

A Local Constant Factor MDS Approximation for Bounded Genus Graphs

Saeed Amiri*
TU Berlin, Germany
saeed.amiri@tu-berlin.de

Stefan Schmid
Aalborg University, Denmark
& TU Berlin, Germany
schmiste@cs.aau.dk

Sebastian Siebertz*
TU Berlin, Germany
sebastian.siebertz@tu-berlin.de

ABSTRACT

The Minimum Dominating Set (MDS) problem is not only one of the most fundamental problems in distributed computing, it is also one of the most challenging ones. While it is well-known that minimum dominating sets cannot be approximated locally on general graphs, over the last years, several breakthroughs have been made on computing local approximations on *sparse graphs*.

This paper presents a deterministic and local constant factor approximation for minimum dominating sets on bounded genus graphs, a large family of sparse graphs. 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. We believe that our techniques can be useful for the study of local problems on sparse graphs beyond the scope of this paper.

CCS Concepts

•Theory of computation → Distributed computing models; Distributed algorithms; •Networks → Network algorithms;

Keywords

Distributed Algorithms, LOCAL Model, Sparse Graphs, Minimum Dominating Sets

1. INTRODUCTION

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

*The research of Saeed Amiri and Sebastian Siebertz has been 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).

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODC'16, July 25 - 28, 2016, Chicago, IL, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-3964-3/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2933057.2933084>

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, even on planar graphs, however, the problem can be approximated well on planar graphs [1].

In this paper, we study the *distributed* time complexity of finding small dominating sets, in the classic *LOCAL model* of distributed computing [4, 11, 13, 14]. It is known that finding small dominating sets locally is hard: Kuhn et al. [7] show that in r rounds the MDS problem on an n -vertex graphs of maximum degree Δ can only be approximated within factor $\Omega(n^{c/r^2})$ and $\Omega(\Delta^{c'/r})$, where c and c' are constants. This implies that, in general, to achieve a constant approximation ratio, every distributed algorithm requires at least $\Omega(\sqrt{\log n})$ and $\Omega(\log \Delta)$ communication rounds. The currently best results for general graphs are by Kuhn et al. [7] who present a $(1 + \epsilon) \ln \Delta$ -approximation in $\mathcal{O}(\log(n)/\epsilon)$ rounds for any $\epsilon > 0$, and by Barenboim et al. [2] who present a deterministic $\mathcal{O}((\log n)^{k-1})$ -time algorithm that provides an $\mathcal{O}(n^{1/k})$ -approximation, for any integer parameter $k \geq 2$.

For sparse graphs, the situation is more promising. For graphs of arboricity a , Lenzen and Wattenhofer [10] present a forest decomposition algorithm achieving a factor $\mathcal{O}(a^2)$ approximation in randomized time $\mathcal{O}(\log n)$, and a deterministic $\mathcal{O}(a \log \Delta)$ approximation algorithm requiring time $\mathcal{O}(\log \Delta)$ 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. Czygrinow et al. [3] show that given any $\delta > 0$, $(1 + \delta)$ -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 [9]. Lenzen et al. [8] 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 [16] showed that the algorithm of Lenzen et al. in fact computes a 52-approximation of a minimum dominating set with small message size. Wawrzyniak [15] also showed that message sizes of $\mathcal{O}(\log n)$ suffice to give a (slightly worse) constant factor approximation on planar graphs. In terms of lower bounds, Hilke et al. [6] 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 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 $K_{3,t}$ ($t \geq 3$) 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 edge density. The most prominent locally embeddable graph classes are classes of bounded genus. Concretely, our result implies that MDS can be $\mathcal{O}(g)$ -approximated locally and deterministically on graphs of (both orientable or non-orientable) genus g . This result generalizes existing constant factor approximation results for planar graphs to a significantly larger graph family. Prior works by Lenzen et al. [8] and Wawrzyniak [16] 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.

Our main *technical* contribution is a new analysis of a slightly modified variant of the elegant algorithm by Lenzen et al. [8] for planar graphs. As we will show, with a slight modification, the algorithm also works for locally embeddable graphs.

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.

We believe that our new analysis and techniques can be useful also for the study of other local problems and on more general sparse graphs, beyond the scope of this paper.

An interesting side contribution of our modified algorithm is that it is *first-order definable*. In particular, the algorithm can again be modified such that it does not rely on any *maximum* operations, such as finding the neighbor of maximal degree. The advent of sub-microprocessor devices, such as biological cellular networks or networks of nano-devices, has recently motivated the design of very simple, “stone-age” distributed algorithms [5], and we believe that our work nicely complements the finite-state machine model assumed in related work, and opens an interesting field for future research.

1.2 Organization

The remainder of this paper is organized as follows. We introduce some preliminaries in Section 2. Our basic local approximation result is presented in Section 3, and the improved approximation, using preprocessing, is presented in Section 4. After discussing a logic perspective on our work in Section 5, we conclude in Section 6.

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 in G 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 . 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 $\{u \in V(G) : \{u, v\} \in E'\}$. For $A \subseteq V(G)$, we write $N[A]$ for $\bigcup_{v \in A} N[v]$. The edge density of G is $\epsilon(G) = |E(G)|/|V(G)|$. 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 $G[V(G) \setminus B]$. A graph H is a subgraph of a graph G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$.

Depth-1 minors and local embeddable graphs. A graph H is a minor of a graph G , written $H \preceq G$, if there is a set $\{G_v : v \in V(H)\}$ of pairwise disjoint 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 and we call $G[\bigcup_{v \in V(H)} V(G_v)]$ a *minor model* of H in G .

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. 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 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 .

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 $\epsilon(H) \leq c$ for some constant c and all depth-1 minors H of G .

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

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 not mention them explicitly in the following. We do not make explicit use of any topological arguments and hence refer to [12] for more background on graphs on surfaces. We will use the following facts about bounded genus graphs.

Graphs of genus g are closed under taking subgraphs and edge contraction.

LEMMA 1. *If $H \preceq G$, 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 (and in particular as depth-1 minors). The lemma

follows immediately from Lemma 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 [12]).

LEMMA 2. *If $g(G) = g$, then G excludes $K_{4g+3,3}$ as a minor and if $\tilde{g}(G) = \tilde{g}$, then G excludes $K_{2\tilde{g}+3,3}$ 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 [12].

LEMMA 3. *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.

LEMMA 4. *We have $|E(G)| \leq 3 \cdot |V(G)| + 6g(G) - 6$ and $|E(G)| \leq 3 \cdot |V(G)| + 3\tilde{g}(G) - 3$*

Distributed complexity. We consider the standard *LOCAL* model of distributed computing [11, 13], see also [14] 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.

3. THE LOCAL MDS APPROXIMATION

Let us start by revisiting the MDS approximation algorithm for planar graphs by Lenzen et al. [8], 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 a set D . It has been shown in [8] that the set D is small in planar graphs. In the second phase, it 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 from below and we replace the condition $|A| \leq 6$ 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

Algorithm 1 Dominating Set Approximation Algorithm for Planar Graphs

```

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 if
19: end for
20: return  $D \cup D'$ 

```

the algorithm does not have to compute the edge density of G , which is not possible in a local manner. Rather, we leverage Lemma 4 which upper bounds $\epsilon(G)$ 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 [8], which in fact uses no topological arguments at all.

LEMMA 5. *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 $\epsilon(H) \leq c$. 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 graph with $V(H) = M \cup N[D \setminus M]$ and where $E(H)$ is a minimal subset of $E(G[V(H)])$ such that all edges with at least one endpoint in $D \setminus M$ are contained in $E(H)$ and such that M is a dominating set in H . By this minimality condition, every vertex $v \in V(H) \setminus (D \cup M)$ has exactly one neighbor $m \in M$, no two vertices of $V(H) \setminus (M \cup D)$ are adjacent, and no two vertices of M are adjacent.

We construct a depth-1 minor \tilde{H} of H by contracting the star subgraphs G_m induced by $N_H[m] \setminus D$ for $m \in M \setminus D$ to a single vertex v_m . 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} . On the other hand, \tilde{H} has at most $c \cdot |V(\tilde{H})|$ edges, and also the subgraph $\tilde{H}[D \setminus M]$ has at most $c \cdot |D \setminus M|$ edges (by assumption on $\epsilon(\tilde{H})$).

Hence

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

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

ASSUMPTION 6. *For the rest of this section, we assume that G satisfies that for all depth-1 minors H of G , $\epsilon(H) \leq c$ for some constant c , and we fix M and D as in Lemma 5. We furthermore assume that for some $t \geq 3$, G excludes $K_{t,3}$ as depth-1 minor.*

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 in the ordering and inductively define a minimal set $E' \subseteq E(G)$ such that M is a dominating set with respect to E' . For this, we add all edges $\{m_1, v\} \in E(G)$ with $v \in N(m_1) \setminus M$ to E' . We then continue inductively by adding for $i > 1$ all edges $\{m_i, v\} \in E(G)$ with $v \in N(m_i) \setminus (M \cup N_{E'}(m_1, \dots, m_{i-1}))$.

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 1. In the figure, solid lines represent edges from E' , lines from $E(G) \setminus E'$ are dashed. We want to count the *endpoints* of directed edges, which represent the dominator function dom .




Figure 1: The graphs G_m . Solid lines represent edges from E' , directed edges 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 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 v .

H has $|M|$ vertices and by our assumption on the density of depth-1 minors of G , it has at most $c|M|$ 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 7 that each edge in H does not represent many such *dom*-edges with distinct endpoints. As H has at most $c|M|$ 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 11 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 $N_{E'}(m)$, note that existence of C_m follows from the definition of the set D . Recall that we assume that G excludes $K_{t,3}$ as depth-1 minor.

LEMMA 7. *Let $1 \leq i < j \leq |M|$. Let $N_i := N_{E'}(m_i)$ and $N_j := N_{E'}(m_j)$.*

1. *If $m_j \notin D$, then*

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

2. *If $m_i \notin D$ (and hence $m_j \notin D$), 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' , it holds that $m_i \notin C_{m_j}$. 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 $K_{t,3}$ as depth-1 minor, a contradiction. See Figure 2 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

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 8. $|Y| \leq 6c^2t|M|$.

PROOF. Each of the $c|M|$ many edges in H represents edges between N_i and N_j , where N_i and N_j are defined as




Figure 2: Visualisation of the proof of Lemma 7 in the case of excluded $K_{3,3}$

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 9. *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 7, 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 $C := \bigcup_{m \in M \setminus D} C_m$. 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 10.

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

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 11.

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

PROOF. Fix some $m \in M \setminus D$ and some $u \in N_{E'}(m)$ with $dom(u) \in N_{E'}(m) \setminus (Y \cup Z)$. Because $dom(u) \notin Y$, $dom(u)$ is not connected to a vertex of a different star, except possibly for vertices from M . Because $dom(u) \notin Z$, 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. By Lemma 9, $dom(u)$ is connected to at most $2c(t-1)$ vertices from the same star. Hence the degree of $dom(u)$ is at most $4c + 2c(t-1)$. Because u preferred to choose $dom(u) \in N_{E'}(m)$ 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 $(2(t-1) + 4)c|M|$ such vertices. \square

We are now ready to put together the numbers.

LEMMA 12. *If all depth-1 minors H of G have edge density at most c and G excludes $K_{t,3}$ 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. The set D has size at most $(c+1)|M|$ according to Lemma 5. Since M is a dominating set also with respect to the edges E' , it suffices to determine $|\{dom(u) : u \in (N_{E'}[M \setminus D] \setminus N[D])\}|$. According to Corollary 8, the set $Y = \{u \in N_{E'}(m_i) : \text{there is } v \in N_{E'}(m_j), i \neq j \text{ and } \{u, v\} \in E(G)\}$ has size at most $6c^2t|M|$. In particular, there are at most so many vertices $dom(u) \in N_{E'}(m_i)$ with $u \in N_{E'}(m_j)$ for $i \neq j$. Clearly, $|dom(R) \cap M| \leq |M|$ and $|dom(M)| \leq |M|$. Together, this bounds the number of endpoints of dom -edges that go from one star to another star. According to Lemma 10, there are only few vertices which are highly connected to $M \cup C$, that is, the set $Z = \{u \in N_{E'}(M \setminus D) : |N(u) \cap (M \cup C)| > 4c\}$ 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 11, this image has size at most $(2(t-1) + 4)c|M|$. In total, we hence find a set of size

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

\square

Our theorem for bounded genus graphs is now a corollary of Lemma 2 and 12.

THEOREM 13. *Let \mathcal{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 $\mathcal{O}(g)$ -approximation ($\mathcal{O}(\tilde{g})$ -approximation resp.) for the dominating set in $\mathcal{O}(g)$ ($\mathcal{O}(\tilde{g})$ resp.) 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 [16] improved analysis (52).

4. IMPROVING THE APPROXIMATION FACTOR WITH PREPROCESSING

We now show that the constant 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 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 say K_h is the v -canonical subgraph of G and we denote it by K_v . If $K = \emptyset$ we set $K_v := \emptyset$.

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

```

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

```

LEMMA 14. *Given a graph G and a vertex $v \in V(G)$. The v -canonical subgraph of G (K_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 $H = G[N^6(v)]$ 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.

THEOREM 15. *Algorithm 2 provides a $24g + \mathcal{O}(1)$ MDS approximation for graphs of genus at most g , and requires $12g + \mathcal{O}(1)$ communication rounds.*

PROOF. The resulting vertex set is clearly a legal dominating set. Moreover, as Phase 1 is unchanged, we do not have to recalculate D .

In the preprocessing phase, if for two vertices $u \neq v$ we choose both K_u, 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 intersect, the distance between u, v can be at most 12. On the other hand, by Lemma 3, there are at most g disjoint such 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.

In order to compute the size of the set in the third phase, we can use the analysis of Lemma 12 for $t = 3$, which together with the first phase and preprocessing phase, results in a $24g + \mathcal{O}(1)$ -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 com-

pute the 12-neighbourhood. Therefore, the number of communication rounds is $12g + \mathcal{O}(1)$. \square

This significantly improves the approximation upper bound of Theorem 13: namely from $4(6c^2 + 2c)g + \mathcal{O}(1)$, which, since $c \leq 3.01$ can be as high as $241g + \mathcal{O}(1)$ in sufficiently large graphs, to $24g + \mathcal{O}(1)$, at the price of $12g$ extra communication rounds.

5. 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. 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 not in the scope of a quantifier, and we write $\phi(x, \dots, x_k)$ 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. Observe however, that the only place in our analysis that refers to the dominator function d explicitly is Lemma 11. 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 vertex 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, however, we have to address how logic should choose a vertex among the candidate vertices. For this, we assume that the graph is equipped with an order relation such that the formula can simply choose the smallest candidate with respect to the order. It is now easy to see that the solution computed by the algorithm is first-order definable.

6. CONCLUSION

This paper presented the first constant round, constant factor local MDS approximation algorithm for locally embeddable graphs, a class of graphs which is much more general than planar graphs. Our proofs are purely combinatorial and avoid any 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.

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 [5]: distributed algorithms making minimal assumptions on the power of a node.

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.

7. REFERENCES

- [1] BAKER, B. S. Approximation algorithms for np-complete problems on planar graphs. *J. ACM* 41, 1 (Jan. 1994), 153–180.
- [2] BARENBOIM, L., ELKIN, M., AND GAVOILLE, C. A fast network-decomposition algorithm and its applications to constant-time distributed computation. In *Proc. SIROCCO*. 2014, pp. 209–223.
- [3] CZYGRINOW, A., HAŃKOWIAK, M., AND WAWRZY尼亚K, W. Fast distributed approximations in planar graphs. In *Proc. 22nd International Symposium on Distributed Computing (DISC)* (2008), pp. 78–92.
- [4] ELKIN, M. An overview of distributed approximation. *ACM SIGACT News Distributed Computing Column* 35, 4 (Mar. 2004), 40–57.
- [5] EMEK, Y., AND WATTENHOFER, R. Stone age distributed computing. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)* (2013), pp. 137–146.
- [6] HILKE, M., LENZEN, C., AND SUOMELA, J. Brief announcement: local approximability of minimum dominating set on planar graphs. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)* (2014), pp. 344–346.
- [7] KUHN, F., MOSCIBRODA, T., AND WATTENHOFER, R. Local computation: Lower and upper bounds. *arXiv preprint arXiv:1011.5470* (2010).
- [8] LENZEN, C., PIGNOLET, Y. A., AND WATTENHOFER, R. Distributed minimum dominating set approximations in restricted families of graphs. *Distributed Computing* 26, 2 (2013), 119–137.
- [9] LENZEN, C., AND WATTENHOFER, R. Leveraging linial’s locality limit. In *Distributed Computing*, vol. 5218 of *Lecture Notes in Computer Science*. 2008, pp. 394–407.
- [10] LENZEN, C., AND WATTENHOFER, R. Minimum dominating set approximation in graphs of bounded arboricity. In *Distributed Computing*, vol. 6343. 2010, pp. 510–524.
- [11] LINIAL, N. Locality in distributed graph algorithms. *SIAM J. Comput.* 21, 1 (Feb. 1992), 193–201.
- [12] MOHAR, B., AND THOMASSEN, C. *Graphs on Surfaces*. Johns Hopkins University Press, 2001.
- [13] PELEG, D. *Distributed Computing: A Locality-sensitive Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [14] SUOMELA, J. Survey of local algorithms. *ACM Comput. Surv.* 45, 2 (Mar. 2013), 24:1–24:40.
- [15] WAWRZY尼亚K, W. 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* (2013), ACM, pp. 406–408.
- [16] WAWRZY尼亚K, W. 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.