Algorithms for Classes of Graphs with Bounded Expansion

Zdeněk Dvořák* and Daniel Král'

Institute for Theoretical Computer Science (ITI),
Faculty of Mathematics and Physics, Charles University,
Malostranské náměstí 25, 118 00 Prague 1, Czech Republic
{rakdver,kral}@kam.mff.cuni.cz

Abstract. We overview algorithmic results for classes of sparse graphs emphasizing new developments in this area. We focus on recently introduced classes of graphs with bounded expansion and nowhere-dense graphs and relate algorithmic meta-theorems for these classes of graphs to their analogues for proper minor-closed classes of graphs, classes of graphs with bounded tree-width, locally bounded tree-width and locally excluding a minor.

1 Introduction

It is well-known that many hard problems are tractable for classes of graphs with restricted structure. A classical example of this phenomenon is the result of Courcelle [5] that every graph property that can be described by a monadic second order logic formula can solved in linear time for graphs with bounded tree-width. In particular, some NP-hard problems including graph coloring or vertex domination can be solved in linear time for graphs with bounded tree-width.

In this paper, we focus on algorithmic meta-theorems for classes of graphs whose structure is limited in some sense. To motivate the results we want to present, let us switch from the algorithmic to the structural point of view and look at the chromatic number. Graphs with bounded tree-width are degenerate and thus their chromatic number is bounded. Similarly, the chromatic number of planar graphs and more generally graphs that can be embedded on a fixed surface is bounded. Graphs with bounded tree-width, planar graphs and graphs that can be embedded on a fixed surface form minor-closed classes of graphs. A general experience says that most structural (and algorithmic) properties that hold both for classes of graphs with bounded tree-width and for classes of graphs embedded on a fixed surface are also true for classes of graphs excluding a fixed minor. The chromatic number being bounded is an example.

* Institute for Theoretical computer science is supported as project 1M0545 by Czech Ministry of Education. This work was also supported by the grant GACR 201/09/0197.

C. Paul and M. Habib (Eds.): WG 2009, LNCS 5911, pp. 17–32, 2010.

However, the chromatic number is not bounded only for classes of graphs excluding a fixed minor. Other classes of graphs with bounded chromatic number include graphs with bounded maximum degree or d-degenerate graphs (for fixed integer d). Since any graph is a minor of a cubic graph, these classes of graphs are clearly not minor-closed. A more tricky example of such class is the class of graphs obtained from planar graphs by adding at most two (not necessarily non-crossing) edges to each face. Still it turns out that some algorithmic properties of planar graphs also hold for the above mentioned graph classes.

Based on these examples, one would maybe guess that the only requirement we need is that the number of edges of a graph is bounded by the function linear in the number of its vertices, i.e., its average degree is bounded. This is however not sufficient since the average degree can be decreased by adding a sparse part to the graph (a set of isolated vertices being the simplest example, but one can easily think of more sophisticated ways which also preserve connectivity or other parameters). Similarly, the maximum average degree is not fine enough since subdividing each edge of an input graph decreases maximum average degree below four but most of the structural properties of an input graph are preserved. So, one needs a more robust structural parameter to capture the common properties of the above graph classes that are essential for the algorithmic results we are interested in.

A framework of classes of graphs with bounded expansion and a more general framework of classes of nowhere-dense graphs that have been introduced in a series of papers by Nešetřil and Ossona de Mendéz [19,20,21,22,23,24,25] seems to be the right one to be considered in this setting. In this paper, we will survey known structural and algorithmic results, including recent results of the authors and Thomas on decidability of first order logic properties, for classes of graphs with bounded expansion and classes of nowhere-dense graphs and relate these results to the earlier results for other graph classes. We will also provide proofs of some easier facts and those that are essential for algorithmic applications.

2 Definitions

In this section, we present definitions and notions important for our exposition. Though some of the notions we present are fairly standard, we decided to include them for the sake of completeness.

2.1 Graph Decompositions, Graph Minors

The graph minor project of Robertson and Seymour is one of the basic stones of modern graph theory. In this subsection, we recall some definitions and results from this area which we need in our further exposition.

A *tree-decomposition* of a graph G is a tree T whose vertices correspond to subsets of vertices of G, referred to as *bags*, and the following three properties hold:

1. every vertex of G is in at least one of the bags,

2. for every edge of G, there is a bag containing both its end-vertices, and
3. if a vertex v of G is contained in the bags associated with vertices u and u' of T, then v is contained in all the bags associated with the vertices on the path between u and u' in T.

The *order* of a tree-decomposition T is the maximum size of a bag associated to a vertex of T decreased by one. The *tree-width* $\mathrm{tw}(G)$ of a graph G is the minimum order of a tree-decomposition of G. Graphs with tree-width zero are edge-less and those with tree-width at most one are forests.

More restricted width parameter is the tree-depth. The *tree-depth* $\mathrm{td}(G)$ of a graph G is the minimum depth of a rooted tree T with the same vertex set as G that for every edge vv' of G, v is an ancestor of v' or v' is an ancestor of v. To fix our terminology, the *depth* of a rooted tree T is the maximum number of vertices in a path from the root to a vertex of T, e.g., the depth of the one-vertex rooted tree is one. Vertices on the path from a vertex v to the root are *ancestors* of v. Those vertices v' such that v is an ancestor of v' are *descendents* of v.

It is not hard to see that the tree-width of a graph G is bounded by its tree-depth decreased by one (consider the optimum tree T from the definition of the tree-depth, form bags as sets of vertices on the paths from the root to the leaves and associate them with vertices of a path in the order in which the leaves of T are visited during the depth-first search). On the other hand, the tree-depth of a graph is not bounded by any function of its tree-width (the tree-depth of the n-vertex path is $\lceil \log_2(n+1) \rceil$). In fact, the tree-depth of a graph G is proportional to the length ℓ of the longest path in G since $\lceil \log_2(\ell+2) \rceil \leq \mathrm{td}(G) \leq \binom{\ell+3}{2} - 1$. It also holds [3] that $\mathrm{td}(G) \leq (\mathrm{tw}(G)+1)\log_2(n+1)$ where n is the number of vertices of G.

An alternative definition of the tree-depth can be given by means of a vertex-coloring [26]. The *ranking number* of a graph G, as defined in [2], is the minimum number k of colors $1, \ldots, k$ needed to color the vertices of G such that any path joining two vertices of the same color contains a vertex with a bigger color. It can be shown that the tree-depth of a graph G is equal to its ranking number (to obtain the coloring, color the vertices of the tree T from the definition of the tree-depth based on their distance from the root, giving the root the largest color; to obtain a decomposition, proceed conversely).

A *minor* of a graph G is a graph obtained by deleting vertices and edges and contracting edges. Recall that the operation of contracting an edge e consists of removing e, identifying its end-vertices and deleting any loops and parallel edges that arise. A class $\mathcal{G}$ of graphs is *minor-closed* if every minor of a graph from $\mathcal{G}$ is also contained in $\mathcal{G}$. Examples of minor-closed classes of graphs include graphs embeddable in a fixed surface, graphs with tree-width at most k for an integer k, graphs with tree-depth at most k and many others. Proper minor-closed classes of graphs are *degenerate*, i.e., for every proper minor-closed class $\mathcal{G}$, there exists an integer k such that every graph $G \in \mathcal{G}$ is k-degenerate which means that G and each of its subgraphs has a vertex of degree at most k.

One of the main results in the graph minor series of Robertson and Seymour [30] asserts that every minor-closed class $\mathcal{G}$ of graphs has a finite list of obstructions, i.e., there exist $G_1, \ldots, G_k$ such that $G \in \mathcal{G}$ if and only if G does not contain any of the graphs $G_1, \ldots, G_k$ as a minor (these graphs are also called *obstructions*). E.g., the tree-width of a graph G is at most one if and only if G does not contain K_3 as a minor, and at most two if and only if G does not contain K_4 as a minor. The complete list containing four obstructions for graphs with tree-width at most three was given in [1]. A minor-closed class of graphs can contain graphs with arbitrary big tree-width (planar graphs being an example), but it is known that the tree-width of graphs in a minor-closed class $\mathcal{G}$ of graphs is bounded if and only if one of the obstructions for $\mathcal{G}$ is planar [29].

2.2 Local Parameters

First order logic graph properties are of localized nature as we discuss in Subsection 4.2. Because of this, graphs with locally restricted structure are important from the algorithmic point of view: classes of graphs with locally bounded tree-width were introduced by Eppstein [12] (using somewhat different notation) and classes of graphs locally excluding a minor were defined by Dawar, Grohe and Kreutzer [6].

Before we define these graph classes, we need to recall several definitions. If G is a graph and v is a vertex of G, then $N_d(v)$ is the *d-neighborhood* of v, i.e., the set of vertices of G at distance at most d from v. If A is a set of vertices of a graph G, then $G[A]$ is the subgraph of G *induced* by A, i.e., the subgraph with vertex set A that contains all edges of G with both end-vertices from A.

We say that a class $\mathcal{G}$ of graphs has *locally bounded tree-width* if there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that the tree-width of $G[N_d(v)]$ is at most $f(d)$ for every graph $G \in \mathcal{G}$, every vertex v of G and every $d \geq 1$. Similarly, a class $\mathcal{G}$ of graphs *locally excludes a minor* if there exists an infinite sequence of graphs $H_1, H_2, \ldots$ such that for every graph $G \in \mathcal{G}$, every vertex v of G and every $d \geq 1$, the graph $G[N_d(v)]$ does not contain H_d as a minor.

Observe that every class of graphs with locally bounded tree-width locally excludes a minor. Similarly, every proper minor-closed class of graphs locally exclude a minor. We later define other locally restricted graph classes.

2.3 Grad and Expansion

We now present the framework of classes of graphs with bounded expansion and classes of nowhere-dense graph which was introduced by Nešetřil and Ossona de Mendéz in [24]. An *r-shallow minor* of a graph G is a graph obtained from G by removing some vertices and edges of G and contracting several vertex-disjoint subgraphs of radius at most r. Recall that the *radius* of a graph is the minimum r such that $G = G[N_r(v)]$ for some vertex v of G, i.e., every vertex of G is at distance at most r from v. If $\mathcal{G}$ is a class of graphs, then $\mathcal{G} \,\triangledown\, r$ is the class of all r-shallow minors of graphs contained in $\mathcal{G}$.

The edge-density of a graph G is $||G||/|G|$, i.e., the ratio of the number of edges of G and the number of its vertices. The *grad* $\triangledown_r(G)$ *with rank* r (greatest reduced average density) of a graph G is the maximum edge-density of an r-shallow minor of G. Observe that if $d \geq 2\nabla_0(G)$, then a graph G is d-degenerate. A class $\mathcal{G}$ of graphs has *bounded expansion* if there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that $\nabla_r(G) \leq f(r)$ for every graph $G \in \mathcal{G}$ and every $r \geq 1$.

Let us give few examples of classes of graphs with bounded expansion. Since every proper minor-closed class $\mathcal{G}$ of graphs is degenerate, the grads of all ranks of graphs contained in $\mathcal{G}$ are bounded by a constant. Hence, all proper minor-closed classes of graphs have bounded expansion. Another example of a class of graphs with bounded expansion are graphs with bounded maximum degree: if G has maximum degree Δ, then $\nabla_r(G) \leq \Delta(\Delta-1)^r/2$. Hence, classes of graphs with bounded maximum degree have also bounded expansion. Another example is the class of graphs that can be embedded to the plane in such a way that each edge is crossed by at most one other edge; this class contains graphs with arbitrary large degrees and is not minor-closed. Other examples of classes of graphs with bounded expansion can be found in [27].

Analogously to already introduced definitions, a class $\mathcal{G}$ of graphs has *locally bounded expansion* if there exists a function $f : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that $\nabla_r(G[N_d(v)]) \leq f(r,d)$ for every graph $G \in \mathcal{G}$, every vertex v of G and any two integers r and d. It can be shown that every class $\mathcal{G}$ of graphs with locally bounded expansion has almost bounded expansion in the following sense: for every $\varepsilon > 0$, there exist functions $f_r(n) : \mathbb{N} \to \mathbb{N}$ such that $f_r(n) \in O(n^\varepsilon)$ for every $r = 0, 1, \ldots$ and $\nabla_r(G) \leq f_r(n)$ for every n-vertex graph $G \in \mathcal{G}$.

This leads us to the definition of nowhere-dense graphs. If $\mathcal{G}$ is a class of graphs and f a real-valued function on the set of all graphs, then

$$\limsup_{G \in \mathcal{G}} f(G)$$

is the supremum of all reals α such that there exists an infinite sequence of distinct graphs $G_1, G_2, \ldots$ from $\mathcal{G}$ with $\alpha = \lim_{k\to\infty} f(G_k)$. The trichotomy theorem of Nešetřil and Ossona de Mendéz [24] asserts the following:

Theorem 1. *For every infinite class $\mathcal{G}$ of graphs, the following holds:*

$$\lim_{r\to\infty} \limsup_{G \in \mathcal{G}\nabla r} \frac{\log ||G||}{\log |G|} \in \{0, 1, 2\} . \tag{1}$$

Let us give the proof of this (at the first sight very surprising) theorem since it gives more insight into the structure of classes of graphs achieving each of the values of the limit.

Proof. If there exists a constant C such that every graph in $\mathcal{G}$ has at most C edges, then $\lim_{k\to\infty} \frac{\log ||G_k||}{\log |G_k|} = 0$ for every infinite sequence $G_1, G_2, \ldots$ of distinct graphs from $\mathcal{G} \nabla r$ (the number of vertices of the graphs G_i must grow to the infinity but the number of their edges is bounded by C).

If there is no constant C bounding the number of edges of every graph in $\mathcal{G}$, proceed as follows: choose G_1 to be K_2, clearly, $K_2 \in \mathcal{G} \nabla 0$. If $G_1, G_2, \ldots, G_k$ have already been fixed, choose G_k to be any graph of $\mathcal{G} \nabla 0$ containing more edges than G_{k-1} and subject to this with the minimum number of vertices. Observe that $|G_k| \leq 2||G_k||$ for every k (otherwise, G_k contains an isolated vertex which contradicts our choice of G_k). It follows that

$$\liminf_{k\to\infty} \frac{\log ||G_k||}{\log |G_k||} \geq 1 .$$

Since $\mathcal{G} \triangledown 0 \subseteq \mathcal{G} \triangledown 1 \subseteq \cdots$, it follows that if the limit in (1) is not equal to zero, then the limit is at least one.

Assume now that the limit given in (1) is greater than 1 for $\mathcal{G}$. Hence, there exist r, $\varepsilon > 0$ and an infinite sequence of graphs $G_1, G_2, \ldots \in \mathcal{G} \triangledown r$ such that $||G_k|| \geq |G_k|^{1+\varepsilon}$. We now apply the following result from [8, Lemma 3.13]: for every $\varepsilon > 0$, there exist an integer d and $\delta > 0$ such that every n-vertex graph with average degree n^ε contains K_{n^δ} as a d-shallow minor. It follows that the class $\mathcal{G} \triangledown rd$ contains complete graphs of arbitrary order and the limit (1) is at least two. Since $||G|| \leq |G|^2/2$ for every graph G, the limit in (1) is at most two and the proof of the theorem is completed.

The classes $\mathcal{G}$ of graphs with the limit (1) equal to 0 or 1 are called classes of *nowhere-dense* graphs. It follows that every class of graphs with locally bounded expansion is a class of nowhere-dense graphs.

We finish this section with Figure 1 where the reader can find inclusions between graph classes we have introduced in this section.

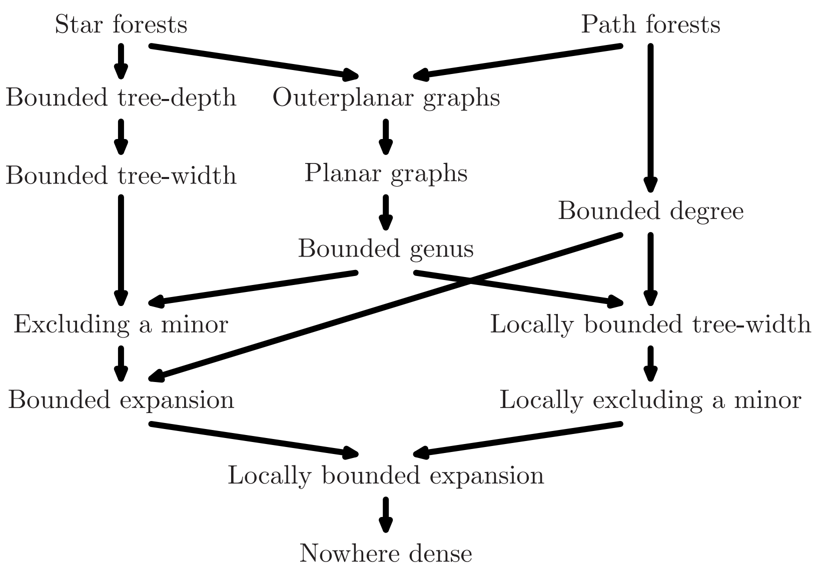

Fig. 1. Overview of inclusions between various graph classes

3 Structural Properties

In this section, we will introduce the notion of fraternal augmentations of orientations of graphs as defined by Nešetřil and Ossona de Mendéz in [20]. Since all the algorithms for classes of graphs with bounded expansion as well as classes of nowhere dense graphs are based on this notion, we decided to present it with full detail. The proofs follow the lines of those given in [20].

3.1 Orientations with Small In-Degree

Every graph $\mathcal{G}$ admits an orientation with maximum in-degree at most $\triangledown_0(G)$. The purpose is to develop a technique of augmenting orientations with small in-degrees preserving the fact that the grads remain small. To achieve this, we will need the following definition and lemma. Two vertices v and v' of a digraph[1] $\vec{G}$ are *k-reachable* for an integer k, if there exists a vertex w and oriented paths P and P' from v and v' to w of lengths ℓ and ℓ', respectively, such that $\ell + \ell' \leq k$. The paths P and P' form an (ℓ, ℓ')-*wedge* between v and v'.

We now state a key lemma for our further considerations.

Lemma 1. *There exist polynomials $P_k(x, y)$, $k = 1, 2, \ldots$, with the following properties. Let G be a graph and $\vec{G}$ an orientation of G with maximum in-degree Δ^-. If H_k is the graph with the vertex set of G and two vertices adjacent if they are k-reachable, then*

$$\triangledown_0(H_k) \leq P_k(\Delta^-, \triangledown_{k-1}(G)) .$$

Proof. The proof proceeds by induction on k. If $k = 1$, then the graphs H_1 and G are the same (observe that two vertices are 1-reachable if and only if they are adjacent in $\vec{G}$). Hence, $\triangledown_0(H_1) = \triangledown_0(G)$ and we can set $P_1(x, y) = y$.

Assume now that $k > 1$. Consider a proper vertex coloring of H_{k-1} with $2P_{k-1}(\Delta^-, \triangledown_{k-2}(G)) + 1$ colors (the existence of this coloring follows from the fact that H_{k-1} is $2P_{k-1}(\Delta^-, \triangledown_{k-2}(G))$-degenerate). Color now the arcs uv of $\vec{G}$ with pairs $[\alpha, \beta]$ of colors where α is the color of v and the color β is chosen in such a way that no two arcs coming to the same vertex have the same color. Since Δ^- choices of colors β suffice at each vertex, the arcs of $\vec{G}$ can be colored with at most $(2P_{k-1}(\Delta^-, \triangledown_{k-2}(G)) + 1)\Delta^-$ colors. Let K be this number of colors.

A *type* of an (ℓ, ℓ')-wedge formed by paths P and P' of lengths ℓ and ℓ' is the pair of two sequences of lengths ℓ and ℓ' formed by the colors of the arcs of P and P', respectively. Fix two integers ℓ and ℓ' such that $\ell + \ell' = k$ and $0 < \ell \leq \ell'$. Observe that the type of any (ℓ, ℓ')-wedge contains mutually distinct colors since the vertices with incoming arcs in an (ℓ, ℓ')-wedge have mutually distinct colors (they are $(k-1)$-reachable).

[1] We allow digraphs to have parallel arcs oriented in the opposite way. If we want to exclude parallel arcs, we will say that a digraph is simple.

Fix now two sequences σ and σ' of arc colors with lengths ℓ and ℓ' such that $\vec{G}$ contains an (ℓ, ℓ')-wedge of type $[\sigma, \sigma']$. Let F be the set of all arcs contained in an oriented path whose arcs are colored with the colors as in σ (respecting the order of the colors) and F' the set of all arcs contained in an oriented path whose arcs are colored as in σ'. Finally, F'' is the set of the arcs of F' that do not have the last color of σ'.

Now consider two paths P and Q of lengths ℓ with arc colors as in σ. Since the colors of all the $\ell+1$ vertices of P are mutually distinct as well as the colors of the $\ell+1$ vertices of Q and no vertex has two incoming arcs with the same color, P and Q are either vertex-disjoint or $P \cap Q$ is the initial sequence of both the paths. Hence, the arcs of F form (vertex-disjoint) out-branchings of depth ℓ in $\vec{G}$. The analogous reasoning also applies to F' and thus the arcs of F' form out-branchings of depth ℓ'.

Consider now two paths P and P' of lengths ℓ and ℓ' with arc colors as in σ and σ', respectively. Since only the pair of the first vertices of P and P' or the pair of last vertices (or both these pairs) can have the same color, either P and P' are vertex-disjoint, or they share their first vertices, or they share their last vertices, or they share both their first and last vertices. Hence, $F \cup F''$ form out-branchings rooted at their original vertices.

Let $\vec{G}'$ be the graph obtained from $\vec{G}$ by removing vertices not incident with arcs of $F \cup F''$ and contracting the out-branchings of $F \cup F''$. Since every leaf of any out-branching of $F \cup F''$ is at distance at most $\max\{\ell, \ell'-1\} \leq k-1$, the graph $\vec{G}'$ is a $(k-1)$-shallow minor of G (after disregarding the orientations of its arcs). If v and v' are k-reachable because of an (ℓ, ℓ')-wedge of type $[\sigma, \sigma']$, then v and v' are roots of out-branchings in $F \cup F''$ and they are adjacent after contracting these out-branchings (through the arc with the last color in σ'). We conclude that the edges between vertices v and v' that are k-reachable because of an (ℓ, ℓ')-wedge of type $[\sigma, \sigma']$ can be oriented in such a way that the in-degree of any vertex is at most $\bigtriangledown_{k-1}(G)$.

Ranging through all choices of $\ell + \ell' = k$ with $\ell > 0$ and $\ell' > 0$ and all choices of σ and σ', we obtain an upper bound of $2(k-1)K^k \bigtriangledown_{k-1}(G)$ on the number of incoming arcs added to H_k. If $\ell = 0$ and $\ell' = k$, then we just orient the new edges based on the direction of the paths they correspond to which increases the in-degree of each vertex by at most $(\Delta^-)^k$. Taking into account the edges present in H_{k-1}, we obtain that H_k has an orientation of its edges with maximum in-degree at most

$$\bigtriangledown_0(H_{k-1}) + (\Delta^-)^k + 2(k-1)K^k \bigtriangledown_{k-1}(G)$$

which is bounded by

$$2(k-1)((2P_{k-1}(\Delta^-, \bigtriangledown_{k-1}(G)) + 2)(\Delta^- + 1))^k \bigtriangledown_{k-1}(G) .$$

The sought polynomial $P_k(x, y)$ can be set to be equal to $4(k-1)((P_{k-1}(x, y) + 1)(x+1))^k y$.

We now define a crucial notion of transitive fraternal augmentations. If $\vec{G}$ is a simple digraph, then the *transitive fraternal augmentation* of $\vec{G}$ is a simple digraph obtained from $\vec{G}$ by adding the following arcs:

1. **transitive arcs**: if uv and vw are arcs of $\vec{G}$, then the arc uw is added unless $\vec{G}$ already contains the arc uw or the arc wu.
2. **fraternal arcs**: if uv and $u'v$ are arcs of $\vec{G}$, then the arc uu' or the arc $u'u$ is added unless $\vec{G}$ already contains the arc uu' or the arc $u'u$.

Our aim is to add fraternal arcs (where it is possible to make a choice which arc to add) in such a way that the maximum in-degree of $\vec{G}$ does not increase significantly. To choose the fraternal arcs, we can apply Lemma 1. However, we would like to iterate the process and thus we need to have a bound on grads of the transitive fraternal augmentation. Such a bound is given in the following theorem from [20]:

Theorem 2. *There exists polynomial $Q_1(x, y), Q_2(x, y), \ldots$ with the following properties. Let G be a graph and $\vec{G}$ an orientation of G with maximum in-degree Δ^-. If H is the graph containing all the edges of the transitive fraternal augmentation of $\vec{G}$, then the following holds for every $r \geq 1$:*

$$\nabla_r(H) \leq Q_r(\Delta^-, \nabla_{2r+1}(G))$$

Proof. Let $V_1, \ldots, V_n$ be subsets of vertices of H such that the radius of $H[V_i]$ is at most r for every $i = 1, \ldots, n$. Let v_i be the center of $H[V_i]$. Consider the shortest distance tree T_i in $H[V_i]$ rooted at v_i and orient the edges of T_i in the direction from v_i. We now modify the simple digraph $\vec{G}$ in another digraph $\vec{G}'$ which need not be simple. If the arc uw of T_i corresponds to an edge of G, add the arc uw to $\vec{G}$. If the arc uw is a transitive edge corresponding to arcs wv and vu, add the arc vw. If the arc uw is a transitive edge corresponding to arcs uv and vw, no action is required. Finally, if the arc uw is a fraternal edge corresponding to arcs uv and wv, add arcs uv and vw.

Observe that the maximum in-degree of $\vec{G}'$ is at most $2\Delta^-+1$: if an arc leading to v is added because of an arc uw of some T_i, then both u and w are in the same V_i and uv or wv is an arc of $\vec{G}$. Since the sets V_i are disjoint, at most Δ^- arcs leading to v can be added. The extra one in the estimate corresponds to an arc added because of the tree containing v.

If the subgraphs $H[V_i]$ and $H[V_j]$ are joined by an edge in H, then v_i and v_j are $2(r+1)$-reachable in $\vec{G}'$. In particular, the subgraph H' obtained from H by removing the vertices not contained in $V_1 \cup \cdots \cup V_n$ and contracting the subgraphs $H[V_i]$ is a subgraph of the graph H_{2r+2} as defined in Lemma 1. Consequently, $\nabla_0(H') \leq P_{2r+2)}(2\Delta^-, \nabla_{2r+1}(G))$ and thus $\nabla_r(H) \leq P_{2(r+1)}(2\Delta^-, \nabla_{2r+1}(G))$.

Theorem 2 guarantees us that there is a choice of fraternal arcs to be added such that the maximum in-degree of the transitive fraternal augmentation of $\vec{G}$ is at most $(\Delta^-)^2 + 2Q_0(\Delta^-, \nabla_1(G))$. Moreover, since the grads of the transitive fraternal augmentations are bounded by polynomials in Δ^- and grads of G, the process can be iterated. In particular, we obtain the following corollaries. Note that since the existence of an orientation with small maximum in-degree is guaranteed by the fact that the grad with rank 0 is bounded, we can use a greedy algorithm to construct it.

Corollary 1. *Let $\mathcal{G}$ be a class of graphs with bounded expansion. There exist $\Delta_0, \Delta_1, \ldots$ such that every graph $G \in \mathcal{G}$ has an orientation $\vec{G}_0$ with maximum in-degree Δ_0 and $\vec{G}_i$ has a transitive fraternal augmentation $\vec{G}_{i+1}$ with maximum in-degree Δ_{i+1} for every $i \geq 0$. Moreover, for every i, $\vec{G}_i$ can be computed in time linear in the number of vertices of G.*

Corollary 2. *Let $\mathcal{G}$ be a class of nowhere dense graphs. For every $\varepsilon > 0$, there exist functions $f_i : \mathbb{N} \to \mathbb{N}$, $i = 0, 1, \ldots$, $f_i(n) \in O(n^\varepsilon)$, such that every n-vertex graph $G \in \mathcal{G}$ has an orientation $\vec{G}_0$ with maximum in-degree $f_0(n)$ and $\vec{G}_i$ has a transitive fraternal augmentation $\vec{G}_{i+1}$ with maximum in-degree $f_{i+1}(n)$ for every $i \geq 0$. Moreover, for every i, $\vec{G}_i$ can be computed in time $O(n^{1+\varepsilon})$ in the number of vertices of G.*

3.2 Low Tree-Width and Low-Tree-Depth Coloring

We now mention one structural result on a special type of vertex colorings of graphs which is important for algorithmic applications and is of independent interest. In [7], DeVos et al. established the existence of low tree-width color with bounded number of colors for proper minor-closed classes of graphs:

Theorem 3. *Let $\mathcal{G}$ be a proper minor-closed class of graphs. For every k, there exists K such that every graph $G \in \mathcal{G}$ has a vertex coloring with K colors such that any k' color classes, $1 \leq k' \leq k$, induce a subgraph of G with tree-width at most $k' - 1$.*

Theorem 3 was strengthened by Nešetřil and Ossona de Mendéz in [20] in two directions: first, the result holds for more general graph classes and second it guarantees the existence of low tree-depth colorings.

Theorem 4. *Let $\mathcal{G}$ be a class of graphs with bounded expansion. For every k, there exists K such that ever graph $G \in \mathcal{G}$ has a vertex coloring with K colors such that any k' color classes, $1 \leq k' \leq k$, induce a subgraph of G with tree-depth at most k'. Moreover, such a coloring can be constructed in linear time for any graph G from $\mathcal{G}$.*

Theorem 4 is implied by the following lemma; we provide its short proof for completeness.

Lemma 2. *For every $p \geq 1$ and $d \geq 1$, the following holds: if $\vec{G}_0$ is an orientation of G, $\vec{G}_1, \vec{G}_2, \ldots$ a series of its transitive fraternal augmentations and H a connected subgraph of G with tree-depth at most d, then $\vec{G}_{3pd}[V(H)]$ either contains a clique of order p or an out-branching T of depth at most p such that*

1. *the end-vertices of every edge of H are joined by a directed path in T, and*
2. *if two vertices u and u' are joined by a directed path in T, then $\vec{G}_{3pd}[V(H)]$ contains the arc uu' or the arc $u'u$.*

Proof. Fix p and let the proof proceed by induction on d. If $d = 1$, H is a single vertex and the claim clearly holds. Assume that $d > 1$ and let v be a vertex of H such that the tree-depth of each component of $H \setminus v$ is at most $d - 1$. Let $V_1, \ldots, V_k$ be the vertex sets of the components of $H \setminus v$. By induction, $\vec{G}_{3pd-3p}[V_i]$ either contains a clique of order p or an out-branching T_i such that any edge of $G[V_i]$ joins a vertex with one of its ancestors and directed paths in T_i give rise to arcs in $\vec{G}_{3pd-3p}[V_i]$. If $\vec{G}_{3pd-3p}[V_i]$ for some i contains a clique of order p, then so does $\vec{G}_{3pd}[V(H)]$. Hence, we assume the existence of out-branchings T_i for all $i = 1, \ldots, k$.

Let r_i be the root vertex of T_i, $i = 1, \ldots, k$. We claim that r_i and v are adjacent in $\vec{G}_{3pd-2p-1}$: since H is connected, v is adjacent to one of the descendants of r_i in T_i, say w. Let $r_i w_1 \ldots w_\ell$ be the oriented path in T_i from r_i to $w = w_\ell$. Since the depth of T_i is at most p, $\ell \leq p - 1$. Applying the transitive or the fraternal rule (depending on the orientation of the arc between v and w_ℓ), we obtain that $\vec{G}_{3pd-3p+1}[V_i]$ contains an arc between v and $w_{\ell-1}$. Repeating the argument, we get that the vertices v and r_i are adjacent in $\vec{G}_{3pd-3p+p-1}[V_i] = \vec{G}_{3pd-2p-1}$. Observe that we have actually proven that if v is adjacent to a vertex u of an out-branching T_i in $\vec{G}_{3pd-3p}$, then v is adjacent in $\vec{G}_{3pd-2p-1}$ to all the vertices on the path from r_i to u.

Let q be the first index such that the in-degree of v is the same in $\vec{G}_{3pd-2p-1+q}$ and $\vec{G}_{3pd-2p-1+q+1}$. If $q \geq p$, then the in-degree of v in $\vec{G}_{3pd-2p-1+p}$ is at least p and thus $\vec{G}_{3pd-p}$ contains a clique of order p (on the in-neighbors of v in $\vec{G}_{3pd-p-1}$). Consequently, $\vec{G}_{3pd}[V(H)]$ contains a clique of order p. Hence, we can assume that $q \leq p - 1$.

Let W be the set of vertices w of H such that $\vec{G}_{3pd-2p-1+q}$ contains the arc wv and all the vertices on the path from the r_i to w in the out-branching T_i containing w are in-neighbors of v. Since $\vec{G}_{3pd-2p+q}$ contains an arc between any two vertices of W by the fraternity rule, $\vec{G}_{3pd-2p+q}[W \cup \{v\}]$ contains a directed Hamilton path, say $w_1, \ldots, w_\ell$. Observe that $w_\ell = v$ because of the choice of W.

Let $T'_1, \ldots, T'_{k'}$ be the out-branchings obtained from T_i by removing the vertices contained in W and let $r'_1, \ldots, r'_{k'}$ be their roots. Consider now the out-branching T in $\vec{G}_{3pd-2p+q}$ formed by the path $w_1 \ldots w_\ell$, the out-branchings $T'_1, \ldots, T'_{k'}$ and the arcs $w_i r'_j$ for $j = 1, \ldots, k'$ where i is the maximum index such that $\vec{G}_{3pd-2p+q}$ contains the arc $w_i r'_j$. Such an index i must exist since either r'_j is a root of one of the out-branchings $T_1, \ldots, T_k$ and thus $\vec{G}_{3pd-2p+q}$ contains the arc $w_\ell r'_j$ or W contains the in-neighbor of r'_j in one of the out-branchings $T_1, \ldots, T_k$. Hence, T is an out-branching contained in $\vec{G}_{3pd-2p+q}$.

We now verify that the end-vertices of every edge uu' of H are joined by a directed path in T. If $u = v$, then either $u' \in W$ (and thus the existence of the path follows) or u' is contained in one of the out-branchings $T'_1, \ldots, T'_{k'}$, say

T'_j. Since $r'_j \notin W$, T contains the arc vr'_j (here, we use that all the vertices between the root of T_i and a vertex of T_i adjacent to v are also adjacent to v in $\vec{G}_{3pd-2p-1} \subseteq \vec{G}_{3pd-2p+q}$). The case $u' = v$ is symmetric and thus we can assume that neither u nor u' is v.

If neither u nor u' is contained in W, then, by induction, they are contained in the same out-branching T'_j and are joined by a directed path in T. If both u and u' are contained in W, then they are clearly joined by a directed path in T since they are both contained in the path $w_1 \dots w_\ell$. It remains to consider the case when $u \in W$ and $u' \notin W$. Let m the index such that $w_m = u$. Since u and u' are adjacent in G, they are contained in the same out-branching T_i. Further assume that u' is contained in an out-branching T'_j. By induction, $\vec{G}_{3pd-3p}$ contains either the arc ur'_j or the arc $r'_j u$. If the arc ur'_j is present in $\vec{G}_{3pd-3p}$, then r'_j is adjacent to a vertex $w_{m'}$ with $m' \geq m$ in T. If the arc $r'_j u$ is present in $\vec{G}_{3pd-3p}$, then $\vec{G}_{3pd-2p+q}$ contains an arc between r'_j and v since $\vec{G}_{3pd-2p-1+q}$ contains the arc uv. If $\vec{G}_{3pd-2p+q}$ contained the arc $r'_j v$, the choice of q would imply that $\vec{G}_{3pd-2p-1+q}$ also contained the arc $r'_j v$ which would imply that r'_j should have been included in W. Otherwise, $\vec{G}_{3pd-2p+q}$ contains the arc vr'_j, thus the arc vr'_j is also contained in T and u and u' are joined by a directed path in T.

We have shown that the out-branching T satisfies that any two end-vertices of an edge of H are joined by a directed path in T. Since $q \leq p - 1$, T is an out-branching in $\vec{G}_{3pd-p}$. We claim that if $u_0, \dots, u_m$ is a directed path in $\vec{G}_{3pd-p}$, then the vertices $u_0, \dots, u_m$ form a clique in $\vec{G}_{3pd-p+m}$. Proceed by induction on m: if $m = 1$, there is nothing to prove. Otherwise, $\vec{G}_{3pd-p+m-1}$ contains a clique on the vertices $u_1, \dots, u_m$. By the fraternity or transitivity rule, $\vec{G}_{3pd-p+m}$ contains an arc between u_0 and each of the vertices $u_1, \dots, u_m$. Hence, the vertices $u_0, \dots, u_m$ form a clique in $\vec{G}_{3pd-p+m}$. We conclude that if the depth of T is at least p, $\vec{G}_{3pd}[V(T)] = \vec{G}_{3pd}[V(H)]$ contains a clique of order p, and if the depth of T is less than p, then any two vertices joined by a directed path in T are adjacent in $\vec{G}_{3pd}$. The proof of the lemma is now finished.

4 Testing Graph Properties

In the final section of the paper, we want to focus on meta-algorithmic results for classes of graphs with restricted structure. Let us remark that the results we present in this section readily translate to relational structures by considering the concept of Gaifman graph. If R is a relational structure with a domain D, then the *Gaifman graph* of R is the graph with vertex set D where two distinct elements x and y of D are joined by an edge if R contains a relation including both x and y. For instance, if a graph G is viewed as a binary relational structure, then the Gaifman graph of G is G itself. Graph concepts we have introduced translate to relational structures by considering corresponding Gaifman graphs; e.g., the class of relational structures has bounded expansion, if the class of their Gaifman graphs has bounded expansion. The results we present further also hold for corresponding classes of relational structures under the assumption that the vocabulary is finite, i.e., the number of different types of relations is finite.

4.1 Σ_1-Properties

Analogously to Σ_1-formulas, which are first-order formulas with existential quantifiers only, a Σ_1-property is a property that can be described by Σ_1-formula. The easiest problem of this kind is testing the existence of a subgraph. Eppstein [10, 11] constructed a linear-time algorithm for deciding the existence of a fixed subgraph for planar graphs. He then extended his algorithm to minor-closed classes of graphs with locally bounded tree-width [12]. All these results were generalized to classes of graphs with bounded expansion by Nešetřil and Ossona de Mendéz in [21, 23]. In fact, they established a more general result on testing arbitrary Σ_1-properties:

Theorem 5. *Let Φ be a Σ_1-property and $\mathcal{G}$ a class of graphs with bounded expansion. There exists a linear time algorithm deciding Φ for graphs $G \in \mathcal{G}$.*

The main idea of the algorithm is that if Φ holds for $G \in \mathcal{G}$, then the witness assignment to variables can use at most k colors where k is the number of quantifiers of Φ. Hence, using Theorem 4, we can color vertices with K colors in such a way that any k colors induce a graph with tree-depth at most k. After finding the coloring (in linear time), the problem is reduced to deciding Φ for $\binom{K}{k}$ subgraphs of an input graph, each subgraph having tree-depth at most k (which can be solved, e.g., using the classical Courcelle's result mentioned at the beginning of the paper).

Following the lines of the above reasoning, we can obtain an analogous results for classes of nowhere-dense graphs, see [24] for further details. An algorithm is *almost linear*, if for any $\varepsilon > 0$ which is part of the input of the algorithm, the algorithm runs in time $O(n^{1+\varepsilon})$ where n is the number of vertices of an input graph.

Theorem 6. *Let Φ be a Σ_1-property and $\mathcal{G}$ a class of nowhere dense graphs. There exists an almost linear time algorithm deciding Φ for graphs $G \in \mathcal{G}$.*

Let us now focus on a particular case of Σ_1-properties, the existence of short paths between two vertices. Kowalik and Kurowski [17,18] designed a data structure with linear build-up time and constant query time answering the existence of a path of length at most d between two vertices of an input planar graph for a fixed integer d. In fact, they approach readily generalize to classes of graphs with bounded expansion. Let us sketch the main idea of the algorithm: let G be an input graph and consider the sequence of its transitive fraternal augmentations $\vec{G}_0, \ldots, \vec{G}_d$ as defined in Corollary 1. If two vertices u and v are joined by a path of length at most d, then they are either adjacent in $\vec{G}_d$ or they have a common in-neighbor in $\vec{G}_d$ (this can easily be proved by induction on d). Since the maximum in-degree of $\vec{G}_d$ is bounded, the existence of an edge joining the two vertices or the existence of their common in-neighbor can be done in constant time.

The data structure can be dynamized using the result of Brodal and Fagerberg [4] on maintaining orientations with small in-degrees of degenerate graphs, see [16]. The arguments readily translate to a setting of classes of graphs with bounded expansion, see [9]:

Theorem 7. *Let $\mathcal{G}$ be a class of graphs with bounded expansion and d a fixed integer. There exists a dynamic data structure for answering the existence of a path of length at most d in a graph $G \in \mathcal{G}$ with the following parameters:*

- *the data structure can be built in linear time,*
- *each query can be answered in constant time,*
- *an edge can be added to the represented graph in time $O(\log^d n)$ where n is its number of vertices, and*
- *an edge can be removed in constant time.*

4.2 First-Order Properties

We now address the complexity of deciding general first-order properties, i.e., those properties that can be described by formulas with quantifications over graph vertices only (quantification over sets of vertices is not allowed). As examples of first-order properties, we can mention deciding the existence of a dominating set of a fixed size or the existence of a vertex cover of a fixed size. First-order properties can always be decided in polynomial time (with degree depending on the property) but we are interested in fixed parameter results. The first result in this direction is the result of Seese [31] that every first-order property can be tested in linear time for any class of graphs with bounded maximum degree. The result is not that surprising after we realize that first-order properties are of very localized nature which is captured in the following classical result of Gaifman [15]:

Theorem 8. *Every first-order formula Φ for graphs is equivalent for some r to a Boolean combination of formulas of the form*

$$\exists x_1 \cdots \exists x_k \left(\wedge_{i \neq j} \mathrm{dist}(x_i, x_j) > 2r \bigwedge \wedge_{i=1,\ldots,k} \Phi_r(x_i) \right)$$

where each $\Phi_r(x_i)$ is r-local with respect to x_i, i.e., all quantifiers contained in $\Phi_r(x_i)$ have domain restricted to the r-neighborhood of x_i.

In the light of Theorem 8, deciding first-order properties for graphs with maximum degree Δ decomposes into a linear number of finite problems (the r-neighborhood of each vertex contains at most $\Delta(\Delta - 1)^{r-1}$ vertices) whose Boolean combination yields the result on whether the formula is satisfied for an input graph.

Frick and Grohe [13, 14] extended this result by considering classes of graphs with locally bounded tree-width. They have shown that any first-order property can be decided in almost linear time for any class of graphs with locally bounded tree-width. In the particular case of planar graphs, they were able to obtain a linear time algorithm using a different covering algorithm. Their algorithm uses the following covering result of Peleg [28] which has applications in other algorithms in the area and thus we would like to mention it explicitly.

Lemma 3. *Let $k \geq 1$ be a fixed integer. There is an algorithm that given $r \geq 1$ and a graph G, outputs sets $A_1, \ldots, A_m$ of vertices of G such that*

- *for every vertex v of G, there exists A_i containing the r-neighborhood of v,*
- *every A_i is contained in the $2kr$-neighborhood of a vertex of G, and*
- *the sum $|A_1| + \cdots + |A_m|$ is at most $O(n^{1+1/k})$.*

The running time of the algorithm is linear in the sum of the numbers of edges contained in $G[A_i]$, $i = 1, \ldots, k$.

Another meta-theorem on graphs with locally restricted structure was obtained by Dawar, Grohe and Kreutzer [6] who showed that deciding first-order properties Φ is fixed-parameter tractable for classes of graphs locally excluding a minor, i.e., there exists a polynomial-time algorithm where the exponent does not depend on Φ. Nešetřil and Ossona de Mendéz [25] gave a linear time algorithm for deciding the existence of a dominating set of a fixed size for classes of graphs with bounded expansion. Their result indicates that the results we mention earlier could hold for classes of graphs with bounded expansion. This turns out to be true as proven by the authors and Thomas in [9]:

Theorem 9. *Let Φ be a first order formula and $\mathcal{G}$ a class of graphs with bounded expansion. There exists a linear-time algorithm deciding Φ for graphs $G \in \mathcal{G}$.*

Theorem 10. *Let Φ be a first order formula and $\mathcal{G}$ a class of nowhere dense graphs. There exists an almost linear time algorithm deciding Φ for graphs $G \in \mathcal{G}$.*

References

1. Arnborg, S., Proskurowski, A., Corneil, D.G.: Forbidden minors characterization of partial 3-trees. Discrete Math. 80, 1–19 (1990)
2. Bodlaender, H.L., Deogun, J.S., Jansen, K., Kloks, T., Kratsch, D., Müller, H., Tuza, Z.: Ranking of graphs. In: Mayr, E.W., Schmidt, G., Tinhofer, G. (eds.) WG 1994. LNCS, vol. 903, pp. 292–304. Springer, Heidelberg (1995)
3. Bodlaender, H.L., Gilbert, J.R., Hafsteinsson, H., Kloks, T.: Approximating treewidth, pathwidth, frontsize and shortest elimination tree. J. Algorithms 25, 1305–1317 (1996)
4. Brodal, G.S., Fagelberg, R.: Dynamic representations of sparse graphs. In: Dehne, F., Gupta, A., Sack, J.-R., Tamassia, R. (eds.) WADS 1999. LNCS, vol. 1663, pp. 342–351. Springer, Heidelberg (1999)
5. Courcelle, B.: The monadic second-order logic of graph I. Recognizable sets of finite graphs. Inform. and Comput. 85, 12–75 (1990)
6. Dawar, A., Grohe, M., Kreutzer, S.: Locally excluding a minor. In: Proc. LICS 2007, pp. 270–279. IEEE Computer Society Press, Los Alamitos (2007)

7. DeVos, M., Ding, G., Oporowski, B., Sanders, D.P., Reed, B., Seymour, P.D., Vertigan, D.: Excluding any graph as a minor allows a low tree-width 2-coloring. J. Combin. Theory Ser. B 91, 25–41 (2004)
8. Dvořák, Z.: Assymptotical structure of combinatorial objects, PhD thesis, Charles University (2007)
9. Dvořák, Z., Král', D., Thomas, R.: Deciding first order properties for nowhere dense graphs (manuscript) (2009)
10. Eppstein, D.: Subgraph isomorphism in planar graphs and related problems. In: Proc. SODA 1995, pp. 632–640. ACM&SIAM (1995)
11. Eppstein, D.: Subgraph isomorphism in planar graphs and related problems. J. Graph Algorithms Appl. 3, 1–27 (1999)
12. Eppstein, D.: Diameter and treewidth in minor-closed graph families. Algorithmica 27, 275–291 (2000)
13. Frick, M., Grohe, M.: Deciding first-order properties of locally tree-decomposable structures. In: Wiedermann, J., Van Emde Boas, P., Nielsen, M. (eds.) ICALP 1999. LNCS, vol. 1644, pp. 331–340. Springer, Heidelberg (1999)
14. Frick, M., Grohe, M.: Deciding first-order properties of locally tree-decomposable structures. J. ACM 48, 1184–1206 (2001)
15. Gaifman, H.: On local and non-local properties. In: Proc. Herbrand Symp. Logic Colloq. North-Holland, Amsterdam (1982)
16. Kowalik, L.: Adjacency queries in dynamic sparse graphs. Inf. Process. Lett. 102, 191–195 (2007)
17. Kowalik, L., Kurowski, M.: Oracles for bounded-length shortest paths in planar graphs. ACM Trans. Algorithms 2, 335–363 (2006)
18. Kowalik, L., Kurowski, M.: Short path queries in planar graphs in constant time. In: Proc. STOC 2003, pp. 143–148 (2003)
19. Nešetřil, J., Ossona de Mendéz, P.: First order properties of nowhere dense structures (manuscript) (2008)
20. Nešetřil, J., Ossona de Mendéz, P.: Grad and classes with bounded expansion I. Decompositions. European J. Combin. 29, 760–776 (2008)
21. Nešetřil, J., Ossona de Mendéz, P.: Grad and classes with bounded expansion II. Algorithmic aspects. European J. Combin. 29, 777–791 (2008)
22. Nešetřil, J., Ossona de Mendéz, P.: Grad and classes with bounded expansion III. Restricted graph homomorphism dualities. European J. Combin. 29, 1012–1024 (2008)
23. Nešetřil, J., Ossona de Mendéz, P.: Linear time low tree-width partitions and algorithmic consequences. In: Proc. STOC 2006, pp. 391–400 (2006)
24. Nešetřil, J., Ossona de Mendéz, P.: On nowhere dense graphs (manuscript) (2008)
25. Nešetřil, J., Ossona de Mendéz, P.: Structural properties of sparse graphs (manuscript) (2008)
26. Nešetřil, J., Ossona de Mendéz, P.: Tree depth, subgraph coloring and homomorphism bounds. European J. Combin. 27, 1022–1041 (2006)
27. Nešetřil, J., Ossona de Mendéz, P., Wood, D.: Characterizations and examples of graph classes with bounded expansion (manuscript) (2008)
28. Peleg, D.: Distance-dependent distributed directories. Info. Computa. 103, 270–298 (1993)
29. Robertson, N., Seymour, P.D.: Graph minors V: Excluding a planar graph. J. Combin. Theory Ser. B 41, 92–114 (1986)
30. Robertson, N., Seymour, P.D.: Graph Minors XX. Wagner's conjecture. J. Combin. Theory Ser. B 92, 325–357 (2004)
31. Seese, D.: Linear time computable problems and first-order descriptions. Math. Structu. Comput. Sci. 6, 505–526 (1996)