

Distributed Dominating Set Approximations beyond Planar Graphs

SAEED AKHOONDIAN AMIRI, Max-Planck-Institut für Informatik, Germany

STEFAN SCHMID, Faculty of Computer Science, University of Vienna, Austria

SEBASTIAN SIEBERTZ, Uniwersytet Warszawski, Poland

The Minimum Dominating Set (MDS) problem is a fundamental and challenging problem in distributed computing. While it is well-known that minimum dominating sets cannot be well approximated locally on general graphs, over the last years, there has been much progress on computing good local approximations on sparse graphs, and in particular on planar graphs. In this paper we study distributed and deterministic MDS approximation algorithms for graph classes beyond planar graphs. In particular, we show that existing approximation bounds for planar graphs can be lifted to bounded genus graphs and more general graphs, which we call locally embeddable graphs, and present

- (1) a local constant-time, constant-factor MDS approximation algorithm on locally embeddable graphs, and
- (2) a local $O(\log^* n)$ -time $(1+\epsilon)$ -approximation scheme for any $\epsilon > 0$ on graphs of bounded genus.

Our main technical contribution is a new analysis of a slightly modified variant of an existing algorithm by Lenzen et al. Interestingly, unlike existing proofs for planar graphs, our analysis does not rely on direct topological arguments, but on combinatorial density arguments only.

Additional Key Words and Phrases: Distributed Computing, Dominating Set, Approximation Algorithms

1 INTRODUCTION

This paper attends to the Minimum Dominating Set (MDS) problem, an intensively studied graph theoretic problem in computer science in general, as well as in distributed computing.

A dominating set D in a graph G is a set of vertices such that every vertex of G either lies in D or is adjacent to a vertex in D . Finding a minimum dominating set is NP-complete [19], even on planar graphs of maximum degree 3 (cf. [GT2] in [14]). Consequently, attention has shifted from computing exact solutions to approximating near optimal dominating sets. The simple greedy algorithm computes an $\ln n$ approximation of a minimum dominating set [18, 25], and for general graphs this algorithm is near optimal – it is NP-hard to approximate minimum dominating sets within factor $(1 - \epsilon) \cdot \ln n$ for every $\epsilon > 0$ [10].

The approach of algorithmic graph structure theory is to exploit structural properties of restricted graph classes for the design of efficient algorithms. For the dominating set problem this has led to

Preliminary versions of this paper appeared as [3] and [4]. The research of Saeed Akhoondian Amiri has been partially supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (ERC consolidator grant DISTRUCT, agreement No 648527). Sebastian Siebertz is partially supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (ERC consolidator grant DISTRUCT, agreement No 648527) and by the National Science Centre of Poland via POLONEZ grant agreement UMO-2015/19/P/ST6/03998, which has received funding from the European Union's Horizon 2020 research and innovation programme (Marie Skłodowska-Curie grant agreement No. 665778).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2017 Copyright held by the owner/author(s). XXXX-XXXX/1997/1-ART1 \$15.00

DOI: 10.475/123_4



a PTAS on planar graphs [5], minor closed classes of graphs with locally bounded tree-width [12], graphs with excluded minors [15], and most generally, on every graph class with subexponential expansion [16]. The problem admits a constant factor approximation on classes of bounded arboricity [6] and an $O(\ln k)$ approximation (where k denotes the size of a minimum dominating set) on classes of bounded VC-dimension [8, 13]. On the other hand, it is unlikely that polynomial-time constant factor approximations exist even on $K_{3,3}$ -free graphs [28]. The general goal of algorithmic graph structure theory is to identify the broadest graph classes on which certain algorithmic techniques can be applied and hence lead to efficient algorithms for problems that are hard on general graphs. These limits of tractability are often captured by abstract notions, such as expansion, arboricity 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 [24]. It is known that finding small dominating sets locally is hard: Kuhn et al. [20] show that in r rounds the MDS problem on an n -vertex graphs of maximum degree Δ can only be approximated within factor $\Omega(n^{c/r^2}/r)$ and $\Omega(\Delta^{1/(r+1)}/r)$, where c is a constant. This implies that, in general, to achieve a constant approximation ratio, every distributed algorithm requires at least $\Omega(\sqrt{\log n/\log \log n})$ and $\Omega(\log \Delta/\log \log \Delta)$ communication rounds. The currently best results for general graphs are by Kuhn et al. [20] who present a $(1 + \epsilon) \ln \Delta$ -approximation in $O(\log(n)/\epsilon)$ rounds for any $\epsilon > 0$, and by Barenboim et al. [7] who present a deterministic $O((\log n)^{k-1})$ -time algorithm that provides an $O(n^{1/k})$ -approximation, for any integer parameter $k \geq 2$.

For sparse graphs, the situation is more promising (an inclusion diagram of the graph classes mentioned in the following paragraph is depicted in Figure 1, for formal definitions we refer to the referenced papers). For graphs of arboricity a , Lenzen and Wattenhofer [23] present a forest decomposition algorithm achieving a factor $O(a^2)$ approximation in randomized time $O(\log n)$, and a deterministic $O(a \log \Delta)$ approximation algorithm requiring $O(\log \Delta)$ rounds. Graphs of bounded arboricity include all graphs which exclude a fixed graph as a (topological) minor and in particular, all planar graphs and any class of bounded genus. Amiri et al. [1] provide a deterministic $O(\log n)$ time constant factor approximation algorithm on classes of bounded expansion (which extends also to connected dominating sets). The notion of bounded expansion offers an abstract definition of uniform sparseness in graphs, which is based on bounding the density of shallow minors (these notions will be defined formally in the next section). Czygrinow et al. [9] show that for any given $\epsilon > 0$, $(1 + \epsilon)$ -approximations of a maximum independent set, a maximum matching, and a minimum dominating set, can be computed in $O(\log^* n)$ rounds in planar graphs, which is asymptotically optimal [22]. Lenzen et al. [21] proposed a constant factor approximation on planar graphs that can be computed locally in a constant number of communication rounds. A finer analysis of Wawrzyniak [31] showed that the algorithm of Lenzen et al. in fact computes a 52-approximation of a minimum dominating set. Wawrzyniak [30] also showed that message sizes of $O(\log n)$ suffice to give a constant factor approximation on planar graphs in a constant number of rounds. In terms of lower bounds, Hilke et al. [17] show that there is no deterministic local algorithm (constant-time distributed graph algorithm) that finds a $(7 - \epsilon)$ -approximation of a minimum dominating set on planar graphs, for any positive constant ϵ .

1.1 Our Contributions

The first and main contribution of this paper is a deterministic and local constant factor approximation for MDS on graphs that we call *locally embeddable graphs*. A locally embeddable graph G excludes the complete bipartite graph $K_{3,t}$, for some $t \geq 3$, as a depth-1 minor, that is, as a minor obtained by star contractions, and furthermore satisfies that all depth-1 minors of G have constant

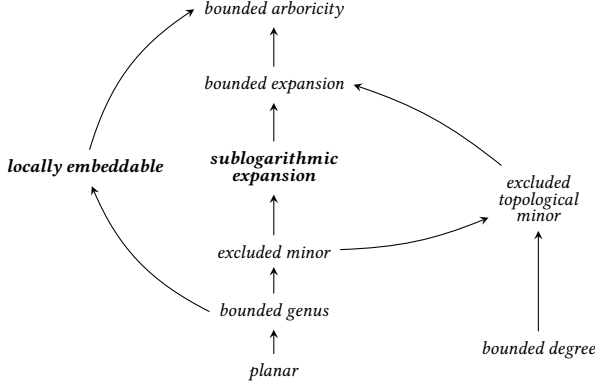


Fig. 1. Inclusion diagram of sparse graph classes.

edge density. The most prominent locally embeddable graph classes are classes of bounded genus. Concretely, our result implies that MDS can be $O(g)$ -approximated locally and deterministically on graphs of (both orientable or non-orientable) genus g . However, also graph classes whose members do not embed into any fixed surface or which do not even have bounded expansion can be locally embeddable, e.g. the class of all 3-subdivided cliques is locally embeddable and this class does not have bounded expansion. However, every locally embeddable class is degenerate. Apart from generalizing the earlier result of Lenzen et al. [21] for planar graphs to a larger graph family, we introduce new techniques by arguing about densities and combinatorial properties of shallow minors only and show that all topological arguments used in [21] can be avoided. The abstract notion of local embeddability yields exactly the ingredients for these arguments to work and therefore offers valuable insights on the limits of algorithmic techniques. This is a contribution going beyond the mere presentation of efficient algorithms for concrete example graph classes.

Our second main contribution is the presentation of a local and deterministic MDS approximation algorithm with the following properties. Given a graph G from a fixed class \mathcal{C} of graphs with sub-logarithmic expansion, a constant factor approximation of an MDS D and any $\epsilon > 0$, the algorithm uses $O(\log^* n)$ rounds and computes from D a $(1 + \epsilon)$ -approximate MDS of G (here, the O -notation hides constants depending on ϵ). Graphs of sub-logarithmic expansion include all proper minor closed classes, and in particular all classes of bounded genus. Our methods are based on earlier work of Czygrinow et al. [9]. In combination with our constant-factor approximation on graphs of bounded genus, we obtain $(1 + \epsilon)$ -approximations in $O(\log^* n)$ communication rounds on graphs of bounded genus. In combination with Amiri et al.'s result [1] on graphs of bounded expansion, we obtain $(1 + \epsilon)$ -approximations in $O(\log n)$ deterministic rounds on graphs of sub-logarithmic expansion. Again, the abstract notion of sub-logarithmic expansion constitutes the border of applicability of these algorithmic techniques.

We observe that the methods of Czygrinow et al. [9] for maximum weighted independent set and maximum matching extend to graphs of sub-logarithmic expansion, however, we focus on the dominating set problem for the sake of a consistent presentation.

1.2 Novelty

Our main technical contribution is a new analysis of a slightly modified variant of the elegant algorithm by Lenzen et al. [21] for planar graphs. As we will show, with a slight modification, the

algorithm also works on locally embeddable graphs, however, the analysis needs to be changed significantly. Prior works by Lenzen et al. [21] and Wawrzyniak [31] heavily depend on topological properties of planar graphs. For example, their analyses exploit the fact that each cycle in a planar graph defines an “inside” and an “outside” region, without any edges connecting the two; this facilitates a simplified accounting and comparison to the optimal solution. In the case of locally embeddable graphs, such global, topological properties do not exist. In contrast, in this paper we leverage the inherent local properties of our low-density graphs, which opens a new door to approach the problem.

A second interesting technique developed in this paper is based on *preprocessing*: we show that the constants involved in the approximation can be further improved by a local preprocessing step.

Another feature of our modified algorithm is that it is *first-order definable*. More precisely, there is a first order formula $\varphi(x)$ with one free variable, such that in every planar graph G the set $D = \{v \in V(G) : G \models \varphi(v)\}$ corresponds exactly to the computed dominating set. In particular, the algorithm can be modified such that it does not rely on any *maximum* operations, such as finding the neighbor of maximal degree.

1.3 Organization

The remainder of this paper is organized as follows. We introduce some preliminaries in Section 2. The constant-factor constant-time approximation result is presented in Section 3, and the $O(\log^* n)$ -time approximation scheme is presented in Section 4. We conclude in Section 5.

2 PRELIMINARIES

Graphs. We consider finite, undirected, simple graphs. Given a graph G , we write $V(G)$ for its vertices and $E(G)$ for its edges. Two vertices $u, v \in V(G)$ are adjacent or neighbors if $\{u, v\} \in E(G)$. The degree $d_G(v)$ of a vertex $v \in V(G)$ is its number of neighbors in G . We write $N(v)$ for the set of neighbors and $N[v]$ for the closed neighborhood $N(v) \cup \{v\}$ of v . For $A \subseteq V(G)$, we write $N[A]$ for $\bigcup_{v \in A} N[v]$. We let $N^1[v] := N[v]$ and $N^{i+1}[v] := N[N^i[v]]$ for $i > 1$. If $E' \subseteq E$, we write $N_{E'}(v)$ for the set $\{u \in V(G) : \{u, v\} \in E'\}$. A graph G has radius at most r if there is a vertex $v \in V(G)$ such that $N^r[v] = V(G)$. The *arboricity* of G is the minimum number of forests into which its edges can be partitioned. A graph H is a subgraph of a graph G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. The edge density of G is the ratio $|E(G)|/|V(G)|$. It is well known that the arboricity of a graph is within factor 2 of its *degeneracy*, that is, $\max_{H \subseteq G} |E(H)|/|V(H)|$. For $A \subseteq V(G)$, the graph $G[A]$ induced by A is the graph with vertex set A and edge set $\{\{u, v\} \in E(G) : u, v \in A\}$. For $B \subseteq V(G)$ we write $G - B$ for the graph $G[V(G) \setminus B]$.

Bounded depth minors and locally embeddable graphs. A graph H is a minor of a graph G , written $H \leq G$, if there is a set $\{G_v : v \in V(H)\}$ of pairwise 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 G_v is *contracted* to the vertex v . If $G_1, \dots, G_k \subseteq V(G)$ are pairwise vertex disjoint and connected subgraphs of G , then we write $G/G_1/\dots/G_k$ for the minor obtained by contracting the subgraphs G_i (observe that the order of contraction does not matter as the G_i 's are vertex disjoint). We call the set $\{G_v : v \in V(H)\}$ a *minor model* of H in G . We say that two minor models $\{G_v^1 : v \in V(H)\}$ and $\{G_v^2 : v \in V(H)\}$ of H in a graph G disjoint if the sets $\bigcup_{v \in V(H)} V(G_v^1)$ and $\bigcup_{v \in V(H)} V(G_v^2)$ are disjoint.

A star is a connected graph G such that at most one vertex of G , called the center of the star, has degree greater than one. A graph H is a *depth-1 minor* of G if H is obtained from a subgraph of G by star contractions, that is, if there is a set $\{G_v : v \in V(H)\}$ of pairwise vertex disjoint stars $G_v \subseteq G$ 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 .

More generally, for a non-negative integer r , a graph H is a *depth- r minor* of G , written $H \leq_r G$, if there is a set $\{G_v : v \in V(H)\}$ of pairwise vertex disjoint connected subgraphs $G_v \subseteq G$ of radius at most r 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 write $K_{t,3}$ for the complete bipartite graph with partitions of size t and 3, respectively. A graph G is a *locally embeddable graph* if it excludes $K_{3,t}$ as a depth-1 minor for some $t \geq 3$ and if $|E(H)|/|V(H)| \leq c$ for some constant c and all depth-1 minors H of G .

More generally, we write $\nabla_r(G)$ for $\max_{H \leq_r G} |E(H)|/|V(H)|$. A class \mathcal{C} of graphs has *bounded expansion* if there is a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $\nabla_r(G) \leq f(r)$ for all graphs $G \in \mathcal{C}$. This is equivalent to demanding that the arboricity of each depth- r minor of G is functionally bounded by r . The class \mathcal{C} has *sub-logarithmic expansion* if the bounding function $f(r) \in o(\log r)$. Note that if every graph $G \in \mathcal{C}$ excludes a fixed minor, then \mathcal{C} has constant expansion, hence classes of sub-logarithmic expansion generalize proper minor closed classes of graphs. We refer to Figure 1 for the inclusion between the above defined classes.

Bounded genus graphs. The (orientable, resp. non-orientable) genus of a graph is the minimal number ℓ such that the graph can be embedded on an (orientable, resp. non-orientable) surface of genus ℓ . We write $g(G)$ for the orientable genus of G and $\tilde{g}(G)$ for the non-orientable genus of G . Every connected planar graph has orientable genus 0 and non-orientable genus 1. In general, for connected G , we have $\tilde{g}(G) \leq 2g(G) + 1$. On the other hand, there is no bound for $g(G)$ in terms of $\tilde{g}(G)$. As all our results apply to both variants, for ease of presentation, and as usual in the literature, we will simply speak of the genus of a graph in the following. We do not make explicit use of any topological arguments and hence refer to [26] for more background on graphs on surfaces. We will use the following facts about bounded genus graphs.

The first lemma states that graphs of genus g are closed under taking minors.

LEMMA 2.1. *If $H \leq G$, then $g(H) \leq g(G)$ and $\tilde{g}(H) \leq \tilde{g}(G)$.*

One of the arguments we will use is based on the fact that bounded genus graphs exclude large bipartite graphs as minors. The lemma follows immediately from Lemma 2.1 and from the fact that $g(K_{m,n}) = \left\lceil \frac{(m-2)(n-2)}{4} \right\rceil$ and $\tilde{g}(K_{m,n}) = \left\lceil \frac{(m-2)(n-2)}{2} \right\rceil$ (see e.g. Theorem 4.4.7 in [26]).

LEMMA 2.2. *If $g(G) = g$, then G excludes $K_{3,4g+3}$ as a minor and if $\tilde{g}(G) = \tilde{g}$, then G excludes $K_{3,2\tilde{g}+3}$ as a minor.*

Graphs of bounded genus do not contain many disjoint copies of minor models of $K_{3,3}$: this is a simple consequence of the fact that the orientable genus of a connected graph is equal to the sum of the genera of its blocks (maximal connected subgraphs without a cut-vertex) and a similar statement holds for the non-orientable genus, see Theorem 4.4.2 and Theorem 4.4.3 in [26].

LEMMA 2.3. *A graph G contains at most $\max\{g(G), 2\tilde{g}(G)\}$ disjoint copies of minor models of $K_{3,3}$.*

Finally, note that graphs of bounded genus have small edge density. It is straightforward to obtain the following from the generalized Euler formula $n - e + f \leq \chi(G)$ [26] for example see [2].

LEMMA 2.4. *Every graph with at least 3 vertices satisfies $|E(G)| \leq 3 \cdot |V(G)| + 6g(G) - 6$ and $|E(G)| \leq 3 \cdot |V(G)| + 3\tilde{g}(G) - 3$.*

LEMMA 2.5. *Let \mathcal{G} be a class of graphs of genus at most g . Then the degeneracy and edge density of every graph $G \in \mathcal{G}$ is bounded by $5\sqrt{g}$.*

PROOF. Recall that the degeneracy of a graph G is defined as $\max_{H \subseteq G} |E(H)|/|V(H)|$, which in particular bounds the edge density $|E(G)|/|V(G)|$. It hence suffices to bound the degeneracy of G .

If $g = 0$ the claim trivially holds, as in this case G is planar and hence $\max_{H \subseteq G} |E(H)|/|V(H)| \leq 3$ by Lemma 2.1 and Lemma 2.4.

Now assume $g \geq 1$ (we prove the lemma for graphs with orientable genus g , the proof for graphs of non-orientable genus g is analogous). We fix any subgraph $H \subseteq G$. We may assume that H has at least $5\sqrt{g}$ vertices, otherwise, the statement is trivially true (as in this case every vertex of H has degree (in H) less than $5\sqrt{g}$). By Lemma 2.1 and Lemma 2.4, we have $|E(H)| \leq 3 \cdot |V(H)| + 6g(H) - 6 \leq 3 \cdot |V(H)| + 6g$. This implies $|E(H)|/|V(H)| \leq 3 + 6g/|V(H)| \leq 3 + 6g/(5\sqrt{g}) \leq 5\sqrt{g}$, as claimed. \square

As an immediate corollary from Lemma 2.1, Lemma 2.2 and Lemma 2.4, we get that if \mathcal{G} is a class of graphs of bounded genus, then \mathcal{G} is a class of locally embeddable graphs.

Dominating sets. Let G be a graph. A set $D \subseteq V(G)$ *dominates* G if all vertices of G lie either in D or are adjacent to a vertex of D , that is, if $N[D] = V(G)$. A minimum dominating set D is a dominating set of minimum cardinality (among all dominating sets). The size of a minimum dominating set of G is denoted $\gamma(G)$.

f -Approximation. Let $f : \mathbb{N} \rightarrow \mathbb{R}^+$. Given an n -vertex graph G and a set $D \subseteq V(G)$, we say that D is an f -approximation for the dominating set problem, if D is a dominating set of G and $|D| \leq f(n) \cdot \gamma(G)$. An algorithm computes an f -approximation for the dominating set problem on a class \mathcal{C} of graphs if for all $G \in \mathcal{C}$ it computes a set D which is an f -approximation for the dominating set problem. If f maps every number to a fixed constant c , we speak of a constant factor approximation.

Distributed complexity. We consider the standard *LOCAL* model of distributed computing [24], see also [29] for a recent survey. A distributed system is modeled as a graph G . At each vertex $v \in V(G)$ there is an independent agent/host/processor with a unique identifier $id(v)$. Initially, each agent has no knowledge about the network, but only knows its own identifier. Information about other agents can be obtained through message passing, i.e., through repeated interactions with neighboring vertices, which happens in synchronous communication rounds. In each round the following operations are performed:

- (1) Each vertex performs a local computation (based on information obtained in previous rounds).
- (2) Each vertex v sends one message to each of its neighbors.
- (3) Each vertex v receives one message from each of its neighbors.

The *distributed complexity* of the algorithm is defined as the number of communication rounds until all agents terminate. We call a distributed algorithm r -local, if its output depends only on the r -neighborhoods $N^r[v]$ of its vertices. Observe that an r -local algorithm can (trivially) be implemented in r rounds in the *LOCAL* model.

3 A CONSTANT LOCAL MDS APPROXIMATION

Let us start by revisiting the MDS approximation algorithm for planar graphs by Lenzen et al. [21], see Algorithm 1. The algorithm works in two phases. In the first phase, it adds all vertices whose (open) neighborhood cannot be dominated by a small number of vertices (to be precise, by at most 6 vertices) to a set D . It has been shown in [21] that the set D is small (at most 4 times larger than a minimum dominating set) in planar graphs. In the second phase, the algorithm defines a dominator function dom which maps every vertex v that is not dominated yet by D to its dominator. The dominator $dom(v)$ of v is chosen arbitrary among those vertices of $N[v]$ which dominate the maximal number of vertices not dominated yet.

We now propose the following small change to the algorithm. As additional input, we require an integer c which bounds the edge density of depth-1 minors of G and we replace the condition $|A| \leq 6$

Algorithm 1 Dominating Set Approximation Algorithm for Planar Graphs [21]

```
1: Input: Planar graph  $G$ 
2: (* Phase 1 *)
3:  $D \leftarrow \emptyset$ 
4: for  $v \in V$  (in parallel) do
5:   if there does not exist a set  $A \subseteq V(G) \setminus \{v\}$  such that  $N(v) \subseteq N[A]$  and  $|A| \leq 6$  then
6:      $D \leftarrow D \cup \{v\}$ 
7:   end if
8: end for
9: (* Phase 2 *)
10:  $D' \leftarrow \emptyset$ 
11: for  $v \in V$  (in parallel) do
12:    $d_{G-D}(v) \leftarrow |N[v] \setminus N[D]|$ 
13:   if  $v \in V \setminus N[D]$  then
14:      $\Delta_{G-D}(v) \leftarrow \max_{w \in N[v]} d_{G-D}(w)$ 
15:     choose any  $dom(v)$  from  $N[v]$  with  $d_{G-D}(dom(v)) = \Delta_{G-D}(v)$ 
16:      $D' \leftarrow D' \cup \{dom(v)\}$ 
17:   end if
18: end for
19: return  $D \cup D'$ 
```

in Line 5 by the condition $|A| \leq 2c$. In the rest of this section, we show that the modified algorithm computes a constant factor approximation on any locally embeddable class of graphs. Note that the algorithm does not have to compute the edge density of G , which is not possible in a local manner. Rather, we leverage Lemma 2.4 which upper bounds the edge density for any fixed class of bounded genus graphs: this upper bound can be used as an input to the local algorithm.

We first show that the set D computed in Phase 1 of the algorithm is small. The following lemma is a straightforward generalization of Lemma 6.3 of [21], which in fact does not use topological arguments at all.

LEMMA 3.1. *Let G be a graph and let M be a minimum dominating set of G . Assume that for some constant c all depth-1 minors H of G satisfy $|E(H)|/|V(H)| \leq c$. Let*

$$D := \{v \in V(G) : \text{there is no set } A \subseteq V(G) \setminus \{v\} \text{ such that } N(v) \subseteq N[A] \text{ and } |A| \leq 2c\}.$$

Then $|D| \leq (c + 1) \cdot |M|$.

PROOF. Let H be the induced subgraph of G with $V(H) = M \cup N[D \setminus M]$. Since M is a dominating set, we can fix for each $v \in N[D \setminus M] \setminus (D \cup M)$ a vertex $m_v \in M$ that is adjacent to v . Then for each $m \in M$, the subgraph G_m which consists of the central vertex m and all $v \in N[D \setminus M] \setminus (D \cup M)$ such that $m = m_v$ and all edges $\{m, v\}$ is a star. Furthermore, observe that for different $m_1, m_2 \in M$ the stars G_{m_1} and G_{m_2} are vertex disjoint. We construct a depth-1 minor \tilde{H} of H by contracting the star subgraphs G_m for $m \in M$ into vertices v_m . Then (all non-trivial inequalities will be explained

below)

$$\begin{aligned}
(c+1) \cdot |D \setminus M| &= (2c+1) \cdot |D \setminus M| - c \cdot |D \setminus M| \\
&\leq \sum_{w \in D \setminus M} d_{\tilde{H}}(w) - |E(\tilde{H}[D \setminus M])| & (1) \\
&\leq |E(\tilde{H})| & (2) \\
&\leq c \cdot |V(\tilde{H})| & (3) \\
&= c \cdot (|D \setminus M| + |M|), & (4)
\end{aligned}$$

and hence $|D \setminus M| \leq c \cdot |M|$, which implies the claim.

- (1) Let $w \in D \setminus M$. As $N_G(w)$ cannot be covered by less than $(2c+1)$ elements from $V(G) \setminus \{w\}$ (by definition of D), w also has at least $(2c+1)$ neighbors in \tilde{H} . Hence $\sum_{w \in D \setminus M} d_{\tilde{H}}(w) \geq (2c+1) \cdot |D \setminus M|$. On the other hand, every subgraph \tilde{H}' of \tilde{H} has at most $c \cdot |V(\tilde{H}')|$ edges (every subgraph of a depth-1 minor is also a depth-1 minor of G and we assume that every depth-1 minor of G has edge density at most c). Hence $\tilde{H}[D \setminus M]$ has at most $c \cdot |D \setminus M|$ edges.
- (2) Every edge $\{v, w\} \in \tilde{H}$ with $v, w \in D \setminus M$ is counted twice in the sum $\sum_{w \in D \setminus M} d_{\tilde{H}}(w)$, once when we count $d_{\tilde{H}}(v)$ and once when counting $d_{\tilde{H}}(w)$. By subtracting the number of edges that run between vertices of $D \setminus M$ we get the second inequality.
- (3) The third inequality holds by assumption on the density of depth-1 minors of G .
- (4) By construction, all vertices of $N[D \setminus M] \setminus D$ disappear into some star G_m , hence \tilde{H} has exactly $|D \setminus M| + |M|$ vertices.

□

ASSUMPTION 3.2. *For the rest of this section, we fix a graph G which is locally embeddable, that is, G excludes $K_{3,t}$ for some t as depth-1 minor and all depth-1 minors H of G satisfy $|E(H)|/|V(H)| \leq c$ for some constant c (hence, Lemma 3.1 can be applied). Furthermore, we fix M and D as in Lemma 3.1.*

Let us write R for the set $V(G) \setminus N[D]$ of vertices which are not dominated by D . The algorithm defines a dominator function $dom : R \rightarrow N[R] \subseteq V(G) \setminus D$. The set D' computed by the algorithm is the image $dom(R)$, which is a dominating set of vertices in R . As R contains the vertices which are not dominated by D , $D' \cup D$ is a dominating set of G . This simple observation proves that the algorithm correctly computes a dominating set of G . Our aim is to find a bound on $|dom(R)|$.

We fix an ordering of M as $m_1, \dots, m_{|M|}$ such that the vertices of $M \cap D$ are first (minimal) in the ordering and inductively define a minimal set $E' \subseteq E(G)$ such that M is also a dominating set with respect to E' as follows. For $i = 1$, we add all edges $\{m_1, v\} \in E(G)$ with $v \in N(m_1) \setminus M$ to E' . If for some $i \geq 1$ we have defined the set of edges E' which are incident with m_1, \dots, m_i , we continue to add for $i+1$ all edges $\{m_{i+1}, v\} \in E(G)$ with $v \in N(m_{i+1}) \setminus (M \cup N_{E'}(\{m_1, \dots, m_i\}))$.

For $m \in M$, let G_m be the star subgraph of G with center m and all vertices v with $\{m, v\} \in E'$. Let H be the depth-1 minor of G which is obtained by contracting all stars G_m for $m \in M$. This construction is visualized in Figure 2. In the figure, solid (undirected) lines represent edges from E' , edges incident with $m \in M$ which are not in E' are dashed. We want to count the *endpoints* of directed edges, which represent the dominator function dom .

In the following, we call a directed edge which represents the function dom a *dom-edge*. We did not draw dom -edges that either start or end in M . When counting $|dom(R)|$, we may simply add a term $2|M|$ to estimate the number of endpoints of those edges. We also did not draw a dom -edge




Fig. 2. The graphs G_m . Solid (undirected) lines represent edges from E' , directed edges represent the dominator function dom . Dashed lines represent edges incident with $m \in M$ which are not in E' .

starting in G_{m_1} . In the figure, we assume that the vertex m_1 belongs to $M \cap D$. Hence every vertex v from $N[m_1]$ is dominated by a vertex from D and the function is thus not defined on such v . However, the vertices of $N(m_1)$ may still serve as dominators, as shown in the figure.

The graph H has $|M|$ vertices and by our assumption on the density of depth-1 minors of G , it has at most $c \cdot |M|$ edges.

Our analysis proceeds as follows. We distinguish between two types of dom -edges, namely those which go from one star to another star and those which start and end in the same star. By the star contraction, all edges which go from one star to another star are represented by a single edge in H . We show in Lemma 3.3 that each edge in H does not represent many such dom -edges with distinct endpoints. As H has at most $c \cdot |M|$ edges, we will end up with a number of such edges that is linear in $|M|$. On the other hand, all edges which start and end in the same star completely disappear in H . In Lemma 3.7 we show that these star contractions “absorb” only few such edges with distinct endpoints.

We first show that an edge in H represents only few dom -edges with distinct endpoints. For each $m \in M \setminus D$, we fix a set $C_m \subseteq V(G) \setminus \{m\}$ of size at most $2c$ which dominates $N_{E'}(m)$. The existence of such a set follows from the definition of the set D . Recall that we assume that G excludes $K_{3,t}$ as depth-1 minor.

LEMMA 3.3. *Let $1 \leq i < j \leq |M|$. Let $N_i := N_{E'}(m_i)$ and $N_j := N_{E'}(m_j)$.*

(1) *If $m_j \in M \setminus D$, then*

$$|\{u \in N_j : \text{there is } v \in N_i \text{ with } \{u, v\} \in E(G)\}| \leq 2ct.$$

(2) *If $m_i \in M \setminus D$ (and hence $m_j \in M \setminus D$), then*

$$|\{u \in N_i : \text{there is } v \in N_j \text{ with } \{u, v\} \in E(G)\}| \leq 4ct.$$

PROOF. By definition of E' , we may assume that $m_i \notin C_{m_j}$ (m_i is not connected to $N_{E'}(m_j)$ and hence it can be safely removed if it appears in C_{m_j}). Let $c \in C_{m_j}$ be arbitrary. Then there are at most $t - 1$ distinct vertices $u_1, \dots, u_{t-1} \in (N_j \cap N(c))$ such that there are $v_1, \dots, v_{t-1} \in N_i$ (possibly not distinct) with $\{u_k, v_k\} \in E(G)$ for all $k, 1 \leq k \leq t - 1$. Otherwise, we can contract the star with center m_i and branch vertices $N(m_i) \setminus \{c\}$ and thereby find $K_{3,t}$ as depth-1 minor, a contradiction. See Figure 3 for an illustration in the case of an excluded $K_{3,3}$. Possibly, $c \in N_j$ and it is connected to a vertex of N_i , hence we have at most t vertices in $N_j \cap N[c]$ with a connection to N_i . As $|C_{m_j}| \leq 2c$, we conclude the first item.

Regarding the second item, let $c \in C_{m_i}$ be arbitrary. If $c \neq m_j$, we conclude just as above, that there are at most $t - 1$ distinct vertices $u_1, \dots, u_{t-1} \in (N_i \cap N(c))$ such that there are $v_1, \dots, v_{t-1} \in N_j$ (possibly not distinct) with $\{u_k, v_k\} \in E(G)$ for all $k, 1 \leq k \leq t - 1$ and hence at most t vertices in $N_i \cap N[c]$ with a connection to N_j . Now assume $c = m_j$. Let $c' \in C_{m_j}$. There are at most $t - 1$ distinct vertices $u_1, \dots, u_{t-1} \in (N_i \cap N_{E'}(m_j))$ such that there are vertices $v_1, \dots, v_{t-1} \in N_j \cap N(c)$

(possibly not distinct) with $\{u_k, v_k\} \in E(G)$ for all k , $1 \leq k \leq t-1$. Again, considering the possibility that $c' \in N_i$, there are at most t vertices in $N_i \cap N_E(m_j)$ with a connection to $N_j \cap N(c)$. As $|C_{m_j}| \leq 2c$, we conclude that in total there are at most $2ct$ vertices in $N_i \cap N_E(m_j)$ with a connection to N_j . In total, there are hence at most $(2c-1)t + 2ct \leq 4ct$ vertices of the described form. \square




Fig. 3. Visualisation of the proof of Lemma 3.3 in the case of excluded $K_{3,3}$

We write Y for the set of all vertices $\{u \in N_{E'}(m_i) : m_i \notin D \text{ and there is } v \in N_{E'}(m_j), j \neq i \text{ and } \{u, v\} \in E(G)\}$.

COROLLARY 3.4. $|Y| \leq 6c^2t \cdot |M|$.

PROOF. Each of the $c \cdot |M|$ many edges in H represents edges between N_i and N_j , where N_i and N_j are defined as above. By the previous lemma, if $i < j$, there are at most $2ct$ vertices in $N_i \cap Y$ and at most $4ct$ vertices in $N_j \cap Y$. Hence in total, each edge accounts for at most $6ct$ vertices in Y . \square

We continue to count the edges which are inside the stars. First, we show that every vertex has small degree inside its own star.

LEMMA 3.5. *Let $m \in M \setminus D$ and let $v \in N_{E'}(m) \setminus C_m$. Then*

$$|\{u \in N_{E'}(m) : \{u, v\} \in E(G)\}| \leq 2c(t-1).$$

PROOF. Let $c \in C_m$. By the same argument as in Lemma 3.3, there are at most $t-1$ distinct vertices $u_1, \dots, u_{t-1} \in (N_{E'}(m) \cap N(c))$ such that $\{u_k, v\} \in E(G)$ for all k , $1 \leq k \leq t-1$. \square

Let $C := \bigcup_{m \in M \setminus D} C_m$. We show that there are only few vertices which are highly connected to $M \cup C$. Let $Z := \{u \in N_{E'}(M \setminus D) : |N(u) \cap (M \cup C)| > 4c\}$.

LEMMA 3.6.

$$|Z| < |M \cup C|.$$

PROOF. Assume that $|Z| > |M \cup C|$. Then the subgraph induced by $Z \cup M \cup C$ has more than $\frac{1}{2}4c|Z|$ edges and $|Z \cup M \cup C|$ vertices. Hence its edge density is larger than $2c|Z|/(|Z \cup M \cup C|) > 2c|Z|/(2|Z|) = c$, contradicting our assumption on the edge density of depth-1 minors of G (which includes its subgraphs). \square

Finally, we consider the image of the *dom*-function inside the stars.

LEMMA 3.7.

$$\left| \bigcup_{m \in M \setminus D} \{u \in N_{E'}(m) : \text{dom}(u) \in (N_{E'}(m) \setminus (Y \cup Z))\} \right| \leq (2(t-1) + 4)c|M|.$$

PROOF. Fix some $m \in M \setminus D$ and some $u \in N_{E'}(m)$ with $\text{dom}(u) \in N_{E'}(m) \setminus (Y \cup Z)$. Because $\text{dom}(u) \notin Y$, $\text{dom}(u)$ is not connected to a vertex of a different star, except possibly for vertices from M . Because $\text{dom}(u) \notin Z$, it is however connected to at most $4c$ vertices from $M \cup C$. Hence it is connected to at most $4c$ vertices from different stars. According to Lemma 3.5, $\text{dom}(u)$ is connected to at most $2c(t-1)$ vertices from the same star. Hence the degree of $\text{dom}(u)$ is at most $4c + 2c(t-1)$. Because u preferred to choose $\text{dom}(u) \in N_{E'}(m)$ over m as its dominator, we conclude that m has at most $4c + 2c(t-1)$ E' -neighbors. Hence, in total there can be at most $(2(t-1) + 4)c \cdot |M|$ such vertices. \square

We are now ready to put together the numbers.

LEMMA 3.8. *If all depth-1 minors H of G have edge density at most c and G excludes $K_{3,t}$ as depth-1 minor, then the modified algorithm computes a $6c^2t + (2t + 5)c + 4$ approximation for the minimum dominating set problem on G .*

PROOF. Since M is a dominating set also with respect to the edges E' , it suffices to bound $|\{\text{dom}(u) : u \in (N_{E'}[M \setminus D] \setminus N[D])\}|$. This set is partitioned into the following (not necessarily disjoint) sets. First, all endpoints of $\text{dom}(R)$ that go from one star to another star are found in one of the sets $Y = \{u \in N_{E'}(m_i) : \text{there is } v \in N_{E'}(m_j), i \neq j \text{ and } \{u, v\} \in E(G)\}$, $\text{dom}(R) \cap M$ and $\text{dom}(M)$. All other dom-edges connect vertices inside individual stars. Here, $\text{dom}(R)$ splits into those vertices which are highly connected to $M \cup C$, that is, the set $Z = \{u \in N_{E'}(M \setminus D) : |N(u) \cap (M \cup C)| > 4c\}$, the set C and the set Y (which will not be counted twice though). All other dom-edges lead to vertices which lie neither in Y nor in Z .

In the previous lemmas we have bounded the sizes of each of the described sets. The set D has size at most $(c+1)|M|$ according to Theorem 3.1. According to Corollary 3.4, the set Y has size at most $6c^2t|M|$. In particular, there are at most so many vertices $\text{dom}(u) \in N_{E'}(m_i)$ with $u \in N_{E'}(m_j)$ for $i \neq j$. Clearly, $|\text{dom}(R) \cap M| \leq |M|$ and $|\text{dom}(M)| \leq |M|$. According to Lemma 3.6, the set Z satisfies $|Z| < |M \cup C|$. We have $|C| \leq 2c|M|$, as each C_m has size at most $2c$. It remains to count the image of dom inside the stars which do not point to Y or Z . According to Lemma 3.7, this image has size at most $(2(t-1) + 4)c|M|$. In total, we can bound $|\text{dom}(R)|$ by

$$(c+1)|M| + 6c^2t|M| + 2|M| + (2c+1)|M| + (2(t-1) + 4)c|M| \leq (6c^2t + (2t+5)c + 4)|M|.$$

\square

Our theorem for bounded genus graphs is now a corollary of Lemma 2.2, 2.5 and 3.8.

THEOREM 3.9. *Let C be a class of graphs of orientable genus at most g (non-orientable genus at most \tilde{g} resp.). The modified algorithm computes an $O(g^2)$ -approximation ($O(\tilde{g}^2)$ -approximation resp.) for the dominating set in a constant number of communication rounds.*

For the special case of planar graphs, our analysis shows that the algorithm computes a 199-approximation. This is not much worse than Lenzen et al.'s original analysis (130), however, off by a factor of almost 4 from Wawrzyniak's [31] improved analysis (52).

Algorithm 2 Dominating Set Approximation for Graphs of Genus $\leq g$

```
1: Input: Graph  $G$  of genus at most  $g$ 
2: Run Phase 1 of Modified Algorithm 1 with density parameter  $10\sqrt{g}$  to obtain set  $D$ 
3: (* Preprocessing Phase *)
4: for  $v \in V - D$  (in parallel) do
5:   compute  $K_v$  in  $G - D$  (see Lemma 3.10)
6: end for
7: for  $i = 1..g$  do
8:   for  $v \in V - D$  (in parallel) do
9:     if  $K_v \neq \emptyset$  then chosen := true
10:    for all  $u \in N^{12}(v)$  do
11:      if  $K_u \cap K_v \neq \emptyset$  and  $u < v$  then chosen := false
12:    end for
13:    if (chosen = true) then  $D := D \cup V(K_v)$ 
14:  end for
15: end for
16: Run Phase 2 of Algorithm 1
```

3.1 Improving the Approximation Factor with Preprocessing

We now show the approximation factors related to the genus g , derived in the previous section, can be improved using a local preprocessing step.

Given a graph G and a vertex $v \in V(G)$, let $K = \{K_1, \dots, K_j\}$ denote the set of minimal subgraphs of G containing v such that for all $1 \leq i \leq j$, $K_{3,3}$ is a depth-1 minor of K_i . Let $K_h \in K$ be the one with lexicographically smallest identifiers in K . We call K_h the v -canonical subgraph of G and we denote it by K_v . If $K = \emptyset$ we set $K_v := \emptyset$.

LEMMA 3.10. *Given a graph G and a vertex $v \in V(G)$. The v -canonical subgraph K_v of v can be computed locally in at most 6 communication rounds. Furthermore, K_v has at most 24 vertices.*

PROOF. The proof is constructive. As $K_{3,3}$ has diameter 2, every minimal subgraph of G containing $K_{3,3}$ as a depth-1 minor has diameter at most 6 (every edge may have to be replaced by a path of length 3). Therefore, it suffices to consider the subgraph induced by the vertices at distance at most 6 from v , $H = G[N^6(v)]$, and find the lexicographically minimal subgraph which contains $K_{3,3}$ as depth-1 minor in H which includes v as a vertex. If this is the case, we output it as K_v ; otherwise we output the empty set. Furthermore, $K_{3,3}$ has 9 edges and hence a minimal subgraph containing it as depth-1 minor has at most 24 vertices (again, every edge is subdivided at most twice and $2 \cdot 9 + 6 = 24$). \square

To improve the approximation factor, we propose the following modified algorithm, see Algorithm 2. We first carry out the first phase of Algorithm 1 with density parameter $10\sqrt{g}$ (the parameter is twice the edge density of the input graph). In the following preprocessing phase we eliminate all copies of depth-1 minor models of $K_{3,3}$ that G possibly contains. By Lemma 2.3 we know that there are at most g (where g is the genus of the graph) disjoint such models. As guaranteed by Lemma 3.10, the vertices can make a canonical local choice on which model to delete. After g elimination rounds we are left with a locally embeddable graph (with the parameter $t = 3$) and we call the second phase of Algorithm 1.

THEOREM 3.11. *Algorithm 2 provides a $24g + O(1)$ MDS approximation for graphs of genus at most g , and requires $12g + O(1)$ communication rounds.*

PROOF. The resulting vertex set is clearly a dominating set. It remains to bound its size.

As Phase 1 is unchanged, the computed set D is at most $6\sqrt{g}$ times larger than an optimal dominating set by Lemma 2.4: the algorithm is called with parameter $2c$ and outputs a set at most $c + 1$ times larger than an optimal dominating set; here, $c = 5\sqrt{g}$ according to Lemma 2.5.

In the following preprocessing phase, if for two vertices $u \neq v$ we choose both K_u and K_v , then they must be disjoint: Since the diameter of any depth-1 minor of $K_{3,3}$ is at most 6, if two such canonical subgraphs K_u and K_v intersect, then the distance between u, v can be at most 12. Hence, each vertex v can decide in its 12-neighborhood whether its canonical subgraph K_v is the smallest among all choices. On the other hand, by Lemma 2.3, there are at most g disjoint such minor models. So in the *preprocessing* phase, we can remove at most g disjoint subgraphs K_v (and add their vertices to the dominating set) and thereby select at most $24g$ extra vertices for the dominating set. Once the *preprocessing* phase is finished, the remaining graph is locally embeddable. Observe that if the input graph G is planar, no vertices will be added to D in the preprocessing phase.

In order to compute the size of the set in the third phase, we can use the analysis of Lemma 3.8 for $t = 3$, which together with the first phase and preprocessing phase, results in a $24g + O(1)$ -approximation guarantee.

To count the number of communication rounds, note that the only change happens in the second phase. In that phase, in each iteration, we need 12 communication rounds to compute the 12-neighborhood. Therefore, the number of communication rounds is $12g + O(1)$. \square

This significantly improves the approximation upper bound of Theorem 3.9: namely from $4(6c^2 + 2c)g + O(1)$, where $c = O(\sqrt{g})$, hence from $O(g^2)$ to $24g + O(1)$, at the price of $12g$ extra communication rounds.

3.2 A Logical Perspective

Interestingly, as we will elaborate in the following, a small modification of Algorithm 1 can be interpreted both from a distributed computing perspective, namely as a local algorithm of constant distributed time complexity, as well as from a logical perspective.

First order logic has atomic formulas of the form $x = y$, $x < y$ and $E(x, y)$, where x and y are first-order variables and E is a binary relation symbol. The set of first order formulas is closed under Boolean combinations and existential and universal quantification over the vertices of a graph. To define the semantics, we inductively define a satisfaction relation \models , where for a graph G , a formula $\phi(x_1, \dots, x_k)$ and vertices $v_1, \dots, v_k \in V(G)$, $G \models \phi(v_1, \dots, v_k)$ means that G satisfies ϕ if the free variables x_1, \dots, x_k are interpreted as v_1, \dots, v_k , respectively. The free variables of a formula are those that have an instance not in the scope of a quantifier, and we write $\phi(x_1, \dots, x_k)$ to indicate that the free variables of the formula ϕ are among x_1, \dots, x_k . For $\phi(x_1, x_2) = x_1 < x_2$, we have $G \models \phi(v_1, v_2)$ if $v_1 < v_2$ with respect to the ordering $<$ of $V(G)$ and for $\phi(x_1, x_2) = E(x_1, x_2)$ we have $G \models \phi(v_1, v_2)$ if $\{v_1, v_2\} \in E(G)$. The meaning of the equality symbol, the Boolean connectives, and the quantifiers is as expected.

A first-order formula $\phi(x)$ with one free variable naturally defines the set $\phi(G) = \{v \in V(G) : G \models \phi(v)\}$. We say that a formula ϕ defines an f -approximation to the dominating set problem on a class \mathcal{C} of graphs, if $\phi(G)$ is an f -approximation of a minimum dominating set for every graph $G \in \mathcal{C}$.

Observe that first-order logic is not able to count, in particular, no fixed formula can determine a neighbor of maximum degree in Line 14 of the algorithm. Also note however that the only place in

our analysis which refers to the dominator function dom explicitly is Lemma 3.7. The proof of the lemma in fact shows that we do not have to choose a vertex of maximal residual degree, but that it suffices to choose a neighbour of degree greater than $4c + 2c(t - 1)$ if such a vertex exists, or any vertex, otherwise. For every fixed class of bounded genus, this number is a constant. We use the binary predicate $<$ to make a unique choice of a dominator in this case.

Then we define D by the following formula

$$\varphi_D(x) = \neg(\exists x_1 \dots \exists x_{2c} \forall y \left(E(x, y) \rightarrow \bigvee_{1 \leq i \leq 2c} E(y, x_i) \right))$$

and D' by

$$\psi_{D'}(x) = \exists y \left(E(x, y) \wedge \forall z (\varphi_D(z) \rightarrow \neg E(y, z)) \wedge \xi_{\max}(x, y) \right),$$

where $\xi_{\max}(x, y)$ states that x is the maximum (residual) degree neighbour of y up to threshold $4c + 2c(t - 1)$. We can express this cumbersome formula with $4c + 2c(t - 1)$ quantifiers. Note that the formulas φ_D and $\psi_{D'}$ are different in spirit. While φ_D directly describes a property of vertices which causes them to be included in the dominating set, in the formula $\psi_{D'}(x)$ we state the existence of an element which is not yet dominated by D and which elects x as a dominator.

4 $(1 + \epsilon)$ -APPROXIMATIONS

In this section we show how to extend techniques developed by Czygrinow et al. [9] to find $(1 + \epsilon)$ -approximate dominating set for planar graphs to graphs of sub-logarithmic expansion. These graphs are very general classes of sparse graphs, including planar graphs and all classes that exclude a fixed minor. We focus on the dominating set problem, however, the approximations for the maximum weight independent set problem and maximum matching problem proposed by Czygrinow et al. can be extended in a similar way.

Our notation in this section closely follows that of Czygrinow et al. [9]. In particular, we will work with vertex and edge weighted graphs, that is, every graph G is additionally equipped with two weight functions $\omega : V(G) \rightarrow \mathbb{R}^+$ and $\bar{\omega} : E(G) \rightarrow \mathbb{R}^+$. If $H \subseteq G$ is a subgraph of G , then we write $\omega(H)$ for $\sum_{v \in V(H)} \omega(v)$ and $\bar{\omega}(H)$ for $\sum_{e \in E(H)} \bar{\omega}(e)$. If $\{G_1, \dots, G_n\}$ is a minor model of a graph H in a weighted graph G , then H naturally inherits a weight function ω_H from G as follows. If $u, v \in V(H)$ are represented by the subgraphs G_u and G_v in the minor model, then $\omega(u) = \sum_{w \in V(G_u)} \omega(w)$ and if $\{u, v\} \in E(H)$, then $\bar{\omega}_H(\{u, v\}) = \sum_{e \in E(G), e \cap V(G_u) \neq \emptyset, e \cap V(G_v) \neq \emptyset} \bar{\omega}(e)$.

4.1 Clustering Algorithm

We first generalize the partitioning algorithm provided by Czygrinow et al. to graphs with sub-logarithmic expansion.

Definition 4.1 (Pseudo-Forest). A pseudo-forest is a directed graph \vec{F} in which every vertex has out-degree at most 1.

For a directed graph \vec{F} , we write F for the underlying undirected graph of \vec{F} . The following lemma is a straightforward generalization of Fact 1 of [9].

LEMMA 4.2. *Let G be a graph of arboricity a with an edge-weight function $\bar{\omega}$. There is a distributed procedure which in two rounds finds a pseudo-forest \vec{F} such that F is a subgraph of G with $\bar{\omega}(F) \geq \bar{\omega}(G)/(2a)$.*

PROOF. We run the following algorithm. For every vertex v , we choose one edge $\{v, w\}$ of largest weight, and direct it from v to w . If we happen to choose an edge $\{v, w\}$ for both vertices v and

w , we direct it from v to w , using the larger identifier as a tie breaker. Hence every vertex has out-degree at most one and the algorithm outputs a pseudo-forest \vec{F} .

Let us show that $\bar{\omega}(F) \geq \bar{\omega}(G)/(2a)$. Without loss of generality we assume that G has no isolated vertices (we make a statement about edge weights only). As G has arboricity at most a , there exists a forest cover \mathcal{F} into at most a forests. So one of the forests $T \in \mathcal{F}$ collects weight $\bar{\omega}(T) \geq \bar{\omega}(G)/a$. Now associate with each vertex v of T the value $w_T(v)$ which is the weight of the edge connecting it to its parent (if it exists). Similarly, write $w_F(v)$ for the weight of the arc (v, w) or (w, v) in \vec{F} that was chosen in the algorithm. Observe that we may be double counting edges here, but only once. Hence we have $\bar{\omega}(F) \geq \sum_{v \in V(G)} w_F(v)/2 \geq \sum_{v \in V(T)} w_F(v)/2 \geq \sum_{v \in V(T)} w_T(v)/2 \geq \bar{\omega}(G)/(2a)$. \square

It is straightforward to generalize also Lemma 2 of [9].

LEMMA 4.3 (HEAVYSTAR). *There is a local algorithm which takes an edge weighted n -vertex graph G of arboricity at most a as input and in $O(\log^* n)$ rounds computes a partition of $V(G)$ into vertex disjoint stars $H_1, \dots, H_x \subseteq V(G)$ such that $H = G/H_1/\dots/H_x$ has total weight $\bar{\omega}_H(H) \leq (1 - 1/(8a)) \cdot \bar{\omega}(G)$.*

We refrain from presenting a proof of this lemma, as the proof is literally a copy of the proof given in [9]. Czygrinow et al. [9] use only the fact that planar graphs have arboricity 3, while we make the statement for graphs of arboricity a . Hence only numbers must be adapted in the proof (from 24 in their work to $8a$ in our case). We refer the reader to the very accessible presentation in [9].

We come to the final clustering algorithm. We fix a function $f(r) \in o(\log r)$ which bounds the expansion (density of depth- r minors) of the input graphs G . Recall that arboricity is within factor 2 of density of subgraphs, hence the depth- r minors of G have arboricity bounded by $2f(r)$. By iteratively taking depth-1 minors, we obtain minors at exponential depth, as stated in the next lemma.

LEMMA 4.4 (PROPOSITION 4.1, STATEMENT (4.4) OF [27]). *Taking a depth-1 minor of a depth-1 minor for r times gives a depth- $(3^r - 1)/2$ minor of G .*

In particular, when iterating the star contraction routine of Algorithm 3, in iteration i we are dealing with a subgraph of arboricity $2f((3^i - 1)/2) =: g(i)$ which is sublinear in i . Hence, we may apply Lemma 4.3 with arboricity parameter i in iteration i . Note that we do not require the arboricity as an input to the algorithm of Lemma 4.3. Note furthermore, that we have

$$\lim_{i \rightarrow \infty} (1 - 1/(8g(i)))^i \leq \lim_{i \rightarrow \infty} e^{-i/g(i)} = 0,$$

hence for every $\epsilon > 0$ there is a constant i_0 depending only on ϵ and g (and not on the graph size n) such that $(1 - 1/(8g(i_0)))^{i_0} \leq \epsilon$ (we may assume that the function g is monotone, as the density of depth- r minors cannot be smaller than the density of depth- r' minors for $r' \leq r$).

Algorithm 3 Clustering

- 1: Input: G with $\nabla_r(G) \leq f(r)$, $\epsilon > 0$ and i_0 with $(1 - 1/(8g(i_0)))^{i_0} \leq \epsilon$
 - 2: **for** $i = 1, \dots, i_0$
 - 3: Call the algorithm of Lemma 4.3 to find vertex disjoint stars H_1, \dots, H_x in G
 - 4: $H \leftarrow G/H_1/\dots/H_x$ with weights modified accordingly
 - 5: **end for**
 - 6: **return** $\{C_i = V(H_i) : 1 \leq i \leq x\}$.
-

LEMMA 4.5 (CLUSTERING). *Let $c \geq 1$ be a constant and let G be a graph with $\nabla_r(G) \leq f(r)$. If the clustering algorithm gets G and $\epsilon > 0$ as input, then it returns a set of clusters C_1, \dots, C_x partitioning G , such that, each cluster has radius at most $(3^{i_0} - 1)/2$ (where i_0 is the number of iterations in the algorithm). Moreover, if we contract each C_i to a single vertex to obtain a graph H , then $\bar{\omega}_H(H) \leq \epsilon \cdot \bar{\omega}(G)$. The algorithm uses $O_\epsilon(\log^* n)$ communication rounds.*

In the above lemma we use the notation O_ϵ to express that we are treating all constants depending on ϵ as constants.

PROOF. As described above, the graph G_i we are dealing with in iteration i has arboricity at most $2f(3^i - 1)/2 = g(i)$, which is sublinear in i . By applying Lemma 4.3 to G_i , we compute in $O(\log^* n)$ rounds a partition of $V(G_i)$ into vertex disjoint stars $H_1, \dots, H_x \subseteq V(G_i)$ such that $H = G_i/H_1/\dots/H_x$ has total weight $\bar{\omega}_H(H) \leq (1 - 1/(8g(i))) \cdot \bar{\omega}(G_i)$. Note that by Lemma 4.4, the graph G_i obtained in iteration i is a depth- $((3^i - 1)/2)$ minor of G . Hence, by induction, after i iterations, the edge weight of the graph G_i is at most $(1 - 1/(8g(i)))^i$. As argued above, there exists i_0 such that $(1 - 1/(8g(i_0)))^{i_0} \leq \epsilon$, at which time we stop the algorithm.

As in each round we invest at most time $O(\log^* n)$, in total we invest at most $O(\log^* n \cdot i_0) = O_\epsilon(\log^* n)$ time to compute the clustering. \square

4.2 Approximation for Minimum Dominating Set

We are ready to prove the main theorem of this section.

THEOREM 4.6. *There exists a deterministic distributed algorithm which gets as input*

- (1) *an n -vertex graph G of sub-logarithmic expansion,*
- (2) *a c -approximation of a minimum dominating set D of G for some constant c , and*
- (3) *a real parameter $\epsilon > 0$.*

The algorithm runs in $O_{\epsilon,c}(\log^ n)$ rounds and outputs a $(1 + \epsilon)$ -approximation of a minimum dominating set of G .*

COROLLARY 4.7. *Let \mathcal{C} be a class of graphs of sub-logarithmic expansion. Assume there exists an algorithm which computes c -approximations of dominating sets on graphs from \mathcal{C} in t rounds. Then there exists an algorithm which for every $\epsilon > 0$ computes a $(1 + \epsilon)$ -approximation of a minimum dominating set on every n -vertex graph $G \in \mathcal{C}$ in $O_{\epsilon,c}(t + \log^* n)$ rounds.*

We have chosen to present this extension of Czygrinow et al. [9] because it connects very well to the results we obtained in the previous section. In particular, Corollary 4.7 in combination with Theorem 3.9 gives a deterministic distributed $(1 + \epsilon)$ -approximation algorithm in $O_{\epsilon,g}(\log^* n)$ rounds for dominating sets on graphs of genus at most g . We can similarly combine the corollary the result of Amiri et al. [1] to obtain $(1 + \epsilon)$ -approximations in $O(\log n)$ rounds on graphs of sub-logarithmic expansion.

PROOF OF THEOREM 4.6. Let G be the given input graph and let D be a dominating set of G with $|D| \leq c \cdot \gamma(G)$, say $D = \{w_1, \dots, w_k\}$ (recall that $\gamma(G)$ denotes the size of a minimum size dominating set of G). Associate each vertex $v \in V(G) \setminus D$ with one of its dominators, say with the one of minimum identifier, to obtain a partition (W_1, \dots, W_k) of G into clusters of radius 1. This partition is obtained in a single communication round. The graph $H = G/W_1/\dots/W_k$ is a depth-1 minor of G with k vertices and at most $\nabla_1(G) \cdot k$ edges. Define an edge weight function on $E(H)$ by assigning unit weight to each edge. Set $\delta = \epsilon/(2c\nabla_1(G))$. Apply the algorithm of Lemma 4.5 with parameter δ to find a partition (V_1, \dots, V_l) of $V(H)$ such that the weight between different clusters is at most $\delta \cdot |E(H)|$. The algorithm runs in $O_\delta(\log^* n) = O_{\epsilon,c}(\log^* n)$ communication

rounds. By uncontracting the partitions V_i and W_i , we obtain a partition (U_1, \dots, U_l) of $V(G)$, where each U_i has constant radius. Find an optimal dominating set S_i of $G[U_i]$ in each subgraph $G[U_i]$ and return the union $S = \bigcup_{1 \leq i \leq l} S_i$ of these dominating sets. As the algorithm has already learned the subgraphs $G[U_i]$, by the infinite computational power of each processor in the *LOCAL* model, we can compute such a dominating set in one round. This completes the description of the algorithm.

Note that instead of solving the dominating set optimally on each $G[U_i]$, which may be considered an abuse of the *LOCAL* model by some, we can compute a sufficiently good approximation of an optimal dominating set. For this, we can use the PTAS [16] for dominating sets on graphs of polynomial expansion.

Since the U_i form a partition of $V(G)$, it is clear that S is a dominating set of G . Denote by S^* a dominating set of cardinality Opt . Let S' be obtained from S^* by adding for each U_i all vertices $w \in U_i$ which have a neighbor in a different cluster U_j . Then $S' \cap U_i$ is a dominating set of $G[U_i]$. Furthermore, we have

$$|S'| \leq |S^*| + 2\delta|E(H)| \leq \gamma(G) + 2c\delta\nabla_1(G) \cdot \gamma(G) = (1 + \epsilon) \cdot \gamma(G).$$

Observe that the local solutions S_i cannot be worse than the solution $S' \cap U_i$, hence

$$|S| = \sum_{1 \leq i \leq l} |S_i| \leq \sum_{1 \leq i \leq l} |S' \cap U_i| = |S'| \leq (1 + \epsilon) \cdot \gamma(G). \quad \square$$

5 CONCLUSION

This paper presented the first constant round, constant factor local MDS approximation algorithm for locally embeddable graphs, a class of graphs which is more general than planar graphs. Moreover, we have shown how our result can also be used to derive a $O(\log^* n)$ -time distributed approximation scheme for bounded genus graphs.

Our proofs are purely combinatorial and avoid all topological arguments. For the family of bounded genus graphs, topological arguments helped to improve the obtained approximation ratio in a preprocessing step. We believe that this result constitutes a major step forward in the quest for understanding for which graph families such local approximations exist. Besides the result itself, we believe that our analysis introduces several new techniques which may be useful also for the design and analysis of local algorithms for more general graphs, and also problems beyond MDS. In particular, we believe that the notion of bounded depth minors and not the commonly used notion of excluded minors will be the right notions in the setting of local, distributed computing.

Moreover, this paper established an interesting connection between distributed computing and logic, by presenting a local approximation algorithm which is first-order logic definable. This also provides an interesting new perspective on the recently introduced notion of stone-age distributed computing [11]: distributed algorithms making minimal assumptions on the power of a node. Avoiding counting in the arising formulas allows for example an implementation of the algorithm in the circuit complexity class AC^0 , that is, an implementation by circuits of polynomial size and constant depth.

It remains open whether the local constant approximation result can be generalized to sparse graphs beyond bounded genus graphs. Also, it will be interesting to extend our study of first-order definable approximations.

REFERENCES

- [2] Saeed Akhondian Amiri. 2017. *Structural graph theory meets algorithms: covering and connectivity problems in graphs*. Ph.D. Dissertation. Logic and Semantic Group, Technical University Berlin.

- [1] Saeed Akhoondian Amiri, Patrice Ossona de Mendez, Roman Rabinovich, and Sebastian Siebertz. 2018. Distributed domination on graph classes of bounded expansion. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures*. ACM, 143–151.
- [3] Saeed Akhoondian Amiri and Stefan Schmid. 2016. Brief Announcement: A Log-Time Local MDS Approximation Scheme for Bounded Genus Graphs. In *Proc. 30th International Symposium on Distributed Computing (DISC)*. Springer.
- [4] Saeed Akhoondian Amiri, Stefan Schmid, and Sebastian Siebertz. 2016. A Local Constant Factor MDS Approximation for Bounded Genus Graphs. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*.
- [5] Brenda S. Baker. 1994. Approximation Algorithms for NP-complete Problems on Planar Graphs. *J. ACM* 41, 1 (Jan. 1994), 153–180.
- [6] Nikhil Bansal and Seun William Umboh. 2017. Tight approximation bounds for dominating set on graphs of bounded arboricity. *Inform. Process. Lett.* 122 (2017), 21–24.
- [7] Leonid Barenboim, Michael Elkin, and Cyril Gavoille. 2014. A Fast Network-Decomposition Algorithm and its Applications to Constant-Time Distributed Computation. In *Proc. SIROCCO*. 209–223.
- [8] Hervé Brönnimann and Michael T Goodrich. 1995. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry* 14, 4 (1995), 463–479.
- [9] Andrzej Czygrinow, Michał Hańcowski, and Wojciech Wawrzyniak. 2008. Fast Distributed Approximations in Planar Graphs. In *Proc. 22nd International Symposium on Distributed Computing (DISC)*. 78–92.
- [10] Irit Dinur and David Steurer. 2014. Analytical approach to parallel repetition. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. ACM, 624–633.
- [11] Yuval Emek and Roger Wattenhofer. 2013. Stone Age Distributed Computing. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*. 137–146.
- [12] David Eppstein. 2000. Diameter and treewidth in minor-closed graph families. *Algorithmica* 27, 3-4 (2000), 275–291.
- [13] Guy Even, Dror Rawitz, and Shimon Moni Shahar. 2005. Hitting sets when the VC-dimension is small. *Inform. Process. Lett.* 95, 2 (2005), 358–362.
- [14] Michael R. Garey and David S. Johnson. 1979. Computers and intractability: a guide to the theory of NP-completeness. *WH Free. Co., San Fr* (1979).
- [15] Martin Grohe. 2003. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica* 23, 4 (2003), 613–632.
- [16] Sarel Har-Peled and Kent Quanrud. 2015. Approximation algorithms for polynomial-expansion and low-density graphs. In *Algorithms-ESA 2015*. Springer, 717–728.
- [17] Miikka Hilke, Christoph Lenzen, and Jukka Suomela. 2014. Brief announcement: local approximability of minimum dominating set on planar graphs. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*. 344–346.
- [18] David S Johnson. 1974. Approximation algorithms for combinatorial problems. *Journal of computer and system sciences* 9, 3 (1974), 256–278.
- [19] Richard M Karp. 1972. Reducibility among combinatorial problems. In *Complexity of computer computations*. Springer, 85–103.
- [20] Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. 2016. Local Computation: Lower and Upper Bounds. *J. ACM* 63, 2, Article 17 (March 2016), 44 pages.
- [21] Christoph Lenzen, Yvonne Anne Pignolet, and Roger Wattenhofer. 2013. Distributed minimum dominating set approximations in restricted families of graphs. *Distributed Computing* 26, 2 (2013), 119–137.
- [22] Christoph Lenzen and Roger Wattenhofer. 2008. Leveraging Linial’s Locality Limit. In *Distributed Computing*. Lecture Notes in Computer Science, Vol. 5218. 394–407.
- [23] Christoph Lenzen and Roger Wattenhofer. 2010. Minimum Dominating Set Approximation in Graphs of Bounded Arboricity. In *Distributed Computing*. Vol. 6343. 510–524.
- [24] Nathan Linial. 1992. Locality in Distributed Graph Algorithms. *SIAM J. Comput.* 21, 1 (Feb. 1992), 193–201.
- [25] László Lovász. 1975. On the ratio of optimal integral and fractional covers. *Discrete mathematics* 13, 4 (1975), 383–390.
- [26] Bojan Mohar and Carsten Thomassen. 2001. *Graphs on Surfaces*. Johns Hopkins University Press.
- [27] Jaroslav Nešetřil and Patrice Ossona De Mendez. 2012. *Sparsity: graphs, structures, and algorithms*. Vol. 28. Springer Science & Business Media.
- [28] Sebastian Siebertz. 2019. Greedy domination on biclique-free graphs. *Inform. Process. Lett.* (2019).
- [29] Jukka Suomela. 2013. Survey of Local Algorithms. *ACM Comput. Surv.* 45, 2, Article 24 (March 2013), 24:1–24:40 pages.
- [30] Wojciech Wawrzyniak. 2013. Brief announcement: a local approximation algorithm for MDS problem in anonymous planar networks. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*. ACM, 406–408.
- [31] Wojciech Wawrzyniak. 2014. A strengthened analysis of a local algorithm for the minimum dominating set problem in planar graphs. *Proc. Information Processing Letters (IPL)* 114, 3 (2014), 94 – 98.