satisfies

- (1) $(\langle \mathcal{B} \rangle, X) \models \varphi_l(a_l)$ for all $l \in I_e$ such that a_l and its r -neighbourhood is contained in \mathcal{B} and
- (2) $(\langle \mathcal{B} \rangle, X) \models \psi_t$ for all $t \in I_u$.

(Here, we use the same convention as in Section 5.1 above.) To simplify the presentation we write $(\langle \mathcal{B} \rangle, X) \models \varphi$ to indicate that X satisfies the conditions (1) and (2) in \mathcal{B} . The

notation is motivated by the fact that for $\mathcal{B} = \mathcal{G}$, $(\mathcal{G}, X) \models \varphi$ for any set satisfying condition (1) and (2) and vice versa.

Inductively, from the leaves to the root, we compute for every node $t \in V^T$ and for every subset $Y \subseteq A_t$ an $X(t, Y)$ such that $X(t, Y) \subseteq \mathcal{B}_t \setminus A_t$, $(\langle \mathcal{B}_t \rangle, X(t, Y) \cup Y) \models \varphi$ and

$$|X(t, Y)| \leq (1 + \varepsilon) \min\{|X| : (\langle \mathcal{B}_t \rangle, X \cup Y) \models \varphi, X \subseteq \mathcal{B}_t \setminus A_t\}$$

or $X(Y) := \perp$ if no such set exists. As tree-decompositions over $\mathcal{L}(\lambda, \mu)$ have adhesion at most $\lambda + \mu + 1$, we have $|A_t| \leq \lambda + \mu + 1$. Hence, we only have to compute a constant number of sets $X(t, Y)$ for each t . Further, for the root r we have $A_r = \emptyset$ and $\langle \mathcal{B}_r \rangle = \mathcal{G}$. Hence, $X(r, \emptyset)$ is an ε -close solution for $\text{MIN}_{\varphi(X)}(\mathcal{C})$ or \perp if no solution exists.

We show next how to compute the sets $X(t, Y)$. Suppose $t \in V^T$ and for every child t' of t we have already computed the family $X(t', \cdot)$. Let $U \subseteq B_t$ such that $|U| \leq \mu$ and $[B_t] \setminus U \in \mathcal{L}(\lambda)$. (Recall that $[B_t]$ denotes the torso of $(T, (B_t)_{t \in V^T})$ at t .) Let $W := U \cup A_t$.

For every $Z \subseteq W$, let $X_{\min}(Z)$ be a set of minimal cardinality such that $X_{\min}(Z) \subseteq \mathcal{B}_t \setminus W$ and $(\langle \mathcal{B}_t \rangle, X_{\min}(Z) \cup Z) \models \varphi$ or $X_{\min}(Z) := \perp$ if no such set exists.

Claim 19. *For every set $Z \subseteq W$ we can compute in polynomial time an $X(Z)$ such that $X(Z) \subseteq \mathcal{B}_t \setminus W$, $(\langle \mathcal{B}_t \rangle, X(Z) \cup Z) \models \varphi$, and $|X(Z)| \leq (1 + \varepsilon)|X_{\min}(Z)|$ or $X(Z) := \perp$ if no such set exists.*

Before we prove Claim 19 let us show how the proof of Lemma 13 can be completed using the claim. For every $Y \subseteq A_t$ choose a $Z \subseteq W$ such that $Z \cap A_t = Y$ and $|X(Z) \cup (Z \setminus Y)| = \min\{|X(Z') \cup (Z' \setminus Y)| : Y \subseteq Z' \subseteq W \text{ and } Z' \cap A_t = Y\}$. Set $X(t, Y) := X(Z) \cup (Z \setminus Y)$. By our choice of Z it follows that $|X(t, Y)| \leq (1 + \varepsilon) \min\{|X| : (\langle \mathcal{B}_t \rangle, X \cup Y) \models \varphi\}$ which concludes the proof.

So all that remains is to prove Claim 19. Fix a $Z \subseteq W$. We show how to compute $X(Z)$ in polynomial time. If $W = B_t$, i.e. $B_t := U \cup A_t$, then let $X(Z) := \bigcup_{(t, t') \in E^T} X(t', A_{t'} \cap Z)$.

Otherwise choose an arbitrary $v \in B_t \setminus W$. For $1 \leq i \leq k$ and $j \geq 0$ let $L_{ij} := L_v^{[B_t] \setminus W}[(j-1) \cdot k + i - r, j \cdot k + i + r]$. Then $\text{tw}(\langle L_{ij} \rangle) \leq \lambda(k+1+2r)$. For every child t' of t and every $1 \leq i \leq k$ there is at least one $j \geq 0$ such that $A_{t'} \setminus W \subseteq L_{ij}$. This follows from that fact that $A_{t'}$ induces a clique in $[B_t]$. Let $j_{\min}(i, t')$ be the least such j and let

$$L_{ij}^* := L_{ij} \cup \bigcup \{ \mathcal{B}_{t'} \setminus A_{t'} : (t, t') \in E^T \text{ and } j_{\min}(i, t') = j \}.$$

Similarly, for every $X \subseteq L_{ij}$ let

$$X^* := X \cup \bigcup \{ X(t', (X \cup Z) \cap A_{t'}) : (t, t') \in E^T, j_{\min}(i, t') = j \}.$$

We compute an $X_{ij} \subseteq L_{ij}$ with minimal $|X_{ij}^*|$ such that $(\langle L_{ij} \rangle, X_{ij} \cup Z) \models \varphi$ or set $X_{ij} := \perp$ if no such X exists. This can be done in polynomial time using the standard dynamic programming techniques on graphs of bounded tree-width,

provided that the numbers $|X(t', Y)|$ for the children t' of t are given. It is important here that every $A_{t'} \setminus W$ is a clique in $\langle L_{ij} \rangle$, as this implies that it is contained in a single block of every tree-decomposition of $\langle L_{ij} \rangle$.

Let $X_i := \bigcup_{j \geq 0} X_{ij}$ and $X_i^* := \bigcup_{j \geq 0} X_{ij}^*$. Then, by monotonicity of φ in X , $(\langle \mathcal{B}_t \rangle, X_i^* \cup Z) \models \varphi$ or $X_i = \perp$ if no set satisfying φ in $\langle \mathcal{B}_t \rangle$ exists. Finally, choose an $i \in \{1, \dots, k\}$ such that $|X_i^*| := \min\{|X_1^*|, \dots, |X_k^*|\}$ and let $X(Z) := X_i^*$. It follows that $X(Z)$ can be computed in polynomial time.

Recall that we defined $X_{\min} := X_{\min}(Z) \subseteq \mathcal{B}_t \setminus W$ to be a set of minimal order such that $(\langle \mathcal{B}_t \rangle, X_{\min} \cup Z) \models \varphi$ or $X_{\min} := \perp$ if no such set exists. It remains to prove that $|X(Z)| \leq (1 + \varepsilon)|X_{\min}|$.

By hypothesis of the algorithm we have for every child t' of t , $|X(t', (X_{\min} \cup Z) \cap A_{t'})| \leq (1 + \varepsilon)|X_{\min} \cap \mathcal{B}_{t'} \setminus A_{t'}|$. Further, the construction of X_{ij} and X_{ij}^* guarantees that for $1 \leq i \leq k$ and $j \geq 0$

$$|X_{ij}^*| \leq |X_{\min} \cap L_{ij}| + \sum_{\substack{(t,t') \in ET \\ j_{\min}(t,t')=j}} |X(t', (X_{\min} \cup Z) \cap A_{t'})|.$$

But then

$$\begin{aligned} k|X(Z)| &\leq \sum_{i=1}^k |X_i^*| \leq \sum_{i=1}^k \sum_{j \geq 0} |X_{ij}^*| \\ &\leq \sum_{i=1}^k \sum_{j \geq 0} (|X_{\min} \cap L_{ij}| + \\ &\quad \sum_{\substack{(t,t') \in ET \\ j_{\min}(t,t')=j}} |X(t', (X_{\min} \cup Z) \cap A_{t'})|) \\ &\leq \sum_{i=1}^k \sum_{j \geq 0} (|X_{\min} \cap L_{ij}| + \\ &\quad \sum_{\substack{(t,t') \in ET \\ j_{\min}(t,t')=j}} (1 + \varepsilon)|X_{\min} \cap \mathcal{B}_{t'} \setminus A_{t'}|) \\ &\leq (k + 2r)|X_{\min} \cap \mathcal{B}_t| + k(1 + \varepsilon)|X_{\min} \cap \mathcal{B}_t \setminus \mathcal{B}_t| \end{aligned}$$

This implies $|X(Z)| \leq (1 + \varepsilon)|X_{\min}|$ and concludes the proof of Lemma 13 and with it also the proof of Theorem 1.

References

- [1] N. Alon, P. Seymour, and R. Thomas. A separator theorem for graphs with an excluded minor and its applications. In *Proc. of STOC'90*, pages 293–299, 1990.
- [2] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *Journal of Algorithms*, 12(2):308–340, 1991.
- [3] A. Atserias, A. Dawar, and M. Grohe. Preservation under extensions on well-behaved finite structures. In *Proc. of ICALP'05*, volume 3580 of *LNCS*, pages 1437–1450. Springer-Verlag, 2005.
- [4] B. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41:153–180, 1994.
- [5] L. Cai, M. Fellows, D. Juedes, and F. Rosamond. The complexity of polynomial-time approximation. *Theory of Computing Systems*. To appear.
- [6] B. Courcelle. Graph rewriting: An algebraic and logic approach. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*, pages 193–242. Elsevier, 1990.
- [7] E. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *Proc. of FOCS'05*, pages 637–646, 2005.
- [8] R. Diestel. *Graph Theory*. Springer-Verlag, 2nd edition, 2000.
- [9] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer-Verlag, 2nd edition, 1999.
- [10] D. Eppstein. Subgraph isomorphism in planar graphs and related problems. *Journal of Graph Algorithms and Applications*, 3:1–27, 1999.
- [11] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer Verlag, 2006.
- [12] H. Gaifman. On local and non-local properties. In J. Stern, editor, *Proceedings of the Herbrand Symposium, Logic Colloquium '81*, pages 105–135. North Holland, 1982.
- [13] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1:237 – 267, 1976.
- [14] M. Grohe. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003.
- [15] M. Grohe and S. Wöhrle. An existential locality theorem. *Annals of Pure and Applied Logic*, 129:131–148, 2004.
- [16] H. J. Keisler and W. B. Lotfallah. Shrinking games and local formulas. *Annals of Pure and Applied Logic*, 128:215–225, 2004.
- [17] S. Khanna and R. Motwani. Towards a syntactic characterization of PTAS. In *Proc. of STOC'96*, pages 329–337, 1996.
- [18] P. Kolaitis and M. Thakur. Logical definability of NP optimization problems. *Information and Computation*, 115(2):321–353, 1994.
- [19] P. Kolaitis and M. Thakur. Approximation properties of NP minimization classes. *Journal of Computer and System Sciences*, 50:391–411, 1995.
- [20] R. Lipton and R. Tarjan. A separator theorem for planar graphs. *SIAM J. on Applied Mathematics*, 36:177–189, 1979.
- [21] D. Marx. Parameterized complexity and approximation algorithms, 2006.
- [22] C. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43:425–440, 1991.
- [23] N. Robertson and P. Seymour. Graph minors I–XX. App. in *Journal of Combinatorial Theory, Series B* since 1982.
- [24] R. Thomas. Recent excluded minor theorems. In *Surveys in Combinatorics*, LMS Lecture Note Series. Cambridge University Press, 1999.

APPENDIX

A. Proofs omitted in Section 3

Proof of Theorem 2:

Let ζ be an X -positive FO($\tau \cup \{X\}$)-sentence, let q be the quantifier rank of ζ , and let $K, R, Q \geq 2$ be chosen according to Lemma 5. It should be clear that, up to logical equivalence, each of the sets $\text{BL}^+(K, R, Q)$ and $\text{BL}^-(K, R, Q)$ only contains a finite number of distinct sentences. Let $M > 0$ be an upper bound on this number and let $\varphi_1^+, \dots, \varphi_M^+$ (respectively, $\varphi_1^-, \dots, \varphi_M^-$) be a list of all (up to logical equivalence) sentences in $\text{BL}^+(K, R, Q)$ (respectively, $\text{BL}^-(K, R, Q)$). For sets $I^+, I^- \subseteq \{1, \dots, M\}$ we let

$$\theta_{I^+, I^-} := \bigwedge_{i \in I^+} \varphi_i^+ \wedge \bigwedge_{j \in I^-} \neg \varphi_j^-.$$

For every $\tau \cup \{X\}$ -structure \mathcal{A} we let

$$I_{\mathcal{A}}^+ := \{i \in \{1, \dots, M\} : \mathcal{A} \models \varphi_i^+\} \quad \text{and} \\ I_{\mathcal{A}}^- := \{j \in \{1, \dots, M\} : \mathcal{A} \models \neg \varphi_j^-\}.$$

Obviously, $\theta_{I_{\mathcal{A}}^+, I_{\mathcal{A}}^-}$ is an X -positive sentence such that $\mathcal{A} \models \theta_{I_{\mathcal{A}}^+, I_{\mathcal{A}}^-}$.

We let ζ' be the disjunction of the sentences θ_{I^+, I^-} , for all those $I^+, I^- \subseteq \{1, \dots, M\}$, for which there exists a $\tau \cup \{X\}$ -structure \mathcal{A} with $\mathcal{A} \models \zeta$, $I_{\mathcal{A}}^+ = I^+$, and $I_{\mathcal{A}}^- = I^-$. Obviously, ζ' is an X -positive Boolean combination of (symmetric) basic local sentences. To complete the proof of Theorem 2, it remains to show that ζ' is equivalent to ζ . To this end, let \mathcal{B} be an arbitrary $\tau \cup \{X\}$ -structure. Our aim is to show that $\mathcal{B} \models \zeta \iff \mathcal{B} \models \zeta'$.

For the direction " \implies ", note that if $\mathcal{B} \models \zeta$, then $\theta_{I_{\mathcal{B}}^+, I_{\mathcal{B}}^-}$ is one of the disjuncts in ζ' , and thus $\mathcal{B} \models \zeta'$.

For the direction " \impliedby ", note that if $\mathcal{B} \models \zeta'$, then \mathcal{B} satisfies (at least) one of the disjuncts of ζ' . I.e., there is a $\tau \cup \{X\}$ -structure \mathcal{A} with $\mathcal{A} \models \zeta$ such that $\mathcal{B} \models \theta_{I_{\mathcal{A}}^+, I_{\mathcal{A}}^-}$. Since $\mathcal{B} \models \theta_{I_{\mathcal{A}}^+, I_{\mathcal{A}}^-}$, we know that every $\varphi \in \text{BL}^+(K, R, Q)$ that holds in \mathcal{A} also holds in \mathcal{B} and every $\varphi \in \text{BL}^-(K, R, Q)$ that holds in \mathcal{B} also holds in \mathcal{A} . Therefore, by applying Lemma 5, we obtain that $\mathcal{A} \models \zeta$ implies $\mathcal{B} \models \zeta$. This completes the proof of Theorem 2. \square

Proof of Proposition 6:

We prove the contrapositive. I.e., assuming that there is an X -positive FO($\tau \cup \{X\}$)-sentence ζ of quantifier rank at most q such that $\mathcal{A} \models \zeta$ and $\mathcal{B} \models \neg \zeta$, we show that the *spoiler* has a winning strategy in the q -round X -positive EF-game on \mathcal{A} and \mathcal{B} . We assume w.l.o.g. that ζ is in negation normal form, i.e., negation symbols only occur directly in front of atomic subformulae of ζ . By induction on $i \in \{0, \dots, q\}$ we show that the spoiler can play in such a way that after the i th round of the game there is

an X -positive formula ξ_i of quantifier rank at most $q-i$ such that $(\mathcal{A}, a_1, \dots, a_i) \models \xi_i$ and $(\mathcal{B}, b_1, \dots, b_i) \models \neg \xi_i$. Note that, in particular, for $i = q$ this implies that the mapping $(a_j \mapsto b_j)_{j=1, \dots, q}$ cannot be an X -positive partial isomorphism from \mathcal{A} to \mathcal{B} , and thus the spoiler will have won the game.

The basis $i = 0$ is trivially satisfied with $\xi_i := \zeta$. For the induction step we know that there is an X -positive formula ξ_i of quantifier rank at most $q-i$ such that $(\mathcal{A}, a_1, \dots, a_i) \models \xi_i$ and $(\mathcal{B}, b_1, \dots, b_i) \models \neg \xi_i$. If ξ_i is of the form $\exists x \chi$ (respectively, $\forall x \chi$), then letting $\xi_{i+1} := \chi$, the spoiler can choose an element a_{i+1} in \mathcal{A} (respectively, b_{i+1} in \mathcal{B}) such that $(\mathcal{A}, a_1, \dots, a_{i+1}) \models \xi_{i+1}$ (respectively, $(\mathcal{B}, b_1, \dots, b_{i+1}) \models \neg \xi_{i+1}$), whereas for any answer b_{i+1} (respectively, a_{i+1}) that the duplicator may give, we have $(\mathcal{B}, b_1, \dots, b_{i+1}) \models \neg \xi_{i+1}$ (respectively, $(\mathcal{A}, a_1, \dots, a_{i+1}) \models \xi_{i+1}$).

In the remaining cases, ξ_i does *not* start with a quantifier, but it is straightforward to find a subformula ξ'_i of ξ_i that (1) starts with a (universal or existential) quantifier, (2) has the same free variables as ξ_i , and (3) satisfies $(\mathcal{A}, a_1, \dots, a_i) \models \xi'_i$ and $(\mathcal{B}, b_1, \dots, b_i) \models \neg \xi'_i$. Now, letting $\xi'_i := \exists x \chi$ or $\forall x \chi$, one can proceed as in the previous case. \square

Proof of Lemma 8:

Let $k, r, q \geq 0$ and let \mathcal{A}, \mathcal{B} be $\tau \cup \{X\}$ -structures. For some $m \geq 0$ only depending on $r, q, \mathcal{A}, \mathcal{B}$, and for $0 \leq i \leq m$, we inductively construct numbers $R_i, K_i, \kappa_i := 2^{K_i^2 - 1}$, sets $C_i^{\mathcal{A}} \subseteq A, C_i^{\mathcal{B}} \subseteq B$, and sets $\Gamma_i^{\mathcal{A}}, \Gamma_i^{\mathcal{B}} \subseteq \Theta_{(r, q)}$ of (r, q) -types such that:

- (i) $R_m, K_m, \kappa_m, C_m^{\mathcal{A}}, C_m^{\mathcal{B}}$ satisfy (1)–(3);
- (ii) $R_i, K_i, \kappa_i, C_i^{\mathcal{A}}, C_i^{\mathcal{B}}$ satisfy (1) and (2), for all $0 \leq i \leq m$;
- (iii) $N_{R_{i-1}}(C_{i-1}^{\mathcal{D}}) \subseteq N_{R_i}(C_i^{\mathcal{D}})$, for each $\mathcal{D} \in \{\mathcal{A}, \mathcal{B}\}$ and for all $1 \leq i \leq m$;
- (iv) each $\theta \in \Gamma_i^{\mathcal{D}}$ is R_i -covered by $C_i^{\mathcal{D}}$, for each $\mathcal{D} \in \{\mathcal{A}, \mathcal{B}\}$ and for all $0 \leq i \leq m$;
- (v) for all $1 \leq i \leq m$, the following is true: for each $\mathcal{D} \in \{\mathcal{A}, \mathcal{B}\}$, $\Gamma_{i-1}^{\mathcal{D}} \subseteq \Gamma_i^{\mathcal{D}}$, and for at least one $\mathcal{D} \in \{\mathcal{A}, \mathcal{B}\}$, the inclusion is *strict*.

By (v), m is bounded by $2 \cdot |\Theta_{(r, q)}|$. Thus the construction terminates in a number of steps that only depends on r and q (which implies that we obtain upper bounds \hat{K}, \hat{R} on K_m, R_m that only depend on r, q , but not on the particular structures \mathcal{A}, \mathcal{B}).

As the induction basis, we let $R_0 = r, K_0 = k, \kappa_0 := 2^{K_0^2 - 1}, C_0^{\mathcal{A}} = C_0^{\mathcal{B}} = \emptyset$, and $\Gamma_0^{\mathcal{A}} = \Gamma_0^{\mathcal{B}} = \emptyset$. Obviously, (ii)–(v) are satisfied.

For the induction step from i to $i+1$, suppose that $R_i, K_i, \kappa_i := 2^{K_i^2 - 1}, C_i^{\mathcal{A}}, C_i^{\mathcal{B}}, \Gamma_i^{\mathcal{A}}, \Gamma_i^{\mathcal{B}}$ satisfy (ii)–(iv). If $R_i, K_i, \kappa_i, C_i^{\mathcal{A}}, C_i^{\mathcal{B}}$ satisfy (1)–(3), then we let $m = i$ and we

are done. Assume otherwise. Then, there is $\mathcal{D} \in \{\mathcal{A}, \mathcal{B}\}$ and $\theta \in \Theta_{(r,q)}$ such that θ is neither R_i -covered by $C_i^{\mathcal{D}}$ nor $(\kappa_i \cdot 10R_i, 10K_i)$ -free over $C_i^{\mathcal{D}}$. Let us assume w.l.o.g. that $\mathcal{D} = \mathcal{A}$ (the case where $\mathcal{D} = \mathcal{B}$ is symmetric). Then, we inductively define elements $a_1, \dots, a_\ell \in A$ as follows: If θ is $\kappa_i \cdot 10R_i$ -covered by $C_i^{\mathcal{A}}$, then we let $\ell := 0$ (i.e., $\{a_1, \dots, a_\ell\} = \emptyset$). Otherwise, there must be an element $a_1 \notin N_{\kappa_i \cdot 10R_i}(C_i^{\mathcal{A}})$ such that a_1 realises type θ . Now suppose that a_1, \dots, a_j are already defined. If θ is $\kappa_i \cdot 10R_i$ -covered by $C_i^{\mathcal{A}} \cup \{a_1, \dots, a_j\}$, then we let $\ell = j$ and stop. Otherwise, we let $a_{j+1} \notin N_{\kappa_i \cdot 10R_i}(C_i^{\mathcal{A}} \cup \{a_1, \dots, a_j\})$ such that a_{j+1} realises type θ . The construction stops in at most $(10K_i - 1)$ steps, because otherwise a_1, \dots, a_{10K_i} would witness that θ is $(\kappa_i \cdot 10R_i, 10K_i)$ -free over $C_i^{\mathcal{A}}$. Thus $\ell \leq 10K_i - 1$. Let $\tilde{C}^{\mathcal{A}} = C_i^{\mathcal{A}} \cup \{a_1, \dots, a_\ell\}$ and note that θ is $\kappa_i \cdot 10R_i$ -covered by $\tilde{C}^{\mathcal{A}}$. We let $\tilde{C}^{\mathcal{B}} := C_i^{\mathcal{B}}$. Furthermore, we let

$$\begin{aligned} K_{i+1} &:= 11K_i - 1 \quad (\text{thus, } |\tilde{C}^{\mathcal{A}}|, |\tilde{C}^{\mathcal{B}}| \leq K_{i+1}), \\ \kappa_{i+1} &:= 2^{K_{i+1}^2 - 1}, \\ \Gamma_{i+1}^{\mathcal{A}} &:= \Gamma_i^{\mathcal{A}} \cup \{\theta\}, \\ \Gamma_{i+1}^{\mathcal{B}} &:= \Gamma_i^{\mathcal{B}}. \end{aligned}$$

We would like to let $R_{i+1} = \kappa_i \cdot 10R_i$ and $C_{i+1}^{\mathcal{A}} = \tilde{C}^{\mathcal{A}}$, $C_{i+1}^{\mathcal{B}} = \tilde{C}^{\mathcal{B}}$, but that does not work because (2) is not necessarily satisfied. Instead, we inductively construct a sequence of radii $p_0 \leq \dots \leq p_s$ and sets $D_0^{\mathcal{A}} \supseteq \dots \supseteq D_s^{\mathcal{A}}$ and $D_0^{\mathcal{B}} \supseteq \dots \supseteq D_s^{\mathcal{B}}$, and then let $R_{i+1} = p_s$, $C_{i+1}^{\mathcal{A}} = D_s^{\mathcal{A}}$, and $C_{i+1}^{\mathcal{B}} = D_s^{\mathcal{B}}$. We start with $p_0 := \kappa_i \cdot 10R_i$, $D_0^{\mathcal{A}} := \tilde{C}^{\mathcal{A}}$, and $D_0^{\mathcal{B}} := \tilde{C}^{\mathcal{B}}$. Now suppose that p_j and $D_j^{\mathcal{A}}, D_j^{\mathcal{B}}$ are already defined. If $\text{dist}(c, c') > \kappa_{i+1} \cdot 10p_j$ for all c, c' with $c \neq c'$ and $c, c' \in D_j^{\mathcal{A}} \cup D_j^{\mathcal{B}}$, we let $s = j$ and stop. Otherwise, there is $\mathcal{D}' \in \{\mathcal{A}, \mathcal{B}\}$ and $c, c' \in D_j^{\mathcal{D}'}$ with $c \neq c'$ yet $\text{dist}(c, c') \leq \kappa_{i+1} \cdot 10p_j$. We let $p_{j+1} = \kappa_{i+1} \cdot 11p_j$, $D_{j+1}^{\mathcal{D}'} = D_j^{\mathcal{D}'} \setminus \{c'\}$, and if $\mathcal{D}' = \mathcal{A}$ then $D_{j+1}^{\mathcal{B}} := D_j^{\mathcal{B}}$, respectively, if $\mathcal{D}' = \mathcal{B}$ then $D_{j+1}^{\mathcal{A}} := D_j^{\mathcal{A}}$.

This completes the construction. Note that in each step one element is removed from $D_j^{\mathcal{A}} \cup D_j^{\mathcal{B}}$, and thus the construction ends after at most $s \leq |\tilde{C}^{\mathcal{A}}| + |\tilde{C}^{\mathcal{B}}| \leq 2 \cdot K_{i+1}$ steps.

We let

$$\begin{aligned} R_{i+1} &:= p_s, \\ C_{i+1}^{\mathcal{A}} &:= D_s^{\mathcal{A}}, \\ C_{i+1}^{\mathcal{B}} &:= D_s^{\mathcal{B}}. \end{aligned}$$

We now prove that these numbers and sets satisfy (ii)–(v).

For (ii), note that for each $\mathcal{D}' \in \{\mathcal{A}, \mathcal{B}\}$,

- $|C_{i+1}^{\mathcal{D}'}| \leq K_{i+1}$, since $C_{i+1}^{\mathcal{D}'} = D_s^{\mathcal{D}'} \subseteq \tilde{C}^{\mathcal{D}'}$ and $|\tilde{C}^{\mathcal{D}'}| \leq K_{i+1}$;
- $\text{dist}(c, c') > \kappa_{i+1} \cdot 10R_{i+1}$ for all $c, c' \in C_{i+1}^{\mathcal{D}'}$ with $c \neq c'$, by the construction of $D_s^{\mathcal{D}'}$.

For (iii) and (iv), note that for $0 \leq j < s$ and each $\mathcal{D}' \in \{\mathcal{A}, \mathcal{B}\}$ we have $N_{p_j}(D_j^{\mathcal{D}'}) \subseteq N_{p_{j+1}}(D_{j+1}^{\mathcal{D}'})$. In particular,

$$\begin{aligned} N_{R_i}(C_i^{\mathcal{D}'}) &\subseteq N_{\kappa_i \cdot 10R_i}(\tilde{C}^{\mathcal{D}'}) = N_{p_0}(D_0^{\mathcal{D}'}) \\ &\subseteq N_{p_s}(D_s^{\mathcal{D}'}) = N_{R_{i+1}}(C_{i+1}^{\mathcal{D}'}). \end{aligned}$$

This directly proves (iii), and it implies (iv) as follows: Each type $\theta' \in \Gamma_i^{\mathcal{D}'}$ is R_i -covered by $C_i^{\mathcal{D}'}$ and hence R_{i+1} -covered by $C_{i+1}^{\mathcal{D}'}$. The type θ is $\kappa_i \cdot 10R_i$ -covered by $\tilde{C}^{\mathcal{A}}$ and hence R_{i+1} -covered by $C_{i+1}^{\mathcal{A}}$. Finally, (v) is trivially satisfied. Thus, the proof of Lemma 8 is complete. \square

Proof of Lemma 10:

Let $N(X)$ denote the set of neighbours of a set X of vertices of \mathcal{G} . By Hall's theorem applied to the induced subgraph of \mathcal{G} with vertex set $C_S \cup N(C_S)$, there is a matching M that covers each vertex of C_S . Choose a matching M of \mathcal{G} that covers all vertices in C_S such that the number k of vertices of D_S that are not covered by M is minimum. If $k = 0$, then all vertices of $C_S \cup D_S$ are covered by M , and the lemma is proved. Assume for contradiction that $k > 0$. Let $d_0 \in D_S$ be a vertex that is not covered by M . By Hall's theorem applied to the induced subgraph of \mathcal{G} with vertex set $D_S \cup N(D_S)$, there is a matching N that covers each vertex of D_S . Let us choose such a matching N . There is a unique maximal path P starting in d_0 with edges alternating between N and M . This path P must end in a vertex not in $C_S \cup D_S$. Let $M' := (M \setminus P) \cup (N \cap P)$ be the edge set obtained from M by removing all M -edges of P and adding all N -edges. Then M' is also a matching that covers all vertices of C_S . Moreover, M' covers all vertices of D_S covered by M and in addition d_0 . This is a contradiction to the minimality of M and thus concludes the proof of Lemma 10. \square

B. Proofs omitted in Section 5

Proof of Lemma 12:

Suppose first that $(\mathcal{G}, X) \models \psi$. Let $B := \{b_1, \dots, b_l\} \subseteq V^{\mathcal{G}}$ be a set of maximal cardinality such that $(\mathcal{G}, X) \models \neg \vartheta(b_i)$ and $\text{dist}(b_i, b_j) > 2r$ for all $i \neq j$. Note that, by maximality of B , every $c \in V^{\mathcal{G}}$ with $(\mathcal{G}, X) \models \neg \vartheta(c)$ belongs to $N_{2r}(B)$. As $(\mathcal{G}, X) \models \psi$ we have $l < k$.

We construct a set $A \subseteq B$ and a radius $p \geq 8r$ such that

- (a) for all $b \in B$ there is an $a \in A$ with $\text{dist}(a, b) < \frac{p}{4}$ and
- (b) $\text{dist}(a, a') > p$ for all $a \neq a' \in A$.

For this, we inductively define a sequence of sets $A_n \subseteq B$ and radii p_n such that $p_n < p_{n+1}$ and $A_n \supseteq A_{n+1}$ for all n and every set A_n satisfies condition (a) above.

Set $p_0 := 8r$ and $A_0 := B$. For the induction step suppose p_n and A_n are already defined. If $\text{dist}(a, a') > p_n$ for all $a \neq a' \in A$ we set $p := p_n$, $A := A_n$ and terminate the induction. Otherwise there must be two elements $a, a' \in A$ with $a \neq a'$ and $\text{dist}(a, a') \leq p_n$. Let $p_{n+1} := 5p_n$ and $A_{n+1} := A_n \setminus \{a'\}$. By induction hypothesis, for every $b \in B$ there is an $a_b \in A_n$ with $\text{dist}(a_b, b) < \frac{p_n}{4}$. If $a_b \neq a'$ then condition (a) remains satisfied for b in A_{n+1} . Otherwise, $\text{dist}(a', b) < \frac{p_n}{4}$. Hence, $\text{dist}(a, b) \leq \text{dist}(a, a') + \text{dist}(a', b) < p_n + \frac{p_n}{4} \leq \frac{5}{4}p_n = \frac{p_{n+1}}{4}$. So A_{n+1} satisfies condition (a).

Note that for all i , $p_i = 5^i \cdot 8r$. Clearly, this process must stop after at most $n \leq l < k$ induction steps resulting in a radius $p := p_n$ and a set $A := A_n \subseteq B$ satisfying both conditions (a) and (b).

Define $j := |A|$ and let $\bar{a} := (a_1, \dots, a_j)$ list the elements of $A := A_n$. Further, let $P_i := \{b \in B : \text{dist}(b, a_i) < \frac{p}{4}\}$. Condition (a) implies that $\bigcup_i P_i = B$. Let $f \in \pi_j$ be defined as $f(i) := |P_i|$. We claim that $(\mathcal{G}, X) \models \chi_{j, \bar{a}, f, n}$.

Towards a contradiction, suppose $(\mathcal{G}, X) \not\models \chi_{j, \bar{a}, f, n}$. By construction, $\text{dist}(a, a') > p = p_n = 5^n \cdot 8r$ for all $a, a' \in A$. Hence, the conjunction in the first line of the definition of $\chi_{j, \bar{a}, f, n}$ is satisfied. Further, $(\mathcal{G}, X) \models \vartheta(c)$ for every element $c \in V^{\mathcal{G}}$ such that $\text{dist}(c, a) \geq 5^n \cdot 4r = \frac{p_n}{2}$. For, as every element $b \in B$ has distance less than $\frac{p_n}{4}$ from some $a \in A$, $\text{dist}(c, b) > \frac{p_n}{2} - \frac{p_n}{4} = 5^n \cdot 4r - 5^n \cdot 2r \geq 2r$. Hence, if there was such an element c with $(\mathcal{G}, X) \models \neg \vartheta(c)$ this would contradict the maximality assumption for B . So the conjunct in the second line of the definition of $\chi_{j, \bar{a}, f, n}$ is satisfied as well. Thus, as $(\mathcal{G}, X) \not\models \chi_{j, \bar{a}, f, n}$ there must be an element $a_s \in A$ and elements $c_1, \dots, c_{f(s)+1}$ such that $\text{dist}(c_i, c_{i'}) > 2r$ for all $i \neq i'$ and for all i , $(\mathcal{G}, X) \not\models \vartheta(c_i)$ and $\text{dist}(c_i, a_s) < \frac{p_n}{2}$. As $\text{dist}(b, a_t) < \frac{p_n}{4}$ for all $b \in P_t$, $t \neq s$ it follows that for all c_i and all $b \in P_t$, $t \neq s$, $\text{dist}(c_i, b) > 2r$.

But then $B' := \{c_1, \dots, c_{f(s)+1}\} \cup \bigcup_{t \neq s} P_t$ would contradict the maximality assumption of B as, by definition of f , $f(s) + 1 > |P_s|$ and therefore $|B'| > |B|$ but $(\mathcal{G}, X) \not\models \vartheta(b')$ for all $b' \in B'$ and also $\text{dist}(b, b') > 2r$ for all $b \neq b' \in B'$. This concludes the first direction of the proof.

Now suppose $(\mathcal{G}, X) \models \psi^*$, e.g. $(\mathcal{G}, X) \models \chi_{j, \bar{a}, f, n}$ for some j, \bar{a}, f, n , but $(\mathcal{G}, X) \not\models \psi$. Then there must be a set $B := \{b_1, \dots, b_k\} \subseteq V^{\mathcal{G}}$ such that $(\mathcal{G}, X) \models \neg \vartheta(b_i)$ and $\text{dist}(b_i, b_{i'}) > 2r$ for all $i \neq i'$. By definition of $\chi_{j, \bar{a}, f, n}$ all $b \in B$ must be within the $5^n \cdot 4r$ -neighbourhood of some a_i . But, again by definition of $\chi_{j, \bar{a}, f, n}$, in the $5^n \cdot 4r$ -neighbourhood around any a_i there can only be at most $f(i)$ elements which are pairwise more than $2r$ apart and all satisfy $\neg \vartheta$. As $\sum_i f(i) < k$ we get a contradiction. \square