Distributed domination on sparse graph classes

Ozan Heyd

University of Bremen, Bremen, Germany

Simeon Kublenz

University of Bremen, Bremen, Germany

Patrice Ossona de Mendez

Centre d'Analyse et de Mathématiques Sociales (CNRS, UMR 8557), Paris, France and Computer Science Institute of Charles University, Praha, Czech Republic

Sebastian Siebertz

University of Bremen, Bremen, Germany

Alexandre Vigny

University of Bremen, Bremen, Germany

Abstract

We show that the dominating set problem admits a constant factor approximation in a constant number of rounds in the LOCAL model of distributed computing on graph classes with bounded expansion. This generalizes a result of Czygrinow et al. for graphs with excluded topological minors to very general classes of uniformly sparse graphs. We demonstrate how our general algorithm can be modified and fine-tuned to compute an $(11+\varepsilon)$-approximation (for any $\varepsilon > 0$) of a minimum dominating set on planar graphs. This improves on the previously best known approximation factor of 52 on planar graphs, which was achieved by an elegant and simple algorithm of Lenzen et al.

Keywords: Dominating set, Distributed LOCAL algorithms, Bounded expansion graph classes, Planar graphs.

This journal paper subsumes the results of the extended abstracts [35] and [29].

⋆This paper is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreements No 810115 – DYNASNET and 948057) and from the German Research Foundation (DFG) with grant agreement No 444419611.

Email addresses: heyd@uni-bremen.de (Ozan Heyd), kublenz@uni-bremen.de (Simeon Kublenz), pom@ehess.fr (Patrice Ossona de Mendez), siebertz@uni-bremen.de (Sebastian Siebertz), vigny@uni-bremen.de (Alexandre Vigny)

1. Introduction

A dominating set in an undirected and simple graph G is a set $D \subseteq V(G)$ such that every vertex $v \in V(G)$ either belongs to D or has a neighbor in D. The dominating set problem has many applications in theory and practice, see e.g. [17, 46], unfortunately however, already the decision problem whether a graph admits a dominating set of size k is NP-hard [33] and this even holds in very restricted settings, e.g. on planar graphs of maximum degree 3 [25].

Consequently, attention shifted from computing exact solutions to approximating near optimal dominating sets. A simple greedy algorithm computes an $\ln n$ approximation (where n is number of vertices of the input graph) of a minimum dominating set [31, 41], and for general graphs this algorithm is near optimal – it is NP-hard to approximate minimum dominating sets within factor $(1-\varepsilon)\ln n$ for every $\varepsilon > 0$ [15].

Therefore, researchers tried to identify restricted graph classes where better (sequential) approximations are possible. For example, the problem admits a PTAS on classes with subexponential expansion [28]. Here, expansion refers to the edge density of bounded depth minors, which we will define formally below. Important examples of classes with subexponential expansion include the class of planar graphs and more generally classes that exclude some fixed graph as a minor. The dominating set problem admits a constant factor approximation on classes of bounded degeneracy (equivalently, of bounded arboricity) [7, 40] and an $\mathcal{O}(\ln \gamma)$ approximation (where γ denotes the size of a minimum dominating set) on classes of bounded VC-dimension [10, 21]. In fact, the greedy algorithm can be modified to yield an $\mathcal{O}(\ln \gamma)$ approximation on biclique-free graphs (graphs that exclude some fixed complete bipartite graph $K_{t,t}$ as a subgraph) [47] and even a constant factor approximation on graphs with bounded degeneracy [32]. However, it is unlikely that polynomial-time constant factor approximations exist even on $K_{3,3}$-free graphs [47]. The general goal in this line of research is to identify the broadest graph classes on which the dominating set problem (or other important problems that are hard on general graphs) can be approximated efficiently with a certain guarantee on the approximation factor. These limits of tractability are often captured by abstract notions, such as expansion, degeneracy or VC-dimension of graph classes.

In this paper we study the distributed time complexity of finding dominating sets in the classic LOCAL model of distributed computing, which can be traced back at least to the seminal work of Gallager, Humblet and Spira [24]. In this model, a distributed system is modeled by an undirected (connected) graph G, in which every vertex represents a computational entity of the network and every edge represents a bidirectional communication channel. The vertices are equipped with unique identifiers. In a distributed algorithm, initially, the nodes have no knowledge about the network graph. They must then communicate and coordinate their actions by passing messages to one another in order to achieve a common goal, in our case, to compute a dominating set of the network graph. The LOCAL model focuses on the aspects of communication complexity and therefore the main measure for the efficiency of a distributed algorithm is the number of communication rounds it needs until it returns its answer.

Kuhn et al. [37] proved that in r rounds on an n-vertex graphs of maximum degree Δ one can approximate minimum dominating sets only within a factor $\Omega(n^{c/r^2}/r)$ and $\Omega(\Delta^{1/(r+1)}/r)$, respectively, where c is a constant. This implies that, in general, to achieve a constant approximation ratio, we need at least $\Omega(\sqrt{\log n/\log\log n})$ and $\Omega(\log\Delta/\log\log\Delta)$ communication rounds, respectively. Kuhn et al. [37] also presented a $(1+\varepsilon)\ln\Delta$-approximation that runs in $\mathcal{O}(\log(n)/\varepsilon)$ rounds for any $\varepsilon > 0$, Barenboim et al. [8] presented a deterministic $\mathcal{O}((\log n)^{k-1})$-time algorithm

that provides an $\mathcal{O}(n^{1/k})$-approximation, for any integer parameter $k \geq 2$. More recently, the combined results of Rozhon, Ghaffari, Kuhn, and Maus [26, 45] provide an algorithm computing a $(1+\varepsilon)$-approximation of the dominating set in $\text{poly}(\log(n)/\varepsilon)$ rounds [45, Corollary 3.11].

Since by the results of Kuhn et al. [37] in general graphs it is not possible to compute a constant factor approximation in a constant number of rounds, much effort has been invested to improve the ratio between approximation factor and number of rounds on special graph classes. For graphs of degeneracy a (equivalent to arboricity up to factor 2), Lenzen and Wattenhofer [40] provided an algorithm that achieves a factor $\mathcal{O}(a^2)$ approximation in randomized time $\mathcal{O}(\log n)$, and a deterministic $\mathcal{O}(a \log \Delta)$ approximation algorithm with $\mathcal{O}(\log \Delta)$ rounds. Graphs of bounded degeneracy include all graphs that exclude a fixed graph as a (topological) minor and in particular, all planar graphs and any class of bounded genus.

Amiri et al. [1] provided a deterministic $\mathcal{O}(\log n)$ time constant factor approximation algorithm on classes of bounded expansion (which extends also to connected dominating sets). Czygrinow et al. [11] showed that for any given $\varepsilon > 0$, $(1+\varepsilon)$-approximations of a maximum independent set, a maximum matching, and a minimum dominating set, can be computed in $\mathcal{O}(\log^* n)$ rounds in planar graphs, which is asymptotically optimal [39].

Lenzen et al. [38] proved that on planar graphs a 130-approximation of a minimum dominating set can be computed in a constant number of rounds. A careful analysis of Wawrzyniak [50] later showed that the algorithm computes in fact a 52-approximation. In terms of lower bounds, Hilke et al. [30] showed that there is no deterministic local algorithm (constant-time distributed graph algorithm) that finds a $(7-\varepsilon)$-approximation of a minimum dominating set on planar graphs, for any positive constant ε. Better approximation ratios are known for some special cases, e.g. 32 if the planar graph is triangle-free [3, Theorem 2.1], 18 if the planar graph has girth five [4] and 5 if the graph is outerplanar (and this bound is tight) [9, Theorem 1]. Wawrzyniak [49] showed that message sizes of $\mathcal{O}(\log n)$ suffice to give a constant factor approximation on planar graphs in a constant number of rounds.

The constant factor approximations in a constant number of rounds for planar graphs were gradually extended to classes with bounded genus [2, 5], classes with sublogarithmic expansion [6] and eventually by Czygrinow et al. [12] to classes with excluded topological minors. Again, one of the main goals in this line of research is to find the most general graph classes on which the dominating set problem admits a constant factor approximation in a constant number of rounds.

We take a step towards this goal and generalize the result of Czygrinow et al. [12] to classes of bounded expansion. The notion of bounded expansion was introduced by Nešetřil and Ossona de Mendez [42] and offers an abstract definition of uniform sparseness in graphs. It is based on bounding the density of shallow minors. Intuitively, while a minor is obtained by contracting arbitrary connected subgraphs of a graph to single vertices, in an r-shallow minor we are only allowed to contract connected subgraphs of radius at most r.

A class of graphs has bounded expansion if for every radius r the set of all r-shallow minors has edge density bounded by a constant depending only on r. We write $\nabla_r(G)$ for the maximal edge density of an r-shallow minor of a graph G. Of course, every class $\mathscr{C}$ that excludes a fixed graph H as a minor has bounded expansion. For such classes there exists an absolute constant c such that for all $G \in \mathscr{C}$ and all r we have $\nabla_r(G) \leq c$. Special cases are the class of planar graphs, every class of graphs that can be drawn with a bounded number of crossings, and every class of graphs that embeds into a fixed surface. Every class of intersection graphs of low density objects in low dimensional Euclidean space has polynomial expansion, that is, the function ∇_r is bounded polynomially in r on $\mathscr{C}$. Also every class $\mathscr{C}$ that excludes a fixed graph H as a topological minor

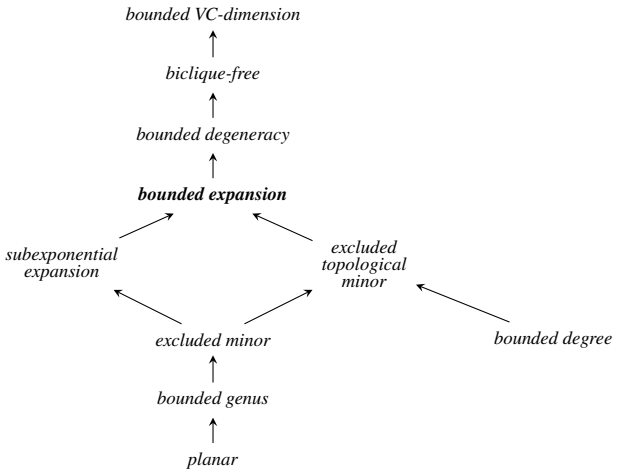

Figure 1: Inclusion diagram of the mentioned graph classes.

has bounded expansion. Important special cases are classes of bounded degree and classes of graphs that can be drawn with a linear number of crossings. Further examples include classes of graphs with bounded queue-number, bounded stack-number or bounded non-repetitive chromatic number. Also, for each constant $d > 0$ there is a bounded expansion class $\mathscr{R}_d$, to which the Erdős-Rényi random graphs $G(n, d/n)$ belong asymptotically almost surely. See [28, 44] for reference to all these examples.

Classes of bounded expansion are more general than classes excluding a topological minor. However, maybe not surprisingly, when performing local computations, it is not properties of minors or topological minors, but rather of shallow minors that enable the necessary combinatorial arguments in the algorithms. This observation was already made in the study of the kernelization complexity of dominating set on classes of sparse graphs [16, 19, 20, 22, 34]. Moreover, bounding the edge density of shallow minors might be needed only up to some depth. For example, degenerate classes are those classes where only $\nabla_0(G)$, the edge density of subgraphs, is bounded, and these classes are more general than classes of bounded expansion.

The algorithm of Czygrinow et al. [12] for classes excluding a topological minor is based on an quite complicated iterative process of choosing dominating vertices from so called *pseudo-covers*. Based on the fact that classes excluding a topological minor in particular exclude some complete bipartite graph $K_{t,t}$ as a subgraph, it is proved that this iterative process terminates after at most t rounds and produces a good approximation of a minimum dominating set.

2. Our contribution

Our contribution is threefold: First, we simplify the arguments used by Czygrinow et al. and give a more accessible description of their algorithm. Second, we identify the boundedness of $\nabla_1(G)$ as the key property that makes the algorithm work. Classes with only this restriction are less general than degenerate classes, but more general than bounded expansion classes. We generalize the algorithm to these general classes and prove that the pseudo-covering method cannot

be extended further, e.g. to classes of bounded degeneracy. Last, we optimize the bounds that arise in the algorithm in terms of several parameters. Czygrinow et al. explicitly stated that they did not aim to optimize any constants, and as presented, the constants in their construction are enormous. Even though the constants in our analysis are still large, they are by magnitudes smaller than those in the original presentation. The following is our first main theorem.

> *Theorem* 2.1. Let $\nabla_1 > 0$ be an integer. There exists a LOCAL algorithm that computes in a constant number of rounds, for any input graph G with $\nabla_1(G) \leq \nabla_1$, a dominating set of size $\mathcal{O}(\gamma(G))$, where $\gamma(G)$ denotes the size of a minimum dominating set of G.

Note that the algorithm depends on the constant ∇_1. The reason for this is that we cannot locally compute or approximate $\nabla_1(G)$ in a constant number of rounds. This is also the case for the algorithm of Czygrinow et al., which works with the assumption that the inputs exclude a complete graph K_t with t vertices as a topological minor, a property that can also not be verified locally. Furthermore, the number of rounds depends on ∇_1.

The algorithm is actually tuned using more parameters, like upper bounds on $\nabla_0(G)$ and on integers s and t such that the complete bipartite graph $K_{s,t}$ is not subgraph of G, in order to improve the approximation ratio of the algorithm. However, all these parameters can be upper bounded in terms of ∇_1. When these parameters are given, the algorithm computes a $(2(\nabla_0+1)((2\nabla_1)^{4s\nabla_1}+2)\gamma$ approximation in a number of rounds that depends on ∇_1 and t.

Then, we modify and fine-tune the algorithm for graphs excluding $K_{3,t}$ as a subgraph (and having ∇_1 bounded). Important examples of graphs with this property are graphs that can be embedded into a fixed surface of bounded genus. We prove the following theorem.

> *Theorem* 2.2. Let $\nabla_1 > 0$ and $t \geq 3$ be integers. There exists a LOCAL algorithm and a function C that for every $K_{3,t}$-free graph G with $\nabla_1(G) \leq \nabla_1$ and every $\varepsilon > 0$, computes in $C(\varepsilon)$ rounds a dominating set of size at most $(6\nabla_1 + 3)\gamma$.

Since planar graphs satisfy $\nabla_1 \leq 3$ and exclude $K3,3$ as a subgraph, from Theorem 2.2 we can derive an approximation factor of 21 for planar graphs. A more careful analysis leads to the following theorem.

> *Theorem* 2.3. There exists a LOCAL algorithm and a function C that for every planar input graph G and $\varepsilon > 0$, computes in $C(\varepsilon)$ rounds a dominating set of size at most $(11 + \varepsilon) \cdot \gamma(G)$.

We further analyze our algorithm on restricted classes of planar graphs and improve the upper bounds in several cases (see Table 1).

Graph class	Lower bound	Previous upper bound	Our upper bound
Planar graphs	$7 - \varepsilon$ [30]	52 [50]	$11 + \varepsilon$
Triangle-free planar graphs		32 [3]	$8 + \varepsilon$
Bipartite planar graphs			$7 + \varepsilon$
Outerplanar graphs	5 [9]	5 [9]	$(8 + \varepsilon)$
Girth ≥ 5 planar graphs		18 [4]	7

Table 1: Approximation factors for a LOCAL approximation of $\gamma(G)$ is a constant number of rounds. Our algorithm improves the approximation factors in all these cases, except for the class of outerplanar graphs.

Before we go into the technical details, let us give an overview of the algorithm. The algorithm works in three phases. Each phase ($i \in \{1, 2, 3\}$) computes a small set D_i that is added to the output dominating set (where by a small set we mean a set whose size is linear in γ).

The first phase is a preprocessing phase, which was similarly employed in the algorithm of Lenzen et al. [38]. In a key lemma, Lenzen et al. proved that for planar graphs there are only few vertices whose open neighborhood cannot be dominated by at most six vertices. This lemma generalizes to graphs G where $\nabla_1(G)$ is bounded as shown in [6] (where six is replaced by a different constant depending on $\nabla_1(G)$). We improve this general lemma and derive in particular that in the case of planar graphs there are only few vertices whose open neighborhood cannot be dominated by *three* other vertices. We pick these few vertices as the set D_1, remove them from G and mark all their neighbors as dominated. Hence, after the first phase the open neighborhoods of all remaining vertices can be dominated by a constant number of other vertices.

In the second phase, we compute concurrently for each vertex v all the so-called *domination sequences* $v_1, \dots, v_s$ starting at v (see Definition 5.3 for a formal definition). The analysis of this phase is based on the construction of *pseudo-covers* as in the work of Czygrinow et al. [12] and in the approach of greedy domination in biclique-free graphs [47]. The domination sequences intuitively provide a tool to carry out a fine-grained analysis of the vertices that can potentially dominate the remaining non-dominated neighborhoods. All the vertices v_s are gathered in the set D_2 and are removed from G, with all their neighbors marked as dominated. For $K_{3,t}$-free graphs, we slightly modify the algorithm and provide an even finer analysis.

Call the number $d_R(v)$ of non-dominated neighbors of a vertex v the *residual degree* of v. We prove that after the second phase we are left with a graph where every vertex has residual degree at most Δ_R for a constant Δ_R. In particular, every vertex from a minimum dominating set of size γ can dominate at most $\Delta_R + 1$ non-dominated vertices (each vertex dominates its neighbors and itself) and we conclude that the set R of non-dominated vertices has size bounded by $(\Delta_R + 1)\gamma$. Hence, we could at this point pick all non-dominated vertices to add at most $(\Delta_R + 1)\gamma$ vertices and conclude. Instead, we study two different ways to proceed with a third phase.

Our first option for the third phase is to apply an LP-approximation based on results of Bansal and Umboh [7], who showed that a very simple selection procedure leads to a constant factor approximation when the solution to the dominating set linear program (LP) is given. As shown by Kuhn et al. [36] we can approximate such a solution in a constant number of rounds when the maximum degree Δ of the graph is bounded. To apply these results, we have to overcome two obstacles. First, note that even though we have established that the maximum residual degree is bounded by a constant Δ_R, we may still have unbounded maximum degree Δ. We overcome this problem by keeping only a few representative potential dominators around the set R of non-dominated vertices. By a simple density argument, there can be only very few high degree vertices left that we simply select into the dominating set. As a result, we are left with a graph where Δ is bounded by a constant. The second obstacle, which is easily overcome, is that we do not need to dominate the whole remaining graph but only the set R. This requires a small adaptation of the LP-formulation of the problem and a proof that the algorithm of Bansal and Umboh still works for this slightly different setting. In total, in this version of the third phase of the algorithm, we add at most $\mathcal{O}(\gamma)$ vertices.

Our second option for the third phase is to design a distributed version of the classical greedy algorithm. We proceed in a greedy manner in d rounds, as follows (where d is a bound on the maximum residual degree Δ_R of the graph after phase 2). In the first round, if a non-dominated vertex has a neighbor of residual degree d, it elects one such neighbor into the dominating set (or if it has residual degree d itself, it may choose itself). The neighbors of the chosen elements are

marked as dominated and the residual degrees are updated. Note that all non-dominated neighbors of a vertex of residual degree d in this round select a dominator. Hence, the residual degrees of all vertices of residual degree d are decreased to 0 and, after this round, there are no vertices of residual degree d left. In the second round, if a non-dominated vertex has a neighbor of residual degree $d-1$, it elects one such vertex into the dominating set, and so on, until after d rounds in the final round every vertex selects a dominator. Unlike in the general case, where nodes cannot learn the current maximum residual degree in a constant number of rounds, by establishing an upper bound on the maximum residual degree and proceeding in exactly this number of rounds, we ensure that we iteratively exactly selects the vertices of maximum residual degree. For the case of planar graphs, we prove that $\Delta_R \leq 30$. It remains to analyze the performance of this algorithm.

A simple density argument shows that there cannot be too many vertices of degree $i \geq 2\nabla_0(G)$ ($i \geq 6$ in a planar graph). At a first glance it seems that the algorithm would perform worst when in every of the d_R rounds it would pick as many vertices as possible, as the constructed dominating set would grow as much as possible. However, this is not the case, as picking many high degree vertices at the same time makes the largest progress towards dominating the whole graph. It turns out that there is a delicate balance between the vertices that we pick in round i and the remaining non-dominated vertices that leads to the worst case. For planar graphs in total, this leads to a 20-approximation. While the greedy algorithm falls short of achieving the best approximation factor, it is much simpler than the LP-based approach, and interesting to analyze in its own right.

3. Preliminaries

In this paper we study the distributed time complexity of finding dominating sets in undirected and simple graphs in the classical LOCAL model of distributed computing. We assume familiarity with this model and refer to the survey [48] for an extensive background on distributed computing and on the LOCAL model.

We use standard notation from graph theory and refer to the textbook [14] for extensive background. All graphs in this paper are undirected and simple. We write $V(G)$ for the vertex set of a graph G and $E(G)$ for its edge set. The *girth* of a graph G is the length of a shortest cycle in G. A graph is *triangle-free* if it does not contain a triangle (that is, a cycle of length three) as a subgraph. Equivalently, a triangle-free graph is a graph with girth at least four.

A graph is *bipartite* if its vertex set can be partitioned into two parts such that all its edges are incident with two vertices from different parts. We write $K_{s,t}$ for the complete bipartite graph with parts of size s and t, respectively. A set A is *independent* if no two vertices $u, v \in A$ are adjacent. We write $\alpha(G)$ for the size of the largest independent set in a graph G. The *Hall ratio* $\rho(G)$ of G is defined as $\max\{|V(H)|/\alpha(H) \mid H \subseteq G\}$. Hence, every subgraph H of G contains an independent set of size at least $|V(H)|/\rho(G)$.

A graph H is a *minor* of a graph G, written $H \preccurlyeq G$, if there is a set $\{G_v : v \in V(H)\}$ of vertex disjoint and connected subgraphs $G_v \subseteq G$ such that if $\{u, v\} \in E(H)$, then there is an edge between a vertex of G_u and a vertex of G_v. We say that $V(G_v)$ is the *branch set* of v and say that it is *contracted* to the vertex v. The set $\{G_v : v \in V(H)\}$ is a *minor model* of H. The *depth* of a minor model is the maximum radius of one of its branch sets. We call H a 1-*shallow minor*, written $H \preccurlyeq_1 G$, if $H \preccurlyeq G$ and there is a minor model $\{G_v : v \in V(H)\}$ with depth at most 1 witnessing this. In other words, $H \preccurlyeq_1 G$ if H is obtained from G by deleting some vertices and edges and then contracting a set of pairwise disjoint stars. We refer to [43] for an in-depth study of the theory of sparsity based on shallow minors.

A graph is *planar* if it can be embedded in the plane, that is, it can be drawn on the plane in such a way that its edges intersect only at their endpoints. By the famous theorem of Wagner, planar graphs can be characterized as those graphs that exclude the complete graph K_5 on five vertices and the complete bipartite $K_{3,3}$ with parts of size three as a minor. In particular, a minor of a planar graph is again planar.

A graph is *outerplanar* if it has an embedding in the plane such that all vertices belong to the unbounded face of the embedding. Equivalently, a graph is outerplanar if it does not contain the complete graph K_4 on four vertices and the complete bipartite graph $K_{2,3}$ with parts of size 2 and 3, respectively, as a minor. Again, a minor of an outerplanar graph is again outerplanar.

By Euler's formula, planar graphs are sparse: every planar n-vertex graph ($n \geq 3$) has at most $3n-6$ edges (and a graph with at most two vertices has at most one edge). The ratio $|E(G)|/|V(G)|$ is called the *edge density* of G. In particular, every planar graph G has edge density strictly smaller than three. We define

$$\nabla_0(G) = \max\left\{ \frac{|E(H)|}{|V(H)|} \;\middle|\; H \subseteq G \right\},$$
$$\nabla_1(G) = \max\left\{ \frac{|E(H)|}{|V(H)|} \;\middle|\; H \preccurlyeq_1 G \right\},$$
$$\nabla_1^B(G) = \max\left\{ \frac{|E(H)|}{|V(H)|} \;\middle|\; H \preccurlyeq_1 G, H \text{ bipartite} \right\}.$$

It is immediate that $\nabla_0(G) \leq \nabla_1^B(G) \leq \nabla_1(G)$. Note that every graph G (and each of its subgraphs) contains a vertex with degree at most $2\nabla_0(G)$. By iteratively removing a minimum degree vertex and its neighbors, we can find a large independent set, as stated in the next lemma. The bounds for graphs on surfaces are well known.

Lemma 3.1. For all graphs G we have $\rho(G) \leq 2\nabla_0(G) + 1$. Furthermore,

1. if G is planar, then $\nabla_0(G) < 3$, $\nabla_1^B(G) < 2$ and $\rho(G) \leq 4$;
2. if G is outerplanar, or planar and triangle free, then $\nabla_0(G) < 2$ and $\rho(G) \leq 3$;
3. if G is planar and bipartite, then $\nabla_0(G) < 2$ and $\rho(G) \leq 2$.

For a graph G and $v \in V(G)$ we write $N(v) = \{u : \{u, v\} \in E(G)\}$ for the *open neighborhood* of v and $N[v] = N(v) \cup \{v\}$ for the *closed neighborhood* of v. For a set $A \subseteq V(G)$ let $N[A] = \bigcup_{v \in A} N[v]$. A *dominating set* in a graph G is a set $D \subseteq V(G)$ such that $N[D] = V(G)$. We write $\gamma(G)$ for the size of a minimum dominating set of G. For a set $R \subseteq V(G)$ we say that a set $Z \subseteq V(G)$ *dominates* or *covers* R if $R \subseteq N[Z]$. For $v \in V(G)$ we let $N_R(v) = N(v) \cap R$ and $d_R(v) = |N_R(v)|$.

An *orientation* of a graph G is a directed graph $\vec{G}$ that has exactly one of the arcs (u, v) and (v, u) for each edge $\{u, v\} \in E(G)$. The *out-degree* $d^+(v)$ of a vertex v in $\vec{G}$ is the number of arcs leaving v. The following lemma is implicit in the work of Hakimi [27], see also [43, Proposition 3.3].

Lemma 3.2. Every graph G has an orientation with maximum out-degree ∇_0.

We immediately deduce the next corollary.

Corollary 3.1. Let G be a planar graph. Then

1. G has an orientation with maximum out-degree 3.
2. If G is triangle-free or outerplanar, then G has an orientation with maximum out-degree 2.

In the following we fix
- the input graph G,
- a minimum dominating set D,
- $\gamma := |D|$,
- the parameter $\nabla_1 \geq \nabla_1(G)$,

and (possibly defined in terms of ∇_1)
- the parameter $\nabla_0 \in [\nabla_0(G), \nabla_1]$,
- the parameter $\nabla \in (\nabla_1^B(G), \nabla_1 + 1]$,
- the parameters $s \leq t$ with $K_{s,t} \not\subseteq G$,

where D and γ are used only to analyze the performance of the algorithm, and where all the parameters are integers.

Above, $[a, b]$ denotes the closed real interval containing all x with $a \leq x \leq b$ and $(a, b]$ denotes the half-open interval containing all x with $a < x \leq b$. We can choose $s = t = 2\lfloor \nabla_0(G) \rfloor + 1$.

4. Phase 1: Preprocessing

4.1. Small neighborhoods dominators

As outlined in the introduction, our algorithm works in three phases. In phase i for $1 \leq i \leq 3$ we select a partial dominating set D_i and estimate its size in comparison to D. In the end we will return $D_1 \cup D_2 \cup D_3$. We will call vertices have been selected into a set D_i *green*, vertices that are dominated by a green vertex but are not green themselves are called *yellow* and all vertices that still need to be dominated are called *red*. In the beginning, all vertices are marked red.

The first phase of our algorithm is similar to the first phase of the algorithm of Lenzen et al. [38] for planar graphs. It is a preprocessing step that leaves us with only vertices whose neighborhoods can be dominated by a few other vertices. Lenzen et al. proved that if G is planar, then there exist less than 3γ many vertices v such that the open neighborhood $N(v)$ of v cannot be dominated by 6 vertices of $V(G) \setminus \{v\}$ [38, Lemma 6.3]. The lemma can be generalized to more general graphs, see [6]. We prove the following lemma, which is stronger in the sense that the number of vertices required to dominate the open neighborhoods is smaller than in [38] and [6], at the cost of having slightly more vertices with that property.

Lemma 4.1. Let $\hat{D}$ be the set of vertices $v \in V(G)$ whose neighborhood cannot be dominated by $(2\nabla - 1)$ vertices of D other than v, that is,

$$\hat{D} := \{v \in V(G) : \text{for all sets } A \subseteq D \setminus \{v\} \text{ with } N(v) \subseteq N[A] \text{ we have } |A| > (2\nabla - 1)\}.$$

Then $|\hat{D} \setminus D| < \rho(G) \cdot \gamma$.

Remember that ∇ is an integer strictly larger than $\nabla_1^B(G)$, the edge density of a densest bipartite 1-shallow minor of G. Additionally $\rho(G) \leq \chi(G) \leq 2\nabla_0(G) + 1 \leq 2\nabla_1 + 1$. The precise values will be relevant for the planar case.

Proof. Assume $D = \{b_1, \dots, b_\gamma\}$. Assume that there are $\rho(G) \cdot \gamma$ vertices $a_1, \dots, a_{\rho(G)\cdot\gamma} \notin D$ satisfying the above condition. Be definition of the Hall ratio we find an independent subset of the a_is of size γ. We can hence assume that $a_1, \dots, a_\gamma$ are not connected by an edge. We proceed towards a contradiction.

We construct a bipartite 1-shallow minor H of G with the following 2γ branch sets. For every $i \leq \gamma$ we have a branch set $A_i = \{a_i\}$ and a branch set $B_i = N[b_i] \setminus \big(\{a_1, \dots, a_\gamma\} \cup \bigcup_{j<i} N[b_j] \cup \{b_{i+1}, \dots, b_\gamma\}\big)$. Note that the B_i are vertex disjoint and hence we define proper branch sets. Intuitively, for each vertex $v \in N(a_i)$ we mark the smallest b_j that dominates v as its dominator. We then contract the vertices that mark b_j as a dominator together with b_j into a single vertex. Note that because the a_i are independent, the vertices a_i themselves are not associated to a dominator as no a_j lies in $N(a_i)$ for $i \neq j$. Denote by $a'_1, \dots, a'_\gamma, b'_1, \dots, b'_\gamma$ the associated vertices of H. Denote by A the set of the a'_is and by B the set of the b'_js. We delete all edges between vertices of B. The vertices of A are independent by construction. Hence, H is a bipartite 1-shallow minor of G. By the assumption that $N(a_i)$ cannot be dominated by $2\nabla - 1$ elements of D, we associate at least 2∇ different dominators with the vertices of $N(a_i)$. Note that this would not necessarily be true if A was not an independent set, as all $a_j \in N(a_i)$ would not be associated a dominator.

Since $\{b_1, \dots, b_\gamma\}$ is a dominating set of G and by assumption on $N(a_i)$, we have that in H, every a'_i has at least 2∇ neighbors in B. Hence, $|E(H)| \geq 2\nabla|V(A)| = 2\nabla\gamma$. As $|V(H)| = 2\gamma$ we conclude $|E(H)| \geq \nabla|V(H)|$. This however is a contradiction, as by assumption ∇ is strictly larger than $\nabla_1^B(G)$ the edge density of a densest bipartite 1-shallow minor of G. □

Let us fix the set $\hat{D}$ for our graph G.

$$\hat{D} := \{v \in V(G) : \text{ for all } A \subseteq D \setminus \{v\} \text{ with } N(v) \subseteq N[A] \text{ we have } |A| > (2\nabla - 1)\}.$$

Note that $\hat{D}$ cannot be computed by a local algorithm as we do not know the set D. It will only serve as an auxiliary set in our analysis.

We define D_1 as the set of all vertices whose neighborhood cannot be dominated by $2\nabla - 1$ other vertices. The first phase of the algorithm is to compute the set D_1, which can be done in 2 rounds of communication.

$$D_1 := \{v \in V(G) : \text{ for all } A \subseteq V(G) \setminus \{v\} \text{ with } N(v) \subseteq N[A] \text{ we have } |A| > (2\nabla - 1)\}.$$

Lemma 4.2. $D_1 \subseteq \hat{D}$ and hence $|D_1 \setminus D| \leq \rho(G) \cdot \gamma$.

Proof. If the open neighborhood of a vertex v cannot be dominated by $2\nabla - 1$ vertices from $V(G) \setminus \{v\}$, then in particular it cannot be dominated by $2\nabla - 1$ vertices from $D \setminus \{v\}$. Hence $D_1 \subseteq \hat{D}$ and we can bound the size of D_1 by that of $\hat{D}$. □

We mark the vertices of D_1 that we add to the dominating set in the first phase of the algorithm as green, the neighbors of D_1 as yellow and leave all other vertices red. Denote the set of red vertices by R, that is, $R = V(G) \setminus N[D_1]$. For $v \in V(G)$ let $N_R(v) := N(v) \cap R$ and $d_R(v) := |N_R(v)|$ be the *residual degree* of v, that is, the number of neighbors of v that still need to be dominated.

By definition of D_1, the neighborhood of every non-green vertex can be dominated by at most 2∇ other vertices. This holds true in particular for the subset $N_R(v)$ of neighbors that still need to be dominated. Let us fix such a small dominating set for the red neighborhood of every non-green vertex.

For every $v \in V(G) \setminus D_1$, we fix $A_v \subseteq V(G) \setminus \{v\}$ such that:

$$N_R(v) \subseteq N[A_v] \text{ and } |A_v| \leq 2\nabla.$$

Additionally, for vertices $v \in V(G) \setminus \hat{D}$, we enforce that $A_v \subseteq D \setminus \{v\}$.

There are potentially many such sets A_v – we fix one such set arbitrarily. Let us stress that we cannot compute these sets in a local algorithm as the sets D and $\hat{D}$ are not known to the algorithm. We only use these sets for our further argumentation.

4.2. Limitations of the method

We can apply the above approach to obtain a small set D_1 only if ∇_1 is bounded by a constant. For example in graphs of bounded degeneracy in general the number of vertices that dominate the neighborhood of a vertex can only be bounded by $\gamma(G)$. Hence, the approach based on covers and pseudo-covers that is employed in the following cannot be extended to degenerate graph classes. Below we show an example where this is the case.

Example 1. Let $G(\gamma, m)$ be the graph with vertices v_i for $1 \leq i \leq \gamma$, w^j for $1 \leq j \leq m$ and s_i^j for $1 \leq i \leq \gamma, 1 \leq j \leq m$. We have the edges $\{v_1, w^j\}$ for $1 \leq j \leq m$, hence v_1 dominates all w^j. We have the edges $\{w^j, s_i^j\}$ for all $1 \leq i \leq \gamma, 1 \leq j \leq m$, hence, the s_i^j are neighbors of w^j. Finally, we have the edges $\{v_i, s_i^j\}$, that is, v_i dominates the ith neighbor of w_j (see Example 1). Hence, for $m > \gamma$, $G(\gamma, m)$ has a dominating set of size γ and m vertices whose neighborhood can be dominated only by $\gamma(G)$ vertices. Note that $G(\gamma, m)$ is 2-degenerate. As we can choose m arbitrary large, we cannot usefully apply the method based on Lemma 4.2 for degenerate classes in general.

5. Phase 2: Reducing residual degrees – pseudo-covers and domination sequences

After the first phase of the algorithm we have established the situation that for every vertex $v \in V(G) \setminus D_1$ the residual neighborhood $N_R(v)$ is dominated by set $A_v \subseteq V(G) \setminus \{v\}$ of size at most 2∇. For all vertices of $V(G) \setminus \hat{D}$ we have chosen $A_v \subseteq D \setminus \{v\}$. Observe that the set $\bigcup_{v \in V(G)} A_v$ has very good domination properties. First, already the sets A_v for $v \in D$ dominate almost all vertices that remain to be dominated, except possibly the vertices of D itself: We have $R \subseteq \bigcup_{v \in D} N_R[v] = D \cup \bigcup_{v \in D} N_R(v) = D \cup \bigcup_{v \in D} A_v$, hence $R \setminus D \subseteq \bigcup_{v \in D} A_v$. Second, $\bigcup_{v \in V(G)} A_v$ is small, as $|\bigcup_{v \in V(G)} A_v| \leq \sum_{v \in V(G)} |A_v| = \sum_{v \in \hat{D}} |A_v| + \sum_{v \in V(G) \setminus \hat{D}} |A_v|$. So $|\bigcup_{v \in V(G)} A_v| \leq (\rho(G) + 1)\gamma \cdot 2\nabla + \gamma \in \mathcal{O}(\gamma)$.

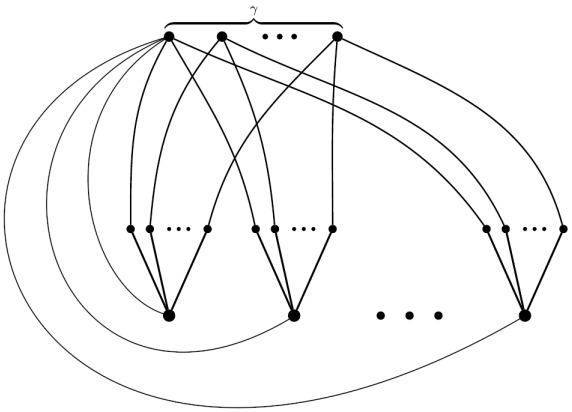

Figure 2: A 2-degenerate graph, where for many $v \in V(G)$ the set $N(v)$ can only be dominated by at least γ vertices different from v.

In the second phase of the algorithm we aim to find a good approximation of the sets A_v. We follow the approach of Czygrinow et al. [12] and define *pseudo-covers*, which describe candidate vertices for the sets A_v. We will then consider a selection process that can be carried out in parallel for all vertices, which is based on the definition of *domination sequences*, and allows to select a bounded number of candidate vertices. The domination properties of the selected vertices are worse than that of the sets A_v, however, at the end of the second phase we will be in the situation that the residual degree of each vertex is bounded by an absolute constant depending only on the graph class under consideration.

5.1. Pseudo-covers

Following the presentation of [12], we name and fix the following constants for the rest of this article. The reason to choose the constants as given will become clear in the course of the proof.

$$\begin{aligned} \kappa &:= \max\{2\nabla_0, 2\nabla\}, & \mu &:= 2\kappa/\lambda = 2\kappa^2, \\ \lambda &:= 1/\kappa, & \nu &:= k\mu = 2\kappa^3. \end{aligned}$$

Definition 5.1. A vertex $z \in V(G)$ is *λ-strong* for a vertex set $W \subseteq V(G)$ if $|N[z] \cap W| \geq \lambda |W|$.

The following is the key definition by Czygrinow et al. [12].

Definition 5.2. A *pseudo-cover* (with parameters $\kappa, \lambda, \mu, \nu$) of a set $W \subseteq V(G)$ is a sequence $(v_1, \ldots, v_m)$ of vertices such that for every $i \leq m$ we have:

- $m \leq \kappa$,
- v_i is λ-strong for $W \setminus \bigcup_{j<i} N[v_j]$,
- $|N[v_i] \cap (W \setminus \bigcup_{j<i} N[v_j])| \geq \mu$, and
- $|W \setminus \bigcup_{j\leq m} N[v_j]| \leq \nu$.

Intuitively, all but at most ν elements of the set W are covered by the $(v_i)_{i\leq m}$. Additionally, each element v_i of the pseudo-cover dominates both an λ-fraction of the part of W that is not yet dominated by the v_j for $j < i$, and at least μ elements. Note that with our choice of constants, if there are more than ν vertices not covered yet, any vertex that covers a λ-fraction of what remains also covers at least μ elements.

The next lemma shows how to derive the existence of pseudo-covers from the existence of small dominating sets.

> *Lemma* 5.1. Let $W \subseteq V(G)$ be of size at least ν and let Z be a dominating set of W with κ elements. There exists an ordering of the vertices of Z as $z_1, \dots, z_\kappa$ and $m \leq \kappa$ such that $(z_1, \dots, z_m)$ is a pseudo-cover of W.

Proof. We build the order greedily by induction. We order the elements by neighborhood size, while removing the neighborhoods of the previously ordered vertices. More precisely, assume that $(z_1, \dots, z_i)$ have been defined for some $i \geq 0$. We then define z_{i+1} as the element that maximizes $|N[z] \cap (W \setminus \bigcup_{j\leq i} N[z_j])|$.

Once we have ordered all vertices of Z, we define m as the maximal integer not larger than κ such that for every $i \leq m$ we have:

- z_i is λ-strong for $W \setminus \bigcup_{j<i} N[z_j]$, and
- $|N[z_i] \cap (W \setminus \bigcup_{j\leq i} N[z_j])| \geq \mu$.

This ensures that $(z_1, \dots, z_m)$ satisfies the first 3 properties of a pseudo-cover of W. It only remains to check the last one. To do so, we define $W' \coloneqq W \setminus \bigcup_{i\leq m} N[z_i]$. We want to prove that $|W'| \leq \nu$. Note that because Z covers W, if $m = \kappa$ we have $W' = \emptyset$ and we are done. We can therefore assume that $m < \kappa$ and $W' \neq \emptyset$. Since Z is a dominating set of W, we also know that $(z_{m+1}, \dots z_\kappa)$ is a dominating set of W', therefore there is an element in $(z_{m+1}, \dots z_\kappa)$ that dominates at least a $1/\kappa$ fraction of W'. Thanks to the previously defined order, we know that z_{m+1} is such an element. Since $\lambda = 1/\kappa$, it follows that z_{m+1} is λ-strong for W'. This, together with the definition of m, we have that $|N[z_i] \cap (W \setminus \bigcup_{j\leq i} N[z_j])| < \mu$ meaning that $|N[z_{m+1}] \cap W'| < \mu$. This implies that $|W'|/\kappa < \mu$. And since $\mu = \nu/\kappa$, we have $|W'| < \nu$. Hence, $(z_1, \dots, z_m)$ is a pseudo-cover of W. □

While there can exist unboundedly many dominating sets for a set $W \subseteq V(G)$, a nice observation of Czygrinow et al. was that the number of pseudo-covers is bounded whenever the input graph excludes some biclique $K_{s,t}$ as a subgraph. We do not state the result in this generality, as it leads to enormous constants. Instead, we focus on the case where small dominating sets A_v exist, implying that $\nabla_0(G)$, and therefore κ, are bounded.

> *Lemma* 5.2. Let $W \subseteq V(G)$ of size at least μ. Then there are less than κ^2 vertices that are λ-strong for W.

Proof. Assume that there is such a set W with $|W| \geq \mu$ and c many vertices that are λ-strong for W. Let H be the subgraph of G induced by W and the λ-strong vertices. We first have that $|V_H| \leq |W| + c$. Second we have that $|E_H| \geq \lambda|W|c - \nabla_0 c$, because there are c vertices that have degree at least $\lambda|W|$ and there are at most $\nabla_0 c$ many vertices between them.

We then have that $|E_H| \leq \nabla_0 |V_H|$ hence $\lambda|W|c - \nabla_0 c \leq \nabla_0(|W| + c)$ from which we derive $c(\lambda|W| - 2\nabla_0) \leq \nabla_0|W|$. Now, using that $|W| > \mu \geq 2k^2$ we have $\lambda|W| > 2k > 4\nabla_0$. Hence $\lambda|W| - 2\nabla_0 \geq \lambda|W|/2$. We can finally deduce that $c(\lambda|W|/2) \leq \nabla_0|W|$ and therefore we have that $c \leq 2\nabla_0/\lambda = \kappa^2$. □

This leads quickly to a bound on the number of pseudo-covers.

Lemma 5.3. For every $W \subseteq V(G)$ of size at least μ, the number of pseudo-covers is less than $\kappa^{2\kappa}$.

The proof of the lemma is exactly as the proof of Lemma 7 in the presentation of Czygrinow et al. [12], we reprove it for the sake of completeness.

Proof. Let W a set of size at least μ. For every $i \leq \kappa$, we define C_i as the set of partial pseudo-covers of W of size at most i, that is, all sets of at most i vertices that can be extended to a pseudo-cover of W. So C_κ is the set of pseudo-covers of W while C_1 only contains λ-strong vertices for W.

Lemma 5.2 implies that $|C_1| < \kappa^2$. Lemma 5.2 also implies that for every $i < \kappa$, we have $|C_{i+1}| < |C_i| \cdot \kappa^2$. We therefore conclude that $|C_k| < (\kappa^2)^\kappa$. □

We write $\mathcal{T}(v)$ for the set of all pseudo-covers of $N(v)$ and $\mathcal{P}(v)$ for the set of all vertices that appear in a pseudo-cover of $N(v)$.

The proof of Lemma 5.3 also bounds the number $\mathcal{P}(v)$.

Corollary 5.1. For every $v \in V(G)$ with $|N_R(v)| > \mu$, we have $|\mathcal{P}(v)| \leq \kappa^{2\kappa}$.

5.2. Domination sequences

We now turn to the use of pseudo-covers. We aim to carry out an iterative process in parallel for all vertices $v \in V(G)$ with a sufficiently large residual neighborhood $N_R(v)$.

Definition 5.3. For any vertex $v \in V(G)$, a *κ-dominating-sequence* of v is a sequence $(v_1, \dots, v_m)$ (without repetition) for which we can define sets $B_1, \dots, B_m$ such that:

- $v_1 = v$, $B_1 \subseteq N_R(v_1)$,
- for every $i \leq m$ we have $B_i \subseteq (N_R(v_i) \cap B_{i-1})$,
- $|B_i| \geq \kappa^{s-i}(t + s - i + (s - i)\nu)$
- and for every $i \leq m$ we have $v_i \in \mathcal{P}(v_{i-1})$.

A κ-dominating-sequence $(v_1, \dots, v_m)$ is *maximal* if there is no vertex u such that $(v_1, \dots, v_m, u)$ is a κ-dominating-sequence.

Note that this definition requires $|N_R(v)| \geq \kappa^{s-1}(t + s - 1 + (s - 1)\nu)$. For a vertex v with a too small residual neighborhood, there are no κ-dominating-sequences. We show two main properties of these dominating-sequences. First, Lemma 5.4 shows that a maximal dominating sequence must encounter $D \cup \hat{D}$ at some point. Second, with Lemmas 5.5 to 5.7, we show that collecting all "end points" v_m of all maximal dominating sequences results in a set D_2 of size linear in the size of D. While D cannot be computed, we can compute D_2.

Lemma 5.4. Let v be a vertex and let $(v_1, \dots, v_m)$ be a maximal κ-dominating-sequence of v. Then $m < s$ and $(D \cup \hat{D}) \cap \{v_1, \dots, v_m\} \neq \emptyset$.

Proof. First, assume that $v_1, v_2, \dots, v_m$ is a maximal κ-dominating-sequence with $m \geq s$. By definition, every v_i with $i \leq s$ is connected to every vertex of B_s. For every $1 \leq i \leq s$ we have $|B_i| \geq t$ and therefore $|B_s| \geq t$. This shows that the two sets $\{v_1, \dots, v_s\}$ and B_s form a $K_{s,t}$ as a subgraph in $N^2[v]$. Since $K_{s,t}$ is excluded as a subgraph in G, the process must stop having performed at most $s - 1$ rounds.

We now have $m < s$ and to prove the second statement we assume, in order to reach a contradiction, that $(D \cup \hat{D}) \cap \{v_1, \dots, v_m\} = \emptyset$. We have that $B_m \subseteq N(v_m)$, and remember that as v_m is not in $\hat{D}$, we have that $N_R(v_m)$ can be dominated by at most κ elements of D.

By Lemma 5.1, we can derive a pseudo-cover $S = (u_1, \dots, u_j)$ of $N_R(v_m)$, where $j \leq \kappa$ and every u_i is an element of D. Let X denote the set (of size at most ν) of vertices not covered by S. As S contains at most κ vertices there must exist a vertex u in S that covers at least a $1/\kappa$ fraction of $B_m \setminus X$. By construction, we have that $|B_m| \geq \kappa^{s-m} \cdot (t + s - m + (s - m)\nu) \geq \kappa(t + \nu)$ because $m < s$. Therefore $|B_m \setminus X| \geq \kappa$ and we have

$$|N_R[u] \cap B_m| \geq \frac{|B_m| - \nu}{\kappa} \geq \frac{\kappa^{s-m}(t + s - m + (s - m)\nu) - \nu}{\kappa},$$

hence

$$|N_R[u] \cap B_m| \geq \frac{\kappa^{s-m}(t + s - m + (t - m - 1)\nu)}{\kappa} \geq \kappa^{s-m-1}(t + s - m + (t - m - 1)\nu),$$

and therefore

$$|N_R(u) \cap B_m| \geq |N_R[u] \cap B_m| - 1 \geq \kappa^{s-m-1}(t + s - m - 1 + (t - m - 1)\nu).$$

So we can continue the sequence $(v_1, \dots, v_m)$ by defining $v_{m+1} := u$; there is no repetition since by hypothesis $D \cap \{v_1, \dots, v_m\} = \emptyset$, and by construction $u \in D$.

In conclusion if $(v_1, \dots, v_m)$ is a maximal sequence, it contains an element of D or $\hat{D}$. □

Our next goal is to show that there are not many elements v_m (which are the elements that we pick into the set D_2).

Lemma 5.5. For any maximal κ-dominating-sequence $(v_1, \dots, v_m)$, and for any $i \leq m - 1$, we have that

- $v_{i+1} \in \mathcal{P}(v_i)$, and
- $|N_R(v_i)| \geq \mu$.

Proof. By construction we have $v_{i+1} \in \mathcal{P}(v_i)$, furthermore $|B_i| \geq \kappa^{t-i}(2t - i + (t - i)\nu) \geq \nu > \mu$, and $B_i \subseteq N_R(v_i)$. □

Now, for every $v \in V(G)$ we compute all maximal κ-dominating-sequences starting with v. Obviously, as every v_i in any κ-dominating-sequences of v dominates some neighbors of G, we can locally compute these steps after having learned the 2-neighborhood $N^2[v]$ of every vertex in two rounds in the LOCAL model of computation.

We define D_2 as the set of all $u \in V(G)$ such that there is some vertex $v \in V(G)$, and some maximal κ-dominating-sequence $(v_1, \dots, v_m)$ of v with $u = v_m$.

We now take a look at the size of D_2. For a set $W \subseteq V(G)$ we write $\mathcal{P}(W) = \bigcup_{v \in W} \mathcal{P}(v)$. Remember that the definition of $\mathcal{P}(v)$ requires that $|N_R(v)| > \mu$. We simply extend the notation with $\mathcal{P}(v) = \emptyset$ if $|N_R(v)| \leq \mu$. We then define:

$$\mathcal{P}^{(1)}(W) := \mathcal{P}(W)$$

for $1 < i < s$

$$\mathcal{P}^{(i)}(W) := \mathcal{P}(\mathcal{P}^{(i-1)}(W))$$

and, for $1 \leq i \leq s$

$$\mathcal{P}^{(\leq i)}(W) := \bigcup_{1 \leq j \leq i} \mathcal{P}^{(j)}(W).$$

We are now ready to prove that D_2 is small.

Lemma 5.6. $D_2 \subseteq \mathcal{P}^{(\leq s)}(D \cup \hat{D})$.

Proof. Using Lemma 5.5 repetitively, for every κ-dominating-sequence $(v_1, \dots, v_m)$ we have that $v_m \in \mathcal{P}^{(\leq s)}(v_1)$, and, more generally, for every $i \leq m$, we have that $v_m \in \mathcal{P}^{(\leq s)}(v_i)$. Now the statement follows from Lemma 5.4. □

Lemma 5.7. $|D_2| \leq (\kappa^{2s\kappa}(\rho(G) + 1)\gamma$.

Proof. Corollary 5.1 gives us that $|\mathcal{P}(v)| \leq \kappa^{2\kappa}$ for every $v \in V(G)$ with $|N(v)| > \mu$. As $\mathcal{P}(W) \leq \sum_{v \in W} |\mathcal{P}(v)|$, we have $P(W) \leq |W| \cdot \kappa^{2\kappa}$. A simple induction yields that for $i \leq t$,

$$|\mathcal{P}^{(\leq i)}(W)| \leq c^i |W|,$$

where $c = \kappa^{2\kappa}$. With Lemma 5.6 we conclude

$$|D_2| \leq \kappa^{2s\kappa} \cdot |D \cup \hat{D}|.$$

We conclude with Lemma 4.1, stating that $|\hat{D} \setminus D| \leq \rho(G) \cdot \gamma$. □

We update the set R of vertices that still need to be dominated as $V(G) \setminus N[D_1 \cup D_2]$ and the residual neighborhoods $N_R(v) = N(v) \cap R$ and residual degrees $d_R(v) = |N_R(v)|$. We prove next that $d_R(v)$ is bounded by a constant.

Lemma 5.8. For every vertex $v \in V(G)$ we have $d_R(v) < \kappa^{s-1}(t + s - 1 + (s-1)\nu)$.

Proof. Assume, for the sake of reaching a contradiction, that there is a vertex v with $d_R(v) \geq \kappa^{s-1}(t + s - 1 + (s-1)\nu)$ and let $B_1 := N_R(v)$. Note that $v \notin D_1 \cup D_2$, as the residual degree of vertices from this set is 0. Exactly as in the proof of Lemma 5.4, since $v \notin D_1$, we have that B_1 can be dominated by at most κ elements. Hence by Lemma 5.1, we can derive a pseudo-cover $S = (u_1, \dots, u_j)$ of B_1, where $j \leq \kappa$. This leads to the existence of some vertex u in S that covers at least a $1/\kappa$ fraction of $B_1 \setminus X$ for some X of size at most ν. This yields a vertex v_2, and a set B_2.

We can then continue and build a maximal k-dominating-sequence $(v_1, \dots v_m)$ of v. By construction, this sequence has the property that every v_i dominates some elements of B_1. This is true in particular for v_m, but also we have that $v_m \in D_2$, hence a contradiction. □

Let $\Delta_R := \kappa^{s-1}(t + s - 1 + (s-1)\kappa^3)$.

As it remains to dominate the set R, let us fix a minimum dominating set D_R of size γ_R for R.

- Let $D_R \subseteq V(G)$ be a minimum dominating set of R and let $\gamma_R := |D_R|$.
- Let $\eta \in [0, 1]$ be such that $|(D_1 \cup D_2) \cap D| = \eta\gamma$.

Lemma 5.9. $\gamma_R \leq (1 - \eta)\gamma$.

Proof. $D \setminus (D_1 \cup D_2)$ is a dominating set for R, hence $|D_R| \leq |D \setminus (D_1 \cup D_2)|$. □

As every vertex of D_R can dominate at most $\Delta_R + 1$ vertices (its Δ_R residual neighbors and itself), we have the following corollary.

Corollary 5.2. $|R| \leq (\Delta_R + 1)\gamma_R$.

6. Phase 3: LP-based approximation in graphs of bounded maximum degree

6.1. LP-based approximation

In the light of Corollary 5.2, we could now simply choose R as the set D_3 to get a constant factor approximation. We can improve the bounds however, by proceeding with an LP-based approximation. The dominating set problem can be formulated as an integer linear program (ILP). Note that it remains to dominate the set R, which leads to the following ILP.

$$\begin{array}{lll} \text{Minimize} & \sum_{v \in V} x_v & \\ \text{Subject to} & \sum_{u \in N[v]} x_u \geq 1 & \forall v \in R \\ & x_v \in \{0, 1\} & \forall v \in V \end{array}$$

By relaxing the condition that $x_v \in \{0, 1\}$ to $x_v \in [0, 1] \subseteq \mathbb{R}$, we obtain the corresponding linear program (LP). By a result of Bansal and Umboh [7] one can obtain a constant factor approximation of a dominating set from a solution to the LP. The proof can easily be adapted to the problem of approximating a dominating set of the set R.

Lemma 6.1. Assume G has an orientation with maximum out-degree d. Let $\left(x_v\right)_{v \in V(G)}$ be a solution to the R-dominating set LP. Let $H := \{v \in V(G) : x_v \geq 1/(2d + 1)\}$ and let $U := \{v \in R : v \notin N[H]\}$. Then $H \cup U$ dominates R and has size at most $(2d + 1) \cdot \gamma_R$.

Observe that when given the solution $\left(x_v\right)_{v \in V(G)}$ to the R-dominating set LP the lemma gives rise to a simple LOCAL algorithm. First select all vertices v with $x_v \geq 1/(2d + 1)$ into a dominating set and mark all their neighbors as dominated. Then select all non-dominated vertices

of R into the dominating set. Clearly, $H \cup U$ is a dominating set of R. The rest of this section is devoted to the proof of the claimed approximation factor. The proof follows the presentation of Bansal and Umboh [7] with the improved bounds of Dvořák [18] (presented in Lemma 6.1). As every solution to the ILP is also a solution to the LP we have $\sum_{v \in V(G)} x_v \leq \gamma_R$.

Consider an orientation of G such that the neighborhood of each vertex v is decomposed into $N^{in}(v)$ and $N^{out}(v)$, where $|N^{out}(v)| \leq d$.

▷ *Claim* 1. For every vertex $v \in U$, we have $\left(\sum_{u \in N^{in}(v)} x_u \right) \geq d/(2d+1)$.

Proof. As v is not in H, $x_v < 1/(2d+1)$. As v is not in $N(H)$, for every vertex $u \in N^{out}(v)$ we have $x_u < 1/(2d+1)$. As $|N^{out}(v)| \leq d$, and by the first LP condition $\left(\sum_{u \in N^{in}(v)} x_u \right) \geq 1 - \frac{1}{2d+1} - \frac{d}{2d+1} \geq \frac{d}{2d+1}$. □

We can now bound the size of U and H

▷ *Claim* 2. $|H \cup U| \leq (2d+1) \sum_{v \in V} x_v$.

Proof. First, observe that $|H| \leq (2d+1) \sum_{v \in H} \frac{1}{2d+1} \leq (2d+1) \sum_{v \in H}(x_v)$. Then observe that $|U| \leq \frac{2d+1}{d} \cdot \sum_{v \in U} \frac{d}{2d+1} \leq \frac{2d+1}{d} \sum_{v \in U} \sum_{u \in N^{in}(v)} x_u \leq \frac{2d+1}{d} \sum_{u \in N^{in}(U)}(d \cdot x_u) \leq (2d+1) \sum_{u \in N^{in}(U)} x_u$.

By definition of U, we have that $N(U)$ and H are disjoint, this also holds for H and $N^{in}(U)$, hence $|H \cup U| \leq (2d+1) \sum_{v \in V} x_v \leq (2d+1)\gamma_R$. □

6.2. *Solving LPs locally*

As shown by Kuhn et al. [36] we can locally approximate general covering LPs, in particular the above R-dominating set LP, when the maximum degree of the graph is bounded. More precisely, they show how to compute a $\Delta^{1/r}$-approximation in $\mathcal{O}(r^2)$ rounds. Assuming for a moment that Δ is bounded by an absolute constant we can choose r such that $\Delta^{1/r} = 1 + \varepsilon$, hence $r = (\log \Delta)/(\log(1+\varepsilon))$, which is a constant depending only on Δ and ε in order to compute a $(1+\varepsilon)$-approximation for the R-dominating set LP.

> *Corollary* 6.1. Assume G has an orientation with maximum out-degree d. For every $\varepsilon > 0$ we can compute a set D' of size at most $(2d+1)(1+\varepsilon)\gamma_R$ that dominates R in $\mathcal{O}(\log \Delta/(\log(1+\varepsilon))$ rounds in the LOCAL model.

6.3. *From bounded residual degree to bounded degree*

It remains to establish the situation that the maximum degree Δ of our graph is bounded. As argued, we have $|R| \leq (\Delta_R + 1)\gamma_R$. As only the vertices of R need to be dominated it suffices to keep only the vertices that have a neighbor in R; other vertices are not useful as dominators. Also, when two vertices $u, v \in V(G) \setminus R$ have exactly the same neighbors in R, that is, $N_R(u) = N_R(v)$, it suffices to keep one of u and v. Note that we can locally decide whether $N_R(u) = N_R(v)$. For every set $N \subseteq R$ such that there is a vertex v with $N_R(v) = N$ we choose the one with the lowest identifier as a representative. We construct the graph G' consisting of R and all edges between vertices in R as well as the set of all representatives and a minimal set of edges such that $N_R(v)$ is equal in G and G' for all representatives v. Hence in G' we have $N_R(u) \neq N_R(v)$ for all $u \neq v \in V(G') \setminus R$. As argued above, every R-dominating set in G can be transformed into an R-dominating set of the same size in G' (by choosing appropriate representatives) and every

R-dominating set in G' is an R-dominating set in G. We can hence continue to work with the graph G'. In order to avoid complicated notation we simply assume that $G = G'$.

Note that in general we could have $|V(G)| \in \Omega(2^{|R|})$. When $\nabla_1(G)$ bounded, however, it follows from Lemma 4.3 of [23] that $|V(G)| \leq (4^{\nabla_1} + 2\nabla_1)|R|$, which is is linear in $|R|$. This is crucial for our further argumentation.

Corollary 6.2. $|V(G)| \leq (4^{\nabla_1} + 2\nabla_1)(\Delta_R + 1)\gamma_R$.

6.4. Conclusion of the algorithm

Given any $\varepsilon > 0$ we now select all vertices with high degree $\Gamma = \Gamma(\varepsilon)$ into our dominating set, where Γ is chosen such that there exist at most $\varepsilon\gamma$ vertices of degree at least Γ.

Let $\Gamma := 4\nabla_1(4^{\nabla_1} + 2\nabla_1)(\Delta_R + 1)/\varepsilon$ and $D_3^1 := \{v \in V(G) \ : \ d(v) > \Gamma\}$.

Lemma 6.2. $|D_3^1| \leq (\varepsilon/2)\gamma_R$.

Proof. We assume the opposite and count the number of edges of G. When we sum the degree of the vertices, we get twice the number of edges. Hence $2 \cdot |E(G)| > 2\nabla_1(4^{\nabla_1} + 2\nabla_1)(\Delta_R + 1)\gamma_R$. Therefore, with Corollary 6.2, $|E(G)| > \nabla_1|V(G)|$, a contradiction. □

After picking D_3^1 into the dominating set, marking the neighbors of D_3^1 as dominated and updating the set R, we can delete the vertices of D_3^1. We are left with a graph of maximum degree Γ.

Given $\varepsilon > 0$, let D_3^2 be the set computed by the LOCAL algorithm of Corollary 6.1 with parameter $\varepsilon/2$.

Let $D_3 = D_3^1 \cup D_3^2$. We already noted that the definition of D_3 implies that $D_1 \cup D_2 \cup D_3$ is a dominating set of G. We now conclude the analysis of the size of this computed set.

Lemma 6.3. We have that $|D_3| \leq (2\nabla_0 + 1)(1 + \varepsilon)\gamma_R$.

Proof. By Lemma 3.2 G has an orientation with out-degree $d \leq \nabla_0$. By Corollary 6.1 and Lemma 6.2 we have $|D_3^2| \leq (2\nabla_0 + 1)(1 + \varepsilon/2)\gamma_R$, and $|D_3^1| \leq (\varepsilon/2)\gamma_R$. □

Now our main theorem, Theorem 2.1, follows by summing the sizes of D_1, D_2 and D_3.

Lemma 6.4. $|D_1 \cup D_2 \cup D_3| \leq 2(\nabla_0 + 1)(\kappa^{2s\kappa} + 2)\gamma$. Hence, putting $c = 2(\nabla_1 + 1)$, we have

$$|D_1 \cup D_2 \cup D_3| \leq \left(c^{2c^2} + c\right)\gamma.$$

Proof. We have

$$\begin{aligned}
&\rho(G) \leq 2\nabla_0(G) + 1 && \text{(by Lemma 3.1)}\\
&|D_1| \leq \rho(G)\gamma \leq (2\nabla_0 + 1)\gamma && \text{(by Lemma 4.2)}\\
&|D_2| \leq \kappa^{2s\kappa}(\rho(G) + 1)\gamma \leq \kappa^{2s\kappa}(2\nabla_0 + 2)\gamma && \text{(by Lemma 5.7)}
\end{aligned}$$

Last, by setting $\epsilon = 1$ in Lemma 6.3 we have

$$|D_3| \leq (2\nabla_0 + 1)(1 + \epsilon)\gamma_R \leq 2(2\nabla_0 + 1)\gamma.$$

We conclude, as $\kappa \leq 2\nabla_1 + 2$ and $s \leq 2\nabla_1 + 1$.

□

7. Alternative Phase 3: Greedy domination

We now consider an alternative approach for the third phase, which does not improve the approximation factor, however, is conceptually much simpler and interesting in its own. Recall that we bounded Δ_R as $\kappa^{s-1}(t + s - 1 + (s - 1)\nu)$, which is a bound on the residual degree $d_R(v)$ of all vertices.

We simulate the classical greedy algorithm, which in each round selects a vertex of maximum residual degree. Here, we let all non-dominated vertices that have a neighbor of maximum residual degree choose such a neighbor as its dominator (or if they have maximum residual degree themselves, they may choose themselves). In general this is not possible for a LOCAL algorithm, however, as we established a bound on the maximum degree we can proceed as follows. We let $i = \Delta_R$. Every red vertex that has at least one neighbor of residual degree Δ_R arbitrarily picks one of them and elects it to the dominating set. Then every vertex recomputes its residual degree and i is set to $\Delta_R - 1$. We continue until i reaches 0 when all vertices are dominated. More formally, we define several sets as follows.

For $\Delta_R \geq i \geq 0$, for every $v \in R$ in parallel:

If there is some u with $d_R(u) = i$ and ($\{u, v\} \in E(G)$ or $u = v$), then
$\mathrm{dom}_i(v) \leftarrow \{u\}$ (pick one such u arbitrarily),
$\mathrm{dom}_i(v) \leftarrow \emptyset$ otherwise.

- $R_i \leftarrow R$ — *What currently remains to be dominated*
- $P_i \leftarrow \bigcup_{v \in R} \mathrm{dom}_i(v)$ — *What we pick in this step*
- $R \leftarrow R \setminus N[P_i]$ — *Update red vertices*

Last, $D_3 \leftarrow \bigcup_{1 \leq i \leq d} P_i$.

Let us first prove that the algorithm in fact computes a dominating set.

Lemma 7.1. When the algorithm has finished the iteration with parameter $i \geq 1$, then all vertices have residual degree at most $i - 1$.

In particular, after finishing the iteration with parameter 1, there is no vertex with residual degree 1 left and in the final round all non-dominated vertices choose themselves into the dominating set. Hence, the algorithm computes a dominating set of G.

Proof. By induction, before the iteration with parameter i, all vertices have residual at most i. Assume v has residual degree i before the iteration with parameter i. In that iteration, all non-dominated neighbors of v choose a dominator (possibly v, then the statement is trivial), hence, are

removed from R. It follows that the residual degree of v after the iteration is 0. Hence, after this iteration and before the iteration with parameter $i-1$, we are left with vertices of residual degree at most $i-1$. □

For the rest of this section analyze the size of D_3 and we prove the following lemma.

Lemma 7.2. We have

$$|D_3| \leq \left(\nabla_0 \ln\left(\frac{2\Delta_R - 4\nabla_0 + 1}{2\nabla_0 + 1}\right) + 3\nabla_0 + 1\right)\gamma_R.$$

Towards establishing the lemma we analyze the sizes of the sets P_i and R_i. The next lemma follows from the fact that every vertex chooses at most one dominator.

Lemma 7.3. For every $0 \leq i \leq \Delta_R$, $\sum_{j \leq i} |P_j| \leq |R_i|$.

Proof. The vertices of R_i are those that remain to be dominated in the last i rounds of the algorithm. As every vertex that remains to be dominated chooses at most one dominator in one of the rounds $j \leq i$, the statement follows. □

As the vertices of D_R that still dominate non-dominated vertices also have bounded residual degree, we can conclude that not too many vertices remain to be dominated.

Lemma 7.4. For every $0 \leq i \leq \Delta_R$, $|R_i| \leq (i+1)\gamma_R$.

Proof. First note that for every i, $D_R \setminus \bigcup_{j>i} P_j$ is a dominating set for R_i; additionally each vertex in this set has residual degree at most i. As every vertex dominates its residual neighbors and itself, we conclude $|R_i| \leq (i+1)\gamma_R$. □

Finally, we show that we cannot pick too many vertices of high residual degree. This follows from a simply density argument. bounded edge density.

Lemma 7.5. For every $2\nabla_0 < i \leq \Delta_R$, $|P_i| \leq \frac{\nabla_0}{i-2\nabla_0}(|R_i| - |R_{i-1}|)$.

Proof. Let $2\nabla_0 < i \leq \Delta_R$ be an integer. We bound the size of P_i by a counting argument, using that G (as well as each of its subgraphs) have edge density at most ∇_0.

Let $J := G[P_i]$ be the subgraph of G induced by the vertices of P_i, which all have residual degree i. Let $K := G[P_i \cup (N[P_i] \cap R_i)]$ be the subgraph of G induced by the vertices of P_i together with the red neighbors that these vertices dominate.

We have $|E(J)| \leq \nabla_0 |V(J)| = \nabla_0 |P_i|$. As every vertex of J has residual degree exactly i, we get $|E(K)| \geq iP_i - |E(J)| \geq (i - \nabla_0)|P_i|$ (we have to subtract $|E(J)|$ to not count twice the edges of K that are between two vertices of J). We also have $|V(K)| \leq |V(J)| + |N[P_i] \cap R_i|$ and $|E(K)| \leq \nabla_0 |V_K|$, hence $(i - 2\nabla_0)|P_i| \leq \nabla_0 |N[P_i] \cap R_i)|$. Now, as $R_{i-1} = R_i \setminus N[P_i]$, that is, $N[P_i] \cap R_i = R_i \setminus R_{i-1}$, we get $|P_i| \leq \frac{\nabla_0}{i-2\nabla_0}(|R_i| - |R_{i-1}|)$. □

Let $r_i := |R_i|/\gamma_R$ and $d_i := |P_i|/\gamma_R$. Our goal is to maximize $S := \sum_{0 \le i \le \Delta_R} d_i$ (which we have to multiply by γ_R in the end) under the constraints $d_i \ge 0$ and

$$r_i \ge \sum_{j \le i} d_j \qquad (0 \le i \le \Delta_R) \tag{1}$$

$$r_i \le i + 1 \qquad (0 \le i \le \Delta_R) \tag{2}$$

$$d_i \le \frac{\nabla_0}{i - 2\nabla_0} (r_i - r_{i-1}) \qquad (2\nabla_0 < i \le \Delta_R) \tag{3}$$

We may assume that $\Delta_R \ge 3\nabla_0$, as otherwise, Lemma 7.2 follows immediately from Lemma 7.3.

Let a be the minimum integer such that $d_a > 0$.

Lemma 7.6. We can assume $r_i = 0$ for all $i < a$.

Proof. Putting $r_i = 0$ for all $i < a$ obviously preserves Equations (1) and (2). It also preserves Equation (3) as the only case to check is $i = a - 1$ (if $a \ge 2\nabla_0$), for which the right hand side was possibly increased. □

Lemma 7.7. If $a \le 3\nabla_0 - 1$, then decreasing d_a to 0 and r_a to $r_a - d_a$ and increasing d_{a+1} to $d_a + d_{a+1}$ preserves all the constraints and the value of S.

Proof. The sum in Equation (1) does not change if $i > a$ and Equation (1) is obviously satisfied after modifications for $i \le a$. Equation (2) is trivially satisfied after modification, as no r_i increases. The only changes for Equation (3) correspond to the case $i = a - 1$ (for which the left hand side decreases, while the right hand side increases) or to the case $i = a$ (for which the left hand side increases by d_a, while the right hand side increases by $\nabla_0/(a + 1 - 2\nabla_0)\, d_a \ge d_a$). □

From the above lemmas, as $r_a \ge d_a$, it follows that we may assume $a \ge 3\nabla_0$ and $r_i = 0$ for all $i < a$.

Note that Equation (3) implies

$$r_{2\nabla_0} \le r_{2\nabla_0+1} \le \cdots \le r_{\Delta_R}. \tag{4}$$

Remark that increasing r_{Δ_R} obviously preserves Equations (1) and (3). Hence, we can assume that $r_{\Delta_R} = \Delta_R + 1$. Let b be minimum with $r_i = i + 1$ for all $i \ge b$. Note that $b \ge a$.

Lemma 7.8. Let $\alpha = \min(b - r_{b-1}, \sum_{j<b-1} d_j)$. If $b \ge 3\nabla_0 + 1$, then increasing d_{b-1} and r_{b-1} by α and decreasing $\sum_{j<b-1}$ by α preserves the constraints and the value of S.

Proof. Equations (1) and (2) are obviously preserved. For Equation (3) we have to check the case where $i = b - 1$ (for which the right hand side decreases by $\nabla_0/(b - 1 - 2\nabla_0)\, \alpha \le \alpha$ and the left hand side decreases by α) and the case $i = b - 2$ (for which the right hand side increases and the left hand side decreases). □

Applying this lemma, either we can reduce b to $3\nabla_0$ (hence $b = a$), or we force $d_i = 0$ for all $i < b - 1$. Thus, $a = b - 1$ or $a = b$.

Lemma 7.9. We can assume that for every $b < i \le \Delta_R$ we have $d_i = \nabla_0/(i - 2\nabla_0)$.

Proof. Indeed, as $b \geq a \geq 3\nabla_0$, for $b < i \leq \Delta_R$, Equations (1) to (3) reduce to $d_i \leq \frac{\nabla_0}{i-2\nabla_0}$. Hence, we can assume $d_i = \nabla_0/(i-2\nabla_0)$ if $i > b$. □

Lemma 7.10. We can assume $b = a$.

Proof. Assume $a = b-1$ and let $\alpha = b - r_a$. We have $d_a \leq \frac{\nabla_0}{a-2\nabla_0}(b-\alpha)$ and $d_b \leq \frac{\nabla_0}{b-2\nabla_0}(1+\alpha)$. If we increase r_a to b, we can increase d_a by $\frac{\nabla_0\alpha}{a-2\nabla_0}$ and decrease d_b by $\frac{\nabla_0\alpha}{b-2\nabla_0}$, while preserving the constraints and increasing S. □

Lemma 7.11. We can assume $a = 3\nabla_0$.

Proof. Assume $a \geq 3\nabla_0 + 1$. By putting $r_{a-1} = a$ we can increase d_{a-1} by $\frac{\nabla_0}{a-1-2\nabla_0}a$ and decrease d_a by $\frac{\nabla_0}{a-2\nabla_0}a$. Note that the condition $a \geq 3\nabla_0 + 1$ implies that $\frac{\nabla_0}{a-1-2\nabla_0} \leq 1$, which is needed to preserve Equation (1). □

Now we have $a = b = 3\nabla_0$ and we can put $d_a = a + 1$. Hence, the optimum is $d_i = 0$ if $i < 3\nabla_0$, $d_{3\nabla_0} = 3\nabla_0 + 1$ and $d_{3\nabla_0+i} = \nabla_0/(\nabla_0 + i)$. Altogether, we get

$$S = 3\nabla_0 + 1 + \nabla_0 \sum_{i=\nabla_0+1}^{\Delta_R-2\nabla_0} \frac{1}{i} = \nabla_0(H_{\Delta_R-2\nabla_0} - H_{\nabla_0}) + 3\nabla_0 + 1,$$

where $H_i = 1 + 1/2 + \cdots + 1/i$ is the ith harmonic number.

It is known [13] that

$$\frac{1}{24(n+1)^2} < H_n - \ln\Big(n + \frac{1}{2}\Big) - \gamma < \frac{1}{24n^2},$$

where γ is the Euler–Mascheroni constant. We deduce that for $n > m$ we have

$$-\frac{1}{24m^2} < \frac{1}{24}\Big(\frac{1}{(n+1)^2} - \frac{1}{m^2}\Big) < (H_n - H_m) - \ln\Big(\frac{2n+1}{2m+1}\Big) < \frac{1}{24}\Big(\frac{1}{n^2} - \frac{1}{(m+1)^2}\Big) \leq 0$$

we deduce

$$-\frac{1}{24\nabla_0^2} < S - \left(\nabla_0 \ln\Big(\frac{2\Delta_R - 4\nabla_0 + 1}{2\nabla_0 + 1}\Big) + 3\nabla_0 + 1\right) < 0.$$

Hence, with a (negative) error less than 0.042 we have

$$S \approx \nabla_0 \ln\Big(\frac{2\Delta_R - 4\nabla_0 + 1}{2\nabla_0 + 1}\Big) + 3\nabla_0 + 1.$$

This finishes the proof of Lemma 7.2. □

8. $K_{3,t}$-free graphs

We now turn our attention to graphs that exclude $K_{3,t}$ for some t (and with bounded ∇_1). The most prominent graphs with these properties are graphs that embed into a surface of bounded genus and in particular planar graphs.

8.1. Phase 1: Preprocessing

The general preprocessing phase described in Section 4 remains unchanged. Recall that we defined D_1 as $\{v \in V(G) : \text{ for all } A \subseteq V(G) \setminus \{v\} \text{ with } N(v) \subseteq N[A] \text{ we have } |A| > (2\nabla - 1)\}$. Recall as well that for every $v \in V(G) \setminus D_1$, we fixed $A_v \subseteq V(G) \setminus \{v\}$ such that $N_R(v) \subseteq N[A_v]$ and $|A_v| \leq 2\nabla - 1$. Furthermore, for $V(G) \setminus \hat{D}$ we assumed $A_v \subseteq D \setminus \{v\}$.

8.2. Phase 2: local dominators in the $K_{3,t}$-free case

In this second phase, things get simpler than in Section 5. Since we now assume that we exclude $K_{3,t}$ the domination sequences of Definition 5.3 only have length two. We can therefore simplify the analysis of the domination sequences. We simply select every pair of vertices with sufficiently many neighbors in common.

- For $v \in V(G)$ let $B_v := \{z \in V(G) \setminus \{v\} : |N_R(v) \cap N_R(z)| \geq (2\nabla - 1)t + 1\}$.
- Let W be the set of vertices $v \in V(G)$ such that $B_v \neq \emptyset$.
- Let $D_2 := \bigcup_{v \in W} (\{v\} \cup B_v)$.

Once D_1 has been computed in the previous phase, 2 more rounds of communication are enough to compute the sets B_v and D_2. Before we update the residual degrees, let us analyze the sets B_v and D_2. First note that the definition is symmetric: since $N_R(v) \cap N_R(z) = N_R(z) \cap N_R(v)$ we have for all $v, z \in V(G)$ if $z \in B_v$, then $v \in B_z$. In particular, if $v \in D_1$ or $z \in D_1$, then $N_R(v) \cap N_R(z) = \emptyset$, which immediately implies the following lemma.

Lemma 8.1. We have $W \cap D_1 = \emptyset$ and for every $v \in V(G)$ we have $B_v \cap D_1 = \emptyset$.

Now we prove that for every $v \in W$, the set B_v cannot be too big, and has nice properties.

Lemma 8.2. For all vertices $v \in W$ we have

- $B_v \subseteq A_v$, (hence $|B_v| \leq (2\nabla - 1)$) and
- if $v \notin \hat{D}$, then $B_v \subseteq D$.

Proof. Assume $A_v = \{v_1, \ldots, v_\ell\}$ (a set of possibly not distinct vertices) and assume there exists $z \in V(G) \setminus \{v, v_1, \ldots v_\ell\}$ with $|N_R(v) \cap N_R(z)| \geq (2\nabla - 1)t + 1$. As $v_1, \ldots, v_\ell$ dominate $N_R(v)$, and hence also $N_R(v) \cap N_R(z)$, and $\ell \leq (2\nabla - 1)$, there must be some v_i, $1 \leq i \leq \ell$, with $|N_R(v) \cap N_R(z) \cap N[v_i]| \geq \lceil ((2\nabla - 1)t + 1)/(2\nabla - 1) \rceil \geq t + 1$. Therefore, $|N_R(v) \cap N_R(z) \cap N(v_i)| \geq t$, which shows that $K_{3,t}$ is a subgraph of G, contradicting the assumption.

If furthermore $v \notin \hat{D}$, by definition of $\hat{D}$, we can find $w_1, \ldots, w_\ell$ from D that dominate $N(v)$, and in particular $N_R(v)$. If $z \in V(G) \setminus \{v, w_1, \ldots, w_\ell\}$ with $|N_R(v) \cap N_R(z)| \geq (2\nabla - 1)t + 1$ we can argue as above to obtain a contradiction. □

Let us now analyze the size of D_2. For this we refine the set D_2 and define

1. $D_2^1 := \bigcup_{v \in W \cap D}(\{v\} \cup B_v)$,
2. $D_2^2 := \bigcup_{v \in W \cap (\hat{D} \setminus D)}(\{v\} \cup B_v)$, and
3. $D_2^3 := \bigcup_{v \in W \setminus (D \cup \hat{D})}(\{v\} \cup B_v)$.

Obviously $D_2 = D_2^1 \cup D_2^2 \cup D_2^3$. We now bound the size of the refined sets D_2^1, D_2^2 and D_2^3.

Lemma 8.3. $|D_2^1 \setminus D| \leq (2\nabla - 1)\gamma$.

Proof. We have

$$|D_2^1 \setminus D| = |\bigcup_{v \in W \cap D}(\{v\} \cup B_v) \setminus D| \leq |\bigcup_{v \in W \cap D} B_v| \leq \sum_{v \in W \cap D} |B_v|.$$

By Lemma 8.2 we have $|B_v| \leq (2\nabla - 1)$ for all $v \in W$ and as we sum over $v \in W \cap D$ we conclude that the last term has order at most $(2\nabla - 1)\gamma$. □

Lemma 8.4. $D_2^2 \subseteq \hat{D}$ and therefore $|D_2^2 \setminus D| < \rho(G)\gamma$.

Proof. Let $v \in \hat{D} \setminus D$ and let $z \in B_v$. By symmetry, $v \in B_z$ and according to Lemma 8.2, if $z \notin \hat{D}$, then $v \in D$. Since this is not the case, we conclude that $z \in \hat{D}$. Hence $B_v \subseteq \hat{D}$ and, more generally, $D_2^2 \subseteq \hat{D}$. Finally, according to Lemma 4.1 we have $|\hat{D} \setminus D| < \rho(G)\gamma$. □

Finally, the set D_2^3, which appears largest at first glance, was actually already counted, as shown in the next lemma.

Lemma 8.5. $D_2^3 \subseteq D_2^1$.

Proof. If $v \notin \hat{D}$, then $B_v \subseteq D$ by Lemma 8.2. Hence $v \in B_z$ for some $z \in D$, and $v \in D_2^1$. □

Recall that we defined $\eta \in [0, 1]$ to be such that $|(D_1 \cup D_2) \cap D| = \eta\gamma$.

Lemma 8.6. We have $|D_1 \cup D_2| < \rho(G)\gamma + 2\nabla\eta\gamma$.

Proof. By Lemma 8.5 we have $D_2^3 \subseteq D_2^1$, hence, $D_1 \cup D_2 = D_1 \cup D_2^1 \cup D_2^2$. By Lemma 4.2 we have $D_1 \subseteq \hat{D}$ and by Lemma 8.4 we also have $D_2^2 \subseteq \hat{D}$, hence $D_1 \cup D_2^2 \subseteq \hat{D}$. Again by Lemma 4.1, $|\hat{D} \setminus D| < \rho(G)\gamma$ and therefore $|(D_1 \cup D_2^2) \setminus D| < \rho(G)\gamma$.

We have $W \cap D \subseteq D_2^1 \cap D$, hence with Lemma 8.2 we conclude that

$$\left|D_2^1 \setminus D\right| \leq \left|\bigcup_{v \in D \cap D_2^1} B_v\right| \leq \sum_{v \in D \cap D_2^1} |B_v| \leq (2\nabla - 1)\rho\gamma,$$

hence $(D_1 \cup D_2) \setminus D < \rho(G)\gamma + (2\nabla - 1)\eta\gamma$. Finally, $D_1 \cup D_2 = (D_1 \cup D_2) \setminus D \cup ((D_1 \cup D_2) \cap D)$ and with the definition of η we conclude $|D_1 \cup D_2| < \rho(G)\gamma + 2\nabla\eta\gamma$. □

We now update the residual degrees, that is, we update R as $V(G) \setminus N[D_1 \cup D_2]$ and for every vertex the number $d_R(v) = |N_R(v)|$ accordingly.

Just as before, we show that after the first two phases of the algorithm we are in the very nice situation where all residual degrees are small.

Lemma 8.7. For all $v \in V(G)$ we have $d_R(v) \leq (2\nabla - 1)^2 t + (2\nabla - 1)$.

Proof. First, every vertex of $D_1 \cup D_2$ has residual degree 0. Assume that there is a vertex v of residual degree at least $(2\nabla - 1)^2 t + (2\nabla - 1) + 1$. As v is not in D_1, its residual neighbors are dominated by a set A_v of at most $(2\nabla - 1)$ vertices. Hence there is a vertex z (not in D_1 nor D_2) with $|N_R(v) \cap N_R[z]| \geq (2\nabla - 1)t + 2 = ((2\nabla - 1)t + 1) + 1$, hence, $|N_R(v) \cap N_R(z)| \geq (2\nabla - 1)t + 1$. This contradicts that v is not in D_2. □

8.3. Phase 3: LP-based approximation

We now proceed with the LP-based approximation as in the general case presented in Section 6. Recall that for any desired $\varepsilon > 0$ we defined Γ as an high degree and defined D_3^1 as the set of all vertices with degree greater than Γ. We added D_3^1 to the dominating set and were able to call the LP-based approximation algorithm of Corollary 6.1. We finally obtained a set D_3 dominating the remaining vertices with $|D_3| \leq (2\nabla_0 + 1)(1 + \varepsilon)\gamma_R$ according to Lemma 6.3.

We now conclude our main theorem, Theorem 2.2, stating that the algorithm on $K_{3,t}$-free graphs computes a dominating set of size at most $(6\nabla_1 + 3)\gamma$.

Proof of Theorem 2.2. First, D_1, D_2, and D_3 are computed locally, in a bounded number of rounds, and additionally the set $D_1 \cup D_2 \cup D_3$ dominates G. Then,

$$\begin{aligned} |D_1 \cup D_2| &< \rho(G)\gamma + 2\nabla\eta\gamma && \text{(by Lemma 8.6)} \\ &\leq (2\nabla_1 + 1)\gamma + 2\nabla_1\eta\gamma && \text{(as } \rho(G) \leq 2\nabla_1 + 1 \text{ and } \nabla \leq \nabla_1\text{)} \end{aligned}$$

and

$$\begin{aligned} |D_3| &\leq (2\nabla_0 + 1)(1 + \varepsilon)(1 - \eta)\gamma && \text{(by Lemma 6.3)} \\ &\leq (2\nabla_1 + 1)(1 + \varepsilon)(1 - \eta)\gamma && \text{(as } \nabla_0 \leq \nabla_1\text{)} \end{aligned}$$

By choosing $\varepsilon = 1$,

$$\begin{aligned} |D_1 \cup D_2 \cup D_3| &\leq (2\nabla_1 + 1 + 2\nabla_1\eta + (4\nabla_1 + 2)(1 - \eta))\gamma \\ &\leq (6\nabla_1 + 3)\gamma && \text{(maximized when } \eta = 0\text{)} \end{aligned}$$

□

8.4. Planar graphs

Finally, we complete the analysis of our algorithm for planar graphs, showing that it computes an $(11 + \varepsilon)$-approximation. In the following we fix a planar graph G.

Proof of Theorem 2.3. We revisit the proof of Theorem 2.2 and plug in the numbers for planar graphs. For planar graphs we have $K_{3,3} \not\subseteq G$, $\nabla_0 = 3$, $\nabla = 2$ and $\rho = 4$, as stated in Lemma 3.1. Therefore, by Lemma 8.6 we have $|D_1 \cup D_2| < 4\gamma + 4\eta\gamma$ and by Lemma 6.3 we have $|D_3| \leq (7 + \varepsilon)(1 - \eta)\gamma$. Hence, $|D_1 \cup D_2 \cup D_3| \leq \gamma(4 + 4\eta + 7 - 7\eta + \varepsilon - \varepsilon\eta) \leq \gamma(11 + \varepsilon - 3\eta - \varepsilon\eta)$. As $\eta \in [0, 1]$, this is maximized when $\eta = 0$. Hence $|D_1 \cup D_2 \cup D_3| \leq \gamma(11 + \varepsilon)$. □

We now further restrict the input graphs, requiring e.g. planarity together with a lower bound on the girth. Our algorithm works exactly as before, however, using different parameters. From the different edge densities and Hall ratio of the restricted graphs we will then derive different constants and as a result a better approximation factor. Throughout this section we use the same notation as in the first part of the paper.

As in the general case in the first phase we begin by computing the set D_1 and analyzing it in terms of the auxiliary set $\hat{D}$.

Corollary 8.1.

1. If G is bipartite, then $|\hat{D} \setminus D| < 2\gamma$.
2. If G is triangle-free, outerplanar, or has girth 5, then $|\hat{D} \setminus D| < 3\gamma$.

Proof. This is immediate from Lemma 3.1 and Lemma 4.1. □

The inclusion $D_1 \subseteq \hat{D}$ continues to hold and the bound on the sizes as stated in the next lemma is again a direct consequence of the corollary.

Lemma 8.8. We have $D_1 \subseteq \hat{D}$, and

1. *if G is bipartite, then $|\hat{D} \setminus D| < 2\gamma$ and $|\hat{D}| < 3\gamma$.*
2. *if G is triangle-free, outerplanar, or has girth 5, then $|\hat{D} \setminus D| < 3\gamma$ and $|\hat{D}| < 4\gamma$.*

In case of triangle-free planar graphs (in particular in the case of bipartite planar graphs) we proceed with the second phase exactly as in the second phase of the general algorithm (Section 8.2), however, the parameter $(2\nabla - 1)t + 1$ is replaced by the parameter 7. In case of planar graphs of girth at least five or outerplanar graphs, we simply set $D_2 = \emptyset$.

If G is triangle-free:

- For $v \in V(G)$ let $B_v := \{z \in V(G) \setminus \{v\} : |N_R(v) \cap N_R(z)| \geq 7\}$.
- Let W be the set of vertices $v \in V(G)$ such that $B_v \neq \emptyset$.
- Let $D_2 := \bigcup_{v \in W} (\{v\} \cup B_v)$.

If G has girth at least 5 or G is outerplanar, let $D_2 = \emptyset$.

Lemma 8.1 is based only on the definition of B_v and W and does not use particular properties of planar graphs, hence, it also holds in the restricted case.

The next lemma uses the triangle-free property.

Lemma 8.9. If G is triangle-free, then for all vertices $v \in W$ we have

- $B_v \subseteq A_v$ (hence $|B_v| \leq 3$), and
- if $v \notin \hat{D}$, then $B_v \subseteq D$.

Proof. Assume $A_v = \{v_1, v_2, v_3\}$ and assume there is $z \in V(G) \setminus \{v, v_1, v_2, v_3\}$ with $|N_R(v) \cap N_R(z)| \geq 7$. As the vertices v_1, v_2, v_3 dominate $N_R(v)$, and hence $N_R(v) \cap N_R(z)$, there must be some v_i, $1 \leq i \leq 3$, with $|N_R(v) \cap N_R(z) \cap N[v_i]| \geq \lceil 7/3 \rceil \geq 3$. Then on of the following holds: either $|N_R(v) \cap N_R(z) \cap N(v_i)| \geq 3$, or $|N_R(v) \cap N_R(z) \cap N(v_i)| = 2$. The first case shows that $K_{3,3}$ is a subgraph of G contradicting the assumption that G is planar. The second case implies that $v_i \in N_R(v)$. In this situation, by picking $w \in N_R(v) \cap N_R(z) \cap N(v_i)$, we get that (v, v_i, w) is a triangle, hence we also reach a contradiction.

If furthermore $v \notin \hat{D}$, by definition of $\hat{D}$, we can find w_1, w_2, w_3 from D that dominate $N(v)$, and in particular $N_R(v)$. If $z \in V(G) \setminus \{v, w_1, w_2, w_3\}$ with $|N_R(v) \cap N_R(z)| \geq 7$ we can argue as above to obtain a contradiction. □

For our analysis we again split D_2 into three sets D_2^1, D_2^2 and D_2^3. The next lemmas hold also for the restricted cases. We repeat them for convenience with the appropriate numbers filled it.

Lemma 8.10. If G is triangle-free, then $|D_2^1 \setminus D| \leq 3\gamma$.

Lemma 8.11. If G is triangle-free, then $D_2^2 \subseteq \hat{D}$ and therefore $|D_2^2 \setminus D| < 3\gamma$.

Lemma 8.12. If G is triangle-free, then $D_2^3 \subseteq D_2^1$.

Recall that $\eta \in [0, 1]$ is such that $|(D_1 \cup D_2) \cap D| = \eta\gamma$.

Lemma 8.13.

1. If G is bipartite, then $|D_1 \cup D_2| < 2\gamma + 4\eta\gamma$.
2. If G is triangle-free, then $|D_1 \cup D_2| < 3\gamma + 4\eta\gamma$.
3. If G has girth at least 5 or is outerplanar, then $|D_1 \cup D_2| < 3\gamma + \eta\gamma$.

Proof. If G is outerplanar or G has girth at least 5, then $D_2 = \emptyset$. By Lemma 8.8 we have $D_1 \subseteq \hat{D}$ and $|\hat{D} \setminus D| < 3\gamma$, hence $(D_1 \cup D_2) \setminus D < 3\gamma$.

If G is triangle-free, by Lemma 8.12 we have $D_2^3 \subseteq D_2^1$, hence, $D_1 \cup D_2 = D_1 \cup D_2^1 \cup D_2^2$. By Lemma 8.8 we have $D_1 \subseteq \hat{D}$ and by Lemma 8.11 we also have $D_2^2 \subseteq \hat{D}$, hence $D_1 \cup D_2^2 \subseteq \hat{D}$. Again by Lemma 8.8, if G is bipartite, then $|\hat{D} \setminus D| < 2\gamma$, therefore $|(D_1 \cup D_2^2) \setminus D| < 2\gamma$, and if G is triangle-free, then $|\hat{D} \setminus D| < 3\gamma$, therefore $|(D_1 \cup D_2^2) \setminus D| < 3\gamma$. We have $W \cap D \subseteq D_2^1 \cap D$, hence with Lemma 8.9 we conclude that

$$\left|D_2^1 \setminus D\right| \leq \left|\bigcup_{v \in D \cap D_2^1} B_v\right| \leq \sum_{v \in D \cap D_2^1} |B_v| \leq 3\eta\gamma,$$

hence $(D_1 \cup D_2) \setminus D < 2\gamma + 3\eta\gamma$ if G is bipartite and $(D_1 \cup D_2) \setminus D < 3\gamma + 3\eta\gamma$ if G is triangle-free.

Finally, $D_1 \cup D_2 = (D_1 \cup D_2) \setminus D \cup (D_1 \cup D_2) \cap D$ and with the definition of η we conclude

1. $|D_1 \cup D_2| < 2\gamma + 4\eta\gamma$ if G is bipartite,
2. $|D_1 \cup D_2| < 3\gamma + 4\eta\gamma$ if G is triangle-free,
3. $|D_1 \cup D_2| < 3\gamma + \eta\gamma$ if G has girth at least 5 or is outerplanar.

□

Again, we update the residual degrees, that is, we update R as $V(G) \setminus N[D_1 \cup D_2]$ and for every vertex the number $d_R(v) = N(v) \cap R$ accordingly and proceed with the third phase.

Lemma 8.14.

1. *If G is triangle-free, then $\Delta_R(v) \leq 18$.*
2. *If G has girth at least 5, then $\Delta_R(v) \leq 3$.*
3. *If G is outerplanar, then $\Delta_R(v) \leq 9$.*

Proof. Every vertex of $D_1 \cup D_2$ has residual degree 0, hence, we need to consider only vertices that are not in D_1 or D_2.

First assume that the graph is triangle-free and assume that there is a vertex v of residual degree at least 19. As v is not in D_1, its 19 non-dominated neighbors are dominated by a set A_v of at most 3 vertices. Hence, there is vertex z (not in D_1 nor D_2) dominating at least $\lceil 19/3 \rceil = 7$ of them. Here, z cannot be one of these 7 vertices, otherwise it would be connected to v and there would be a triangle in the graph. Therefore we have $|N_R(v) \cap N_R(z)| \geq 7$, contradicting that v is not in D_2.

Now assume that G has girth at least 5 and assume that there is a vertex v of residual degree at least 4. As v is not in D_1, its 4 non-dominated neighbors are dominated by a set A_v of at most 3 vertices. Hence, there is vertex z (not in D_1 nor D_2) dominating at least $\lceil 4/3 \rceil = 2$ of them. Here, z cannot be one of these 2 vertices, otherwise it would be connected to v and there would be a triangle in the graph. However, z can also not be any other vertex, as otherwise we find a cycle of length 4, contradicting that G has girth at least 5.

Finally, assume that G is outerplanar and assume that there is a vertex v of residual degree at least 10. As v is not in D_1, its 10 non-dominated neighbors are dominated by a set A_v of at most 3 vertices. Hence, there is vertex z (not in D_1 nor D_2) dominating at least $\lceil 10/3 \rceil = 4$ of them. Therefore $|N(v) \cap N(z)| \geq 3$, and we find a $K_{2,3}$ as a subgraph, contradicting that G is outerplanar. □

We proceed with the LP-based approximation as in the general case presented in Section 6. Recall that for any desired $\varepsilon > 0$ we defined Γ as an high degree and defined D_3^1 as the set of all vertices with degree greater than Γ. We added D_3^1 to the dominating set and were able to call the LP-based approximation algorithm of Corollary 6.1. We finally obtained a set D_3 dominating the remaining vertices with $|D_3| \leq (2\nabla_0 + 1)(1 + \varepsilon/5)\gamma_R$ according to Lemma 6.3 (by applying it with $\varepsilon' = \varepsilon/5$). The second item does not follow by the LP approximation, but simply from Lemma 8.14 and Corollary 5.2.

Lemma 8.15.

1. *If G is triangle-free, then $|D_3| \leq 5(1 + \varepsilon/5)\gamma_R$.*
2. *If G has girth at least 5, then $|D_3| \leq 4\gamma_R$.*
3. *If G is outerplanar, then $|D_3| \leq 5(1 + \varepsilon/5)\gamma_R$.*

We conclude.

Theorem 8.1. There exists a distributed LOCAL algorithm that, for every triangle free planar graph G, computes in a constant number of rounds a dominating set of size at most $(8+\varepsilon)\gamma(G)$.

Proof. We have

$$|D_1 \cup D_2| < 3\gamma + 4\eta\gamma \qquad \text{(by Lemma 8.13)}$$
$$|D_3| \leq 5(1+\varepsilon/5)\gamma_R \leq 5(1+\varepsilon/5)(1-\eta)\gamma \qquad \text{(by Lemma 8.15)}$$

Thus,

$$\begin{aligned} |D_1 \cup D_2 \cup D_3| &< 8\gamma - 5\eta\gamma + \varepsilon\gamma - \varepsilon\eta\gamma \\ &\leq (8+\varepsilon)\gamma \qquad \text{(maximized when } \eta = 0) \end{aligned}$$

□

The remaining theorems are proved analogously.

Theorem 8.2. There exists a distributed LOCAL algorithm that, for every bipartite planar graph G, computes in a constant number of rounds a dominating set of size at most $(7+\varepsilon)\gamma(G)$.

Theorem 8.3. There exists a distributed LOCAL algorithm that, for every planar graph G of girth at least 5, computes in a constant number of rounds a dominating set of size at most $7\gamma(G)$.

Theorem 8.4. There exists a distributed LOCAL algorithm that, for every outerplanar graph G, computes in a constant number of rounds a dominating set of size at most $(8+\varepsilon)\gamma(G)$.

9. Conclusion

We simplified the presentation and generalized the algorithm of Czygrinow et al. [12] from graph classes that exclude some topological minor to graph classes $\mathscr{C}$ where $\nabla_1(G)$ is bounded by an absolute constant for all $G \in \mathscr{C}$. This is a property in particular possessed by classes with bounded expansion, which include many commonly studied sparse graph classes. The obtained general bounds are still large, but by magnitudes smaller than those obtained in the original work of Czygrinow et al.[12].

It is an interesting and important question to identify the most general graph classes on which certain algorithmic techniques work. The key argument of Lemma 4.1 works only for classes with $\nabla_1(G)$ bounded by an absolute constant. We need different methods to push towards classes with only $\nabla_0(G)$ bounded, which are the degenerate classes.

We then provided a fine-tuned LOCAL algorithm that computes an $(11+\varepsilon)$-approximation of a minimum dominating set in a planar graph in a constant number of rounds. Started with different parameters, the algorithm works also for several restricted cases of planar graphs. We showed that it computes an $(8+\varepsilon)$-approximation for triangle-free planar graphs, a $(7+\varepsilon)$-approximation for bipartite planar graphs, a 7-approximation for planar graphs of girth 5 and an $(8+\varepsilon)$-approximation for outerplanar graphs. In all cases except for the outerplanar case, where an optimal bound of 5 was already known, our algorithm improves on the previously best known approximation factors. This improvement is most significant in the case of general planar graphs, where the previously best known factor was 52.

References

[1] S. Akhoondian Amiri, P. Ossona de Mendez, R. Rabinovich, and S. Siebertz. Distributed domination on graph classes of bounded expansion. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures*, pages 143–151, 2018.

[2] S. Akhoondian Amiri, S. Schmid, and S. Siebertz. A local constant factor mds approximation for bounded genus graphs. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 227–233, 2016.

[3] S. Alipour, E. Futuhi, and S. Karimi. On distributed algorithms for minimum dominating set problem, from theory to application. *arXiv preprint arXiv:2012.04883*, 2020.

[4] S. Alipour and A. Jafari. A local constant approximation factor algorithm for minimum dominating set of certain planar graphs. In *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures*, pages 501–502, 2020.

[5] S. A. Amiri and S. Schmid. Brief announcement: A log*-time local mds approximation scheme for bounded genus graphs. In *Proc. DISC*, 2016.

[6] S. A. Amiri, S. Schmid, and S. Siebertz. Distributed dominating set approximations beyond planar graphs. *ACM Transactions on Algorithms (TALG)*, 15(3):1–18, 2019.

[7] N. Bansal and S. W. Umboh. Tight approximation bounds for dominating set on graphs of bounded arboricity. *Information Processing Letters*, 122:21–24, 2017.

[8] L. Barenboim, M. Elkin, and C. Gavoille. A fast network-decomposition algorithm and its applications to constant-time distributed computation. *Theoretical Computer Science*, 751:2–23, 2018.

[9] M. Bonamy, L. Cook, C. Groenland, and A. Wesolek. A tight local algorithm for the minimum dominating set problem in outerplanar graphs. In *DISC*, volume 209 of *LIPIcs*, pages 13:1–13:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[10] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite vc-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.

[11] A. Czygrinow, M. Hańćkowiak, and W. Wawrzyniak. Fast distributed approximations in planar graphs. In *International Symposium on Distributed Computing*, pages 78–92. Springer, 2008.

[12] A. Czygrinow, M. Hanckowiak, W. Wawrzyniak, and M. Witkowski. Distributed approximation algorithms for the minimum dominating set in k_h-minor-free graphs. In *29th International Symposium on Algorithms and Computation (ISAAC 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[13] D. DeTemple. The non-integer property of sums of reciprocals of consecutive integers. *Math. Gaz.*, 75:193–194, 1991.

[14] R. Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.

[15] I. Dinur and D. Steurer. Analytical approach to parallel repetition. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 624–633, 2014.

[16] P. G. Drange, M. S. Dregi, F. V. Fomin, S. Kreutzer, D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, F. Reidl, F. S. Villaamil, S. Saurabh, S. Siebertz, and S. Sikdar. Kernelization and sparseness: the case of dominating set. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016*, pages 31:1–31:14, 2016.

[17] D.-Z. Du and P.-J. Wan. *Connected dominating set: theory and applications*, volume 77. Springer Science & Business Media, 2012.

[18] Z. Dvořák. On distance-dominating and-independent sets in sparse graphs. *Journal of Graph Theory*, 91(2):162–173, 2019.

[19] E. Eiben, M. Kumar, A. E. Mouawad, F. Panolan, and S. Siebertz. Lossy kernels for connected dominating set on sparse graphs. *SIAM Journal on Discrete Mathematics*, 33(3):1743–1771, 2019.

[20] K. Eickmeyer, A. C. Giannopoulou, S. Kreutzer, O. Kwon, M. Pilipczuk, R. Rabinovich, and S. Siebertz. Neighborhood complexity and kernelization for nowhere dense classes of graphs. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, pages 63:1–63:14, 2017.

[21] G. Even, D. Rawitz, and S. M. Shahar. Hitting sets when the vc-dimension is small. *Information Processing Letters*, 95(2):358–362, 2005.

[22] G. Fabianski, M. Pilipczuk, S. Siebertz, and S. Torunczyk. Progressive algorithms for domination and independence. In *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, pages 27:1–27:16, 2019.

[23] J. Gajarskỳ, P. Hliněnỳ, J. Obdržálek, S. Ordyniak, F. Reidl, P. Rossmanith, F. S. Villaamil, and S. Sikdar. Kernelization using structural parameters on sparse graph classes. *Journal of Computer and System Sciences*, 84:219–242, 2017.

[24] R. G. Gallager, P. A. Humblet, and P. M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Transactions on Programming Languages and systems (TOPLAS)*, 5(1):66–77, 1983.

[25] M. R. Garey and D. S. Johnson. *Computers and intractability*, volume 174. freeman San Francisco, 1979.

[26] M. Ghaffari, F. Kuhn, and Y. Maus. On the complexity of local distributed graph problems. In *STOC*, pages 784–797. ACM, 2017.
[27] S. Hakimi. On the degree of the vertices of a directed graph. *J. Franklin Inst.*, 279(4):290–308, 1965.
[28] S. Har-Peled and K. Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. *SIAM Journal on Computing*, 46(6):1712–1744, 2017.
[29] O. Heydt, S. Siebertz, and A. Vigny. Local planar domination revisited. *arXiv preprint arXiv:2111.14506*, 2021. Accepted at SIROCCO 2022.
[30] M. Hilke, C. Lenzen, and J. Suomela. Brief announcement: local approximability of minimum dominating set on planar graphs. In *Proceedings of the 2014 ACM symposium on Principles of distributed computing*, pages 344–346, 2014.
[31] D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of computer and system sciences*, 9(3):256–278, 1974.
[32] M. Jones, D. Lokshtanov, M. Ramanujan, S. Saurabh, and O. Suchý. Parameterized complexity of directed steiner tree on sparse graphs. *SIAM Journal on Discrete Mathematics*, 31(2):1294–1327, 2017.
[33] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
[34] S. Kreutzer, R. Rabinovich, and S. Siebertz. Polynomial kernels and wideness properties of nowhere dense graph classes. *ACM Transactions on Algorithms (TALG)*, 15(2):1–19, 2018.
[35] S. Kublenz, S. Siebertz, and A. Vigny. Constant round distributed domination on graph classes with bounded expansion. *arXiv preprint arXiv:2012.02701*, 2020.
[36] F. Kuhn, T. Moscibroda, and R. Wattenhofer. The price of being near-sighted. In *SODA*, volume 6, pages 1109557–1109666. Citeseer, 2006.
[37] F. Kuhn, T. Moscibroda, and R. Wattenhofer. Local computation: Lower and upper bounds. *J. ACM*, 63(2):17:1–17:44, 2016.
[38] C. Lenzen, Y.-A. Pignolet, and R. Wattenhofer. Distributed minimum dominating set approximations in restricted families of graphs. *Distributed computing*, 26(2):119–137, 2013.
[39] C. Lenzen and R. Wattenhofer. Leveraging linial's locality limit. In *International Symposium on Distributed Computing*, pages 394–407. Springer, 2008.
[40] C. Lenzen and R. Wattenhofer. Minimum dominating set approximation in graphs of bounded arboricity. In *International symposium on distributed computing*, pages 510–524. Springer, 2010.
[41] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete mathematics*, 13(4):383–390, 1975.
[42] J. Nešetřil and P. Ossona de Mendez. Grad and classes with bounded expansion I. decompositions. *European Journal of Combinatorics*, 29(3):760–776, 2008.
[43] J. Nešetřil and P. Ossona de Mendez. *Sparsity: graphs, structures, and algorithms*, volume 28. Springer Science & Business Media, 2012.
[44] J. Nešetřil, P. Ossona de Mendez, and D. R. Wood. Characterisations and examples of graph classes with bounded expansion. *European Journal of Combinatorics*, 33(3):350–373, 2012.
[45] V. Rozhon and M. Ghaffari. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In *STOC*, pages 350–363. ACM, 2020.
[46] A. Sasireka and A. N. Kishore. Applications of dominating set of a graph in computer networks. *Int. J. Eng. Sci. Res. Technol*, 3(1):170–173, 2014.
[47] S. Siebertz. Greedy domination on biclique-free graphs. *Information Processing Letters*, 145:64–67, 2019.
[48] J. Suomela. Survey of local algorithms. *ACM Computing Surveys (CSUR)*, 45(2):1–40, 2013.
[49] W. Wawrzyniak. Brief announcement: a local approximation algorithm for mds problem in anonymous planar networks. In *Proceedings of the 2013 ACM symposium on Principles of distributed computing*, pages 406–408, 2013.
[50] W. Wawrzyniak. A strengthened analysis of a local algorithm for the minimum dominating set problem in planar graphs. *Information Processing Letters*, 114(3):94–98, 2014.