Distributed Approximation Algorithms for the Minimum Dominating Set in K_h-Minor-Free Graphs

Andrzej Czygrinow[1]
School of Mathematical and Statistical Sciences, Arizona State University, Tempe, AZ, 85287-1804, USA
aczygri@asu.edu

Michał Hanćkowiak
Faculty of Mathematics and Computer Science, Adam Mickiewicz University, Poznań, Poland
mhanckow@amu.edu.pl

Wojciech Wawrzyniak
Faculty of Mathematics and Computer Science, Adam Mickiewicz University, Poznań, Poland
wwawrzy@amu.edu.pl

Marcin Witkowski
Faculty of Mathematics and Computer Science, Adam Mickiewicz University, Poznań, Poland
mw@amu.edu.pl

Abstract

In this paper we will give two distributed approximation algorithms (in the Local model) for the minimum dominating set problem. First we will give a distributed algorithm which finds a dominating set D of size $O(\gamma(G))$ in a graph G which has no topological copy of K_h. The algorithm runs L_h rounds where L_h is a constant which depends on h only. This procedure can be used to obtain a distributed algorithm which given $\epsilon > 0$ finds in a graph G with no K_h-minor a dominating set D of size at most $(1+\epsilon)\gamma(G)$. The second algorithm runs in $O(\log^* |V(G)|)$ rounds.

2012 ACM Subject Classification Theory of computation → Distributed computing models

Keywords and phrases Distributed algorithms, minor-closed family of graphs, MDS

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2018.22

1 Introduction

The minimum dominating set (MinDS) problem is one of the classical graph-theoretic problems which is of theoretical and practical importance. A subset D of the vertex set in a graph G is called a *dominating set* in G if every vertex of G is either in D or has a neighbor in D. In the *minimum dominating set problem* the objective is to find a dominating set D of the smallest size. In this paper we will study distributed approximation algorithms in the *Local* model for the MDS problem in K_h-minor-free graphs.

Although the MDS problem is NP-complete even in planar graphs, there are efficient approximation algorithms. Significant progress has been made in recent years in understanding distributed complexity of many classical graph-theoretic problems in some classes of sparse graphs. In the case of the maximum independent set problem, (MaxIS Problem), it is

[1] Research supported in part by Simons Foundation Grant # 521777.

known ([2]) that finding deterministically a constant approximation of $\alpha(G)$ even in the case when G is a cycle on n vertices requires $\Omega(\log^* n)$ rounds. At the same time, it is possible to find an independent set in a planar graph G of size at least $(1-\epsilon)\alpha(G)$ in $O(\log^* n)$ rounds ([2]). In fact, this result immediately extends to graphs G with no K_h-minor using the same approach. Even more can be proved for the maximum matching problem (MaxM Problem). On one hand, the above lower bound for the independence number extends to matchings and on the other hand, there is a distributed algorithm which finds in $O(\log^* n)$ rounds a matching M of size at least $(1-\epsilon)\beta(G)$ even in graphs of bounded arboricity ([1]). This procedure relies on augmenting paths and is very specific to the maximum matching problem. At the same time, the lower bound for the approximation of maximum independent set, does not extend to the minimum dominating set problem, and a constant-approximation which runs in a constant number of rounds is known for planar graphs and graphs with bounded genus. Specifically, Lenzen et. al. gave in [4] a distributed algorithm which in $O(1)$ rounds finds a dominating set of size at most $126\gamma(G)$ in a planar graph G and Amiri et. al. ([6]) gave $O(g)$-approximation for graphs of genus bounded by g which runs in $O(1)$ rounds. The landscape changes when randomization is allowed. It can be shown that there is a randomized algorithm which in $O(1)$ rounds finds with high probability an independent set I of size at least $(1-\epsilon)\alpha(G)$ in $O(1)$ rounds in a planar graph G ([2]) and similar results can be obtained for the maximum matching. In addition, Lenzen and Wattenhofer [3] showed that there is a $O(a^2)$-approximation of a minimum dominating set can be found in the randomized time $O(\log\Delta)$ in a graph of arboricity a.

In this paper, we will propose deterministic distributed approximation algorithms for the MinDS problem in K_h-minor-free graph.

Recall that H is called a *minor* of G if H can be obtained from a subgraph of G by a sequence of edge contractions. A graph is called K_h-minor-free if it doesn't contains the complete graph K_h as a minor. For a graph H, TH, a topological copy of H, is a graph obtained from H by subdividing each edge $e \in H$ l_e times, for some $l_e \in N_0$. Many important classes of graphs, like for example planar graphs, graphs of bounded genus, or bounded tree-width are K_h-minor-free for some h. As a result our work on K_h-minor-free graphs generalizes previous results on planar and bounded-genus graphs.

Our algorithms work in the *Local* model. This is a synchronous model where a network is modeled as an undirected graph. Each vertex corresponds to a computational unit and an edge to a communication link between two units. Computations proceed in synchronous rounds and in each round a vertex can send and receive messages from its neighbors and can perform local computations. Neither the amount of local computations nor the size of messages is restricted in any way. In addition, we shall assume that vertices of G have unique identifiers from $\{1,\dots,n\}$, where n is the order of G. Consequently, in the case of the MinDS problem, for an underlying network G the objective is to find a set $D \subseteq V(G)$, in the above model, which is a minimum dominating set and has size $O(\gamma(G))$.

We will prove two results on distributed algorithms for the minimum dominating set problem. Our main theorem shows that in the case of graphs H which no TK_h it is possible to find a constant approximation of $\gamma(H)$ in $O(1)$ rounds. Specifically, we have the following theorem.

▶ **Theorem 1.** *Let $h \geq 2$. There exists a distributed algorithm which in a graph H of order n which has no TK_h finds in L_h rounds a dominating set D such that $|D| \leq C_h\gamma(H)$, where L_h and C_h are depending on h only.*

We didn't try to optimize constants C_h and L_h and especially in the case of C_h our proof gives a very big constant. In addition, its value depends on the constant from one of the facts from [5] (see Lemma 3), which in turn, depends on h.

Theorem 1 generalizes results for planar graphs from [4] and bounded genus graphs from [6] to graphs with no topological copy of K_h. In addition, it gives a deterministic $O(1)$-approximation of the MinDS Problem which runs in a constant number of rounds in an important subclass of graphs of bounded arboricity. The proof of Theorem 1 is split into two main steps. In the first step, an algorithm finds certain partitions of $N(v)$ for vertices v, and in the second, for every set W in these partitions, W finds a vertex v such that $W \subseteq N(v)$ and v is a "safe" choice to be added to a dominating set. The proof uses a fact from [5], in addition with some new ideas.

Clearly, if G is a graph which is K_h-minor-free then it contains no TK_h. Consequently, Theorem 1 can also be used, in connection with methods from [2], to obtain a much better approximation ratio in $O(\log^* n)$ rounds when restricted to graphs with no K_h-minor.

▶ **Theorem 2.** *Let $h \geq 2$. There exists a distributed algorithm which given $\epsilon > 0$ finds in a K_h-minor-free graph H of order n a dominating set D such that $|D| \leq (1+\epsilon)\gamma(H)$. The algorithm runs $C \log^* n$ rounds where C depends on h and ϵ only.*

Basically, to prove Theorem 2, we first find a constant approximation of $\gamma(H)$ in H using Theorem 1 and then apply a procedure from [2] to find a better approximation. This in turn, generalizes a corresponding result from [2] to graphs which are K_h-minor free. Note that once Theorem 1 is established, Theorem 2 can be proved by appealing to a fact from [2] which extends to graphs with no K_h-minors in a straightforward way and the rest of the paper is focused on proving Theorem 1. On the other hand, proving Theorem 1 requires new approach as the ideas from [4] and [6] are specific to planar graphs and graphs of bounded genus.

The rest of the paper is structured as follows. In Section 2, in addition to preliminary observations, we will discuss our main tool of building certain partitions of sets $N(v)$ arising from the so-called pseudo-covers. In Section 3 we will prove Theorem 1.

2 Preliminaries

In this section, we will introduce auxiliary concepts which are used in our algorithm. We will start with some definitions. Let $G = (V, E)$ be a graph. For two sets $A, B \subseteq V$, an A, B-path is a path which starts in a vertex from A, ends in a vertex from B and has all internal vertices from $V \setminus (A \cup B)$. In the case $A = \{a\}$, we will simplify the notation to an a, B-path.

Let $D \subset V$ and let $v \in V \setminus D$. A v, D-fan is a set of v, D-paths $P_1, \ldots, P_s$ such that for $i \neq j$, $V(P_i) \cap V(P_j) = \{v\}$. For $k, l \in Z^+$ and set D, let $D_{k,l}$ be the set of vertices w such that there is w, D-fan consisting of k paths each of length at most l. We have the following fact from [5].

▶ **Lemma 3.** *For $h, l \in Z^+$ there is c such that the following holds. Let G be a graph with no TK_h and let D be a dominating set in G. Then $|D_{h-1,l}| \leq c|D|$.*

To describe the intuition behind our approach for approximating a minimum dominating set D^* consider a planar graph G. Let $v \in G$ be an arbitrary vertex. Then $N(v)$ is dominated by some vertices from D^*. It is possible that the minimum number of vertices needed to dominate $N(v)$ is "big", in this case, in view of Lemma 3 (with $l = 1$), adding v or any other vertex of a similar type, will yield a dominating set of size $O(|D^*|)$. Consequently, we have to address the case when the number of such vertices in D^* is small. The main idea of how to build on this intuition is as follows. Supposing $|N(v)| \geq 3$ it is not possible to have three different vertices which dominate the whole set $N(v)$, as it gives a copy of $K_{3,3}$. A similar fact should be true if instead of three vertices dominating $N(v)$, we have for some constant t sets $U_1, \dots, U_t$ such that the size of each each U_i is constant and vertices from U_i dominate $N(v)$. This however is not exactly true. Indeed, for example, it is possible that there is one vertex u dominating all but one vertex from $N(v)$ and many vertices $v_1, \dots, v_s$, each dominating the remaining vertex from $N(v)$, so that $U_i = \{u, v_i\}$. However the contribution, i.e. the number of vertices covered in $N(v)$, by each v_i is minimal and, as it turns out, it can be ignored. Building on this, our approach is to find for $N(v)$, a family of the so-called pseudo-covers, that is sets of a constant size which cover almost all vertices from $N(v)$ and such that each vertex makes a substantially contribution. Note that $\{v\}$ is a choice for such a pseudo-cover. We will argue that the number of such pseudo-covers must be constant and will use these covers to partition $N(v)$ into a constant number of sets of which each, but the exceptional class, will be big and will be covered by a constant number of vertices. Of course, it is not clear which of these vertices should be included in a dominating set, but suppose initially that there are only two vertices v and u which cover a set in this partition. We claim that adding u is a reasonable choice. As indicated above, we will be able to assume that one of u and v is in D^*. If we are lucky then $u \in D^*$; If however $v \in D^*$, then since there is a constant number of vertices in pseudo-partitions, adding u or any other vertex because of v yields a constant approximation.

Describing these ideas more formally requires a little bit of preparation. Let $G = (V, E)$ be a graph. We say that $Z \subseteq V$ is a *cover* of $W \subseteq V$ if $Z \cap W = \emptyset$ and $W \subseteq \bigcup_{x \in Z} N(x)$.

Let $W \subseteq V$ and let $x \in V \setminus W$. We say that x is *α-strong* for W if $|N(x) \cap W| \geq \alpha|W|$. Using the same idea as in a proof of the Kövári-Sos-Turán theorem we have the following fact.

▶ **Fact 4.** *Let $\alpha \in (0, 1)$, $t, s \in Z^+$ and let $M := \frac{(t-1)e^s}{\alpha^s}$. If $G = (V, E)$ is a graph with no $K_{t,s}$ and $W \subseteq V$ is such that $|W| \geq s/\alpha$, then there are at most M α-strong vertices in $V \setminus W$ for W.*

Proof. Let U denote the set of α-strong vertices for W. We will count the number of claws $K_{1,s}$ in the graph $G[U, W]$ with centers in U. On one hand, the number of claws is at least $|U|\binom{\alpha|W|}{s}$, and on the other hand, since $G[U, W]$ has no $K_{t,s}$, every s-element subset of W can be involved in at most $t - 1$ claws. Thus

$$|U|\binom{\alpha|W|}{s} \leq \binom{|W|}{s}(t-1)$$

and so

$$|U|\left(\frac{\alpha|W|}{s}\right)^s \leq \left(\frac{e|W|}{s}\right)^s (t-1)$$

which gives $|U| \leq \frac{(t-1)e^s}{\alpha^s}$. ◀

We are now ready to define the main concept which is used in our algorithm, the notion of an (α, q, l, K)-pseudo-cover.

▶ **Definition 5.** An (α, q, l, K)-pseudo-cover of a set $W \subseteq V$ is a vector of vertices $(x_1, \dots, x_m)$ such that for every i, $x_i \notin W$, and the following conditions are satisfied.

(a) $|W \setminus \bigcup_{i=1}^m N(x_i)| \leq q$;
(b) x_i is α-strong for $W \setminus \bigcup_{j<i} N(x_j)$;
(c) $|N(x_i) \cap (W \setminus \bigcup_{j<i} N(x_j))| \geq l$;
(d) $m \leq K$.

When using the concept, α will be a constant from $(0,1)$, and q, l, K will be constants which depend on h when we consider graphs with no TK_h. To be more precise,

$$K := 2h - 2, \alpha := \frac{1}{K}, l := \frac{h}{\alpha} + 1, q := K \cdot l. \tag{1}$$

In addition, we will have

$$s := h, t := \binom{h}{2} + h.$$

Also note, that in the degenerate case when $|W| \leq q$, we will allow the empty vector.

It is not difficult to see that any cover of a set W with at most K vertices contains an (α, q, l, K)-pseudo-cover with $\alpha = 1/K$ and $l = q/K$.

▶ **Fact 6.** *For every q and every cover Z of W such that $|Z| = K$ there is an ordering of vertices of Z, $(x_1, \ldots, x_K)$, such that for some $m \leq K$, $(x_1, \ldots, x_m)$ is an (α, q, l, K)-pseudo-cover of W with $\alpha = \frac{1}{K}$ and $l = \frac{q}{K}$.*

Proof. Let $l := q/K$. If $|W| \leq q$, then the pseudo-cover is empty. Otherwise, let $x_1 \in Z$ be such that $|N(x_1) \cap W|$ is maximum. Then $|N(x_1) \cap W| \geq |W|/K \geq l$. For the general step. Suppose $|W \setminus \bigcup_{j<i} N(x_j)| > q$. Then there exists $y \in Z \setminus \{x_1, \ldots, x_{i-1}\}$ such that $|N(y) \cap (W \setminus \bigcup_{j<i} N(x_j))| \geq |W \setminus \bigcup_{j<i} N(x_j)|/K$. Set $x_i := y$. We have $|N(y) \cap (W \setminus \bigcup_{j<i} N(x_j))| > q/K$. ◀

One of the key observations used in the proof is that the number of (α, q, l, K)-pseudo-cover of a set W does not depend on $|W|$.

▶ **Lemma 7.** *Let $\alpha \in (0,1)$ $s, t, K \in Z^+$, let $l > s/\alpha$ and $q := l \cdot K$. Then for every graph G with no $K_{s,t}$ and every $W \subseteq V(G)$ such that $|W| \geq l$, the number of (α, q, l, K)-pseudo-covers of W is at most $2\left(\frac{(t-1)e^s}{\alpha^s}\right)^K$.*

Proof. Suppose the number of (α, q, l, K)-pseudo-covers is bigger than $C := 2\left(\frac{(t-1)e^s}{\alpha^s}\right)^K$. Since the first positions are α-strong for W and $|W| \geq s/\alpha$, by Fact 4 there can be at most $M := \frac{(t-1)e^s}{\alpha^s}$ of (α, q, l, K)-pseudo-covers with distinct first positions. Let x_1 be a vertex which appears most often in the first position of these covers. Then at least $C/M > 1$ of the covers have x_1 in the first position and out of these there can be at most one (α, q, l, K)-pseudo-cover which contains only x_1. If $|W \setminus N(x_1)| \leq l$ then no vertex can cover more than l vertices of $W \setminus N(x_1)$. Thus we may assume otherwise. Now we can iterate the above argument restricting attention to those (α, q, l, K)-pseudo-covers which have x_1 in the first position. We have $|W \setminus N(x_1)| > l > s/\alpha$, and so by Fact 4 at least $(C/M - 1)/M = (C - M)/M^2 > 1$ vectors have the second position equal to some x_2. Iterating the above gives that there are at least

$$(C - M - M^2 - \cdots - M^{i-1})/M^i$$

(α, q, l, K)-pseudo-covers starting with $x_1, x_2, \ldots, x_i$ for some $x_1, \ldots, x_i$. Since a pseudo-cover has at most K vertices, $C < 2M^K$ for the above quantity to be at most one when $i = K$; a contradiction. ◀

Let v be such that there exist K vertices $x_1, \ldots, x_K \in V \setminus \{v\}$ with the property $N(v) \subseteq \bigcup_{j \leq K} N(x_j)$. The number of such covers of $N(v)$ can be "large" but in view of Fact 6 and Lemma 7 the number of (α, q, l, K)- pseudo-covers such that $l > s/\alpha$ obtained from covers of

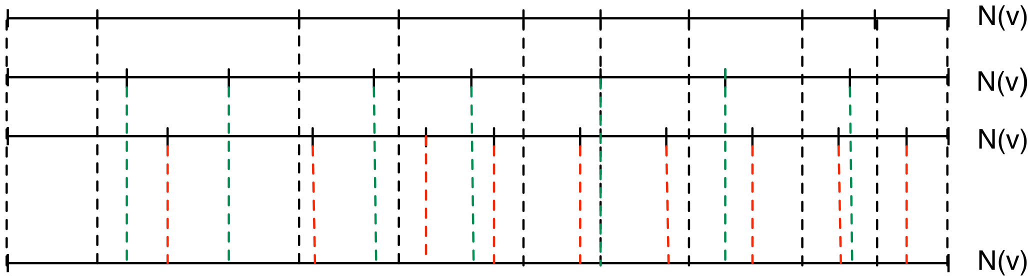

Figure 1 An illustration of refining partitions for three pseudo-covers of $N(v)$. For simplicity the sets in partitions are depicted as intervals but they can be arbitrary.

$N(v)$ is a constant independent of $|N(v)|$. In the rest of the section we will use the fact that the number of (α, q, l, K)-pseudo-covers is constant to refine partitions determined by the covers into a constant number of sets. Fix $0 < \alpha < 1$ and $s, K \in Z^+$ and l so that $l > s/\alpha$. Let v be such that $|N(v)| \geq l$.

Let $\mathcal{T}(v)$ denote the set of (α, q, l, K)-pseudo-covers of $N(v)$. By Lemma 7, we have $|\mathcal{T}(v)| \leq C$ where $C := 2\left(\frac{(t-1)e^s}{\alpha^s}\right)^K$. For $S := (x_1, \dots, x_m) \in \mathcal{T}(v)$ consider the following partition $\mathcal{P}_S = \{W_0, W_1, \dots, W_m\}$ of $N(v)$. Let $W_1 := N(x_1) \cap N(v)$, $W_i := (N(x_i) \cap N(v)) \setminus \bigcup_{1 \leq j < i} W_j$ for $i > 1$, and let $W_0 := N(v) \setminus \bigcup_{j \leq m} N(x_j)$. Since S is an (α, q, l, K)-pseudo-covers of $N(v)$, we have $|W_0| \leq q$.

Let $\mathcal{Q}(v)$ be the minimal partition which refines partitions $\mathcal{P}_S$ over all (α, q, l, K)-pseudo-covers S from $\mathcal{T}(v)$. For example, if there are only two partitions, $\mathcal{P}_S = \{W_0, W_1, \dots, W_m\}$ and $\mathcal{P}_T = \{U_0, U_1, \dots, U_m\}$, then $\mathcal{Q}(v)$ contains all non-empty intersections $W_i \cap U_j$.

▶ **Fact 8.** $|\mathcal{Q}(v)| \leq 2^{(K+1)C}$

Proof. For $S = (x_1, \dots, x_m)$, we have $|\mathcal{P}_S| \leq m + 1 \leq K + 1$ and so there are at most $(K+1)C$ different subsets of $N(v)$ over all $S \in \mathcal{T}(v)$. Taking the refinement of these partitions results in at most $2^{(K+1)C}$ sets. ◀

We will now modify $\mathcal{Q}(v)$ as follows. Let V_0 be the union of these partition classes in $\mathcal{Q}(v)$ which are subsets of $W_0 = N(v) \setminus \bigcup_{j=1}^{m} N(x_j)$ for at least one $S = (x_1, \dots, x_m)$. Let $\{V_1, \dots, V_s\}$ denote the remaining partition classes. Then $\{V_0, V_1, \dots, V_s\}$ is a partition of $N(v)$ (See Figure 1 for an illustration). In addition, we have the following fact.

▶ **Fact 9.** *The following conditions are satisfied.*

(1) $|V_0| \leq Cq$.

(2) *For $i \geq 1$ and for every $(x_1, \dots, x_m) \in \mathcal{T}(v)$, $V_i \subset N(x_j)$ for some $j \in [m]$.*

Proof. The number of vertices which belong to at least one set W_0 is at most Cq. For part (2), fix V_i for $i \geq 1$ and let $(x_1, \dots, x_m) \in \mathcal{T}(v)$. Then vertices from V_i are covered by $\bigcup_{j=1}^{m} N(x_j)$ because V_i doesn't intersect $N(v) \setminus \bigcup_{j=1}^{m} N(x_j)$. Let j_1 be the smallest index such that $V_i \cap N(x_{j_1}) \neq \emptyset$ and suppose that for some $j_2 > j_1$, we have $V_i \cap (N(x_{j_2}) \setminus N(x_{j_1})) \neq \emptyset$. Then V_i intersects W_{j_1}, but since it is not contained in W_{j_1}, it intersects another set in the (α, q, l, K)-pseudo-cover determined by $(x_1, \dots, x_m)$, and so it cannot belong to $\{V_0, V_1, \dots, V_s\}$. ◀

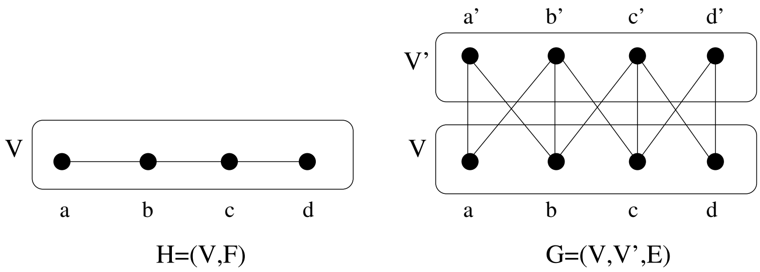

Figure 2 The bipartite graph G associated with H.

We will end this section with some more notation which will be used later. Recall that $\mathcal{T}(v)$ denotes the set of (α, q, l, K)-pseudo-covers of $N(v)$. For set of vertices U we will define $\mathcal{T}(U) := \bigcup_{v\in U} \mathcal{T}(v)$. For a set $\mathcal{S}$ of (α, q, l, K)-pseudo-covers we let $V_{\mathcal{S}}$ be the set of vertices which belong to at least one pseudo-cover from $\mathcal{S}$. We will slightly abuse above notation and define

$$\mathcal{T}(\mathcal{S}) := \mathcal{T}(V_{\mathcal{S}}).$$

Using the above convention, we will use $\mathcal{T}^{(i)}(U) := \mathcal{T}(\mathcal{T}^{(i-1)}(U))$ with $\mathcal{T}^{(1)}(U) := \mathcal{T}(U)$ and $\mathcal{T}^{(\leq k)}(U) := \bigcup_{1\leq i\leq k} \mathcal{T}^{(i)}(U)$.

3 Algorithm

In this section we will give the main algorithm. The algorithm consists of two phases. In the first phase we simply add to a dominating set D vertices v which have only one vector in $\mathcal{T}(v)$, namely (v). In the second phase, we analyze sets in $\mathcal{Q}(v)$ and argue that if a set V_i is big enough then we will be able to find a "good" choice among a constant number of vertices from vectors in $\mathcal{T}(v)$ to dominate V_i.

Let $H = (V, F)$ be a graph with no TK_h and recall that K, α, l, q are given in (1). It will be convenient to work in the double-cover of H which we are going to define next. We say that the bipartite graph $G = (V, V', E)$ is *associated with* H if $V' = \{v' : v \in V\}$ and we have $vu' \in E$ if and only if $vu \in F$ or $u = v$. In other words, edge $vu \in F$ corresponds to two edges vu', uv' in E and $vv' \in E$ for every v from V. Let $\gamma'(G)$ denote the minimum size of a set $S \subseteq V$ which dominates V' in G. Before discussing the first phase of the algorithm, we will mention a few facts on the relation between H and G.

▶ **Fact 10.** *X is a dominating set in H if and only if $N_G(X) = V'$.*

Proof. Suppose X is a dominating set in H. Then every vertex $u \in V(H) \setminus X$ is adjacent to a vertex $v \in X$, and so $u'v \in E(G)$, and for every $u \in X$, $uu' \in E(G)$. Now, let $Y \subset V$ and let $u \in V(H) \setminus Y$. Since $N_G(Y) = V'$, there is a vertex $v \in Y$ such that $u'v \in E(G)$. Then $uv \in E(H)$ as $u' \neq v'$. ◀

In particular, we have

$$\gamma(H) = \gamma'(G).$$

Rather than studying topological minors in G in relation to topological minors in H, we note the following simple fact.

▶ **Fact 11.** *Let $D \subseteq V$, $v \in V \setminus D$, and suppose there is a v, D-fan in G of size $2t-1$ such that every path has length two. Then $v \in D_{t,2}$ in H.*

Proof. A v, D-fan in G consists of paths of the form v, u', w, where $w \in D$. Such paths are mapped to paths of length one (if the corresponding copy of u' is in D) or paths of length two (if the corresponding copy of u' is not in D). Thus for every vertex from D there are at most two paths that are mapped to ones that contain this vertex in H. As a result we can always choose t vertex disjoint paths of a v, D-fan in G out of the $2t-1$ paths of v, D-fan in H. ◀

Lemma 3 and Fact 11 give the following corollary.

▶ **Corollary 12.** *If H has no TK_h and $D \subset V$ is such that $N_G(D) = V'$, then the number of vertices $v \in V$ such that there is a v, D-fan in G of size $2h-1$ such that each path has length two is $O(|D|)$.*

Since we will partition sets $N_G(v) \subseteq V'$ (for $v \in V$), all the partition classes will be subsets of V'. It will be convenient to introduce the following notion.

▶ **Definition 13.** We say that the partition $\{V'_0, V'_1, \ldots, V'_s\}$ of $N_G(v)$ from Fact 9 is associated with v.

Let $D^* \subset V$ be an optimal set which dominates V' in G and let V^* be the set of vertices $v \in V \setminus D^*$ such that there is a v, D^*-fan consisting of $K+1$ paths, each of length at most two. Then, by Corollary 12, $|V^*| = O(|D^*|)$. In view of the previous discussion adding vertices from $V^* \cup D^*$ to our solution results in a dominating set of size $O(|D^*|)$. The remaining vertices v from $V \setminus (V^* \cup D^*)$ will have $N(v)$ dominated by a "few" vertices from D^*. Suppose $v \in V \setminus (V^* \cup D^*)$. Then $N(v)$ is dominated by vertices from D^* and so there exist $d_1, \ldots, d_m \in D^*$ for some $m \leq K$, such that $N(v) \subseteq \bigcup N(d_i)$. Therefore $\{d_1, \ldots, d_m\}$ is a cover of $N(v)$ and by Fact 6, $\{d_1, \ldots, d_m\}$ gives an (α, q, l, K)-pseudo-cover which belongs to $\mathcal{T}(v)$. In addition, by Fact 9, if $\{V'_0, V'_1, \ldots, V'_s\}$ is the partition associated with v, then for every $i \geq 1$, $V'_i \subset N(d_j)$ for some $j \leq m$. In view of the previous discussion, we have the following fact.

▶ **Fact 14.** *Let $D^* \subseteq V$ be an optimal set which dominates V' in G and let $v \in V \setminus (D^* \cup V^*)$. In addition, let $V'_0, V'_1, \ldots, V'_s$ be the partition associated with v, and for $i \geq 1$, let $U_i = \{x \in S : S \in \mathcal{T}(v) \wedge V'_i \subseteq N(x)\}$. Then $U_i \cap D^* \neq \emptyset$.*

In the main part of our algorithm, we will find a set $D \subseteq V$ which dominates some vertices from V' so that when $N_G(D)$ is removed from V', then every vertex in V has its degree bounded by a constant. To extend D and find a set which dominates all vertices from V' we will rely on the following simple observation.

▶ **Fact 15.** *Let $s \in Z^+$ and let $X = (V, V', E)$ be a bipartite graph such that for every vertex $v \in V$, $d(v) \leq s$ and for every $v' \in V'$, $d(v') \geq 1$. Then $\gamma'(X) \geq |V'|/s$.*

Proof. If $D \subseteq V$ is an optimal set which dominates V', then $s|D| \geq |E(D, V')| \geq |V'|$. ◀

Recall that for every $v \in V$, $\mathcal{T}(v) \neq \emptyset$ because $(v) \in \mathcal{T}(v)$. To motivate our discussion suppose first that $|\mathcal{T}(v)| = 1$. Then either $v \in V^*$ (defined above) or otherwise, in view of Fact 14, $v \in D^*$. In either case we can add such a vertex to our solution. In fact a stronger observation is true, if for some V_i in the partition associated with v, v is the only vertex in some $S \in \mathcal{T}(v)$ such that $V_i \subseteq N(v)$, then $v \in V^* \cup D^*$. Unfortunately, in many cases there will be more than one vertex u such that $V_i \subseteq N(u)$ and the challenge is to select one which will lead to a constant approximation of $\gamma'(G)$. The assumption that H has no TK_h implies that the number of choices of u is bounded by a constant which depends on h only but only some of these vertices will be good choices. The algorithm will consist of two main phases. The first one deals with those vertices v for which $|\mathcal{T}(v)| = 1$, and the second one addresses the more difficult case.

Phase 1

Input: Graph $H = (V, F)$ with no TK_h.

1. Consider the double cover $G = (V, V', E)$ of H. Let $D_1 := \emptyset$.
2. Compute $\mathcal{T}(v)$ for every $v \in V$. If $|\mathcal{T}(v)| \geq 2$ then mark v. Add all unmarked vertices to D_1 and delete all vertices from V' dominated by D_1.

Next fact follows from previous discussion.

▶ **Lemma 16.** *Let D_1 be the set obtained in Phase 1. Then $|D_1| = O(\gamma(H))$.*

We will now continue our analysis assuming we have set D_1 obtained in Phase 1 and G has been modified by possibly deleting some vertices from V'. Let V'' denote the remaining vertices in V', that is $V'' := V' \setminus N_G(D_1)$. In addition, we shall use V_i'' to denote $V_i' \cap V''$.

Consider a sequence of constants $M_0, M_1, \ldots, M_h$ such that $M_h \geq \binom{h}{2} + h$ and for every $1 \leq i \leq h$, we have

$$M_{i-1} > (2^{(K+1)C} + 1)(Cq + M_i 2^{(K+1)C}),$$

where $C = 2\left(\frac{(t-1)e^s}{\alpha^s}\right)^K$ and K, q, s, t are defined in (1). In fact, we will only need that $M_{h-1} \geq \binom{h}{2} + h$ but in the process described below, which uses constants M_i, we will allow it to continue more than $h - 1$ times. Let $v \in V \setminus D_1$ and let V_i' be a set in the partition associated with v which satisfies

$$|V_i''| \geq M_0.$$

We set $v_0 := v$ and consider V_i''. For every $v_1 \in V \setminus (D_1 \cup \{v\})$ which belongs to some $S \in \mathcal{T}(v)$ and is such that $V_i'' \subseteq N(v_0)$ take partition $W_0', W_1', \ldots, W_p'$ associated with v_1 and let $W_i'' := W_i' \cap V''$. We have $p \leq 2^{(K+1)C}$ by Fact 8, $|W_0'| \leq Cq$ by Fact 9. Let $\mathcal{P}_{v_0v_1}^{V_i''} = \{W_j'' \cap V_i'' : |W_j'' \cap V_i''| \geq M_1 \wedge j \geq 1\}$. We have $\bigcup \mathcal{P}_{v_0v_1}^{V_i''} \subseteq V_i''$ and $\bigcup_{j \geq 0} V_i'' \cap W_j' = V_i''$. Therefore,

$$|\bigcup \mathcal{P}_{v_0v_1}^{V_i''}| \geq |V_i''| - Cq - 2^{(K+1)C} \cdot M_1.$$

We call sets $W_j'' \cap V_i'' \in \mathcal{P}_{v_0v_1}^{U}$ *fragments*. Now we iterate the process for every fragment $W_j'' \cap V_i'' \in \mathcal{P}_{v_0v_1}^{V_i''}$, that is, we consider $v_2 \in V \setminus (D_1 \cup \{v_1\})$ such that for some $S \in \mathcal{T}(v_1)$ we have $v_2 \in S$, $W_j' \subseteq N_G(v_2)$ and $v_2 \neq v_i$ for $i < 2$. Define $\mathcal{P}_{v_0v_1v_2}^{W_j''\cap V_i''} = \{Z_k'' \cap W_j'' \cap V_i'' : |Z_k'' \cap W_j'' \cap V_i''| \geq M_2 \wedge k \geq 1\}$. We have

$$|\bigcup \mathcal{P}_{v_0v_1v_2}^{W_j''\cap V_i''}| \geq |W_j'' \cap V_i''| - 2Cq - 2^{(K+1)C} \cdot M_2.$$

We repeat the process for as long as possible (See Figure 3 for an illustration). We will now establish three claims about the above process. First claim states that the process must end after $h - 1$ steps or the original graph H contains a TK_h. The second claim shows that in

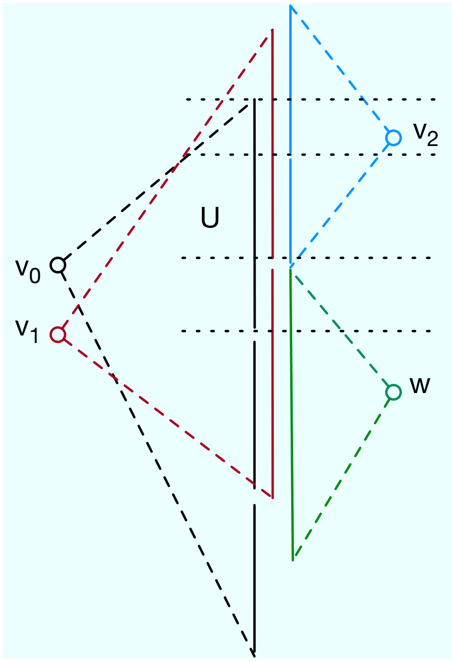

Figure 3 Constructing fragments from a set U as vertices are added to a sequence starting with v_0.

the sequences of vertices obtained by the process, the last vertices are a good choice to be added to a solution. Finally, the last claim states that when last vertices are added then all but a constant number of vertices from V_i'' are dominated.

▶ **Claim 17.** *If $v_0, v_1, \ldots, v_i$ is a sequence obtained in the above process, then $i \leq h-2$.*

Proof. Suppose $i \geq h-1$. Then $\{v_0, \ldots, v_{h-1}\}$ contains distinct vertices, every fragment $U \in \mathcal{P}^W_{v_0 \ldots v_{h-1}}$ has size $|U| \geq M_{h-1}$, and $U \subseteq N(v_0) \cap \cdots \cap N(v_{h-1})$. Since $M_{h-1} \geq \binom{h}{2} + h$, G contains $K_{h,\binom{h}{2}+h}$, and as a result H contains TK_h. ◀

For every maximal sequence $v_0, v_1, \ldots, v_j$ obtained in the process above for V_i'', we add v_j to $D_2(V_i')$.

Let Z be the set of vertices $z \in V$ that belong to some S where $S \in \mathcal{T}^{(\leq h)}(d)$ for some $d \in D^*$. Then we have $|Z| = O(|D^*|)$ because there is a constant number of vertices which belong to some $S \in \mathcal{T}^{(\leq h)}(d)$. In addition, we have the following.

▶ **Claim 18.** *We have $D_2(V_i') \subseteq V^* \cup D^* \cup Z$.*

Proof. Suppose $v_0 \ldots v_i$ is a maximal sequence. Let U be a fragment in $\mathcal{P}^W_{v_0 \ldots v_i}$ and let $X_0', \ldots, X_l'$ denote the partition associated with v_i. Then $U \subseteq X_j'$ for some $j \geq 1$. By Fact 14, either $v_i \in V^* \cup D^*$ or for at least one $d \in D^*$, we have d in some $S \in \mathcal{T}(v_i)$ and $X_j' \subseteq N_G(d)$. Recall that $\mathcal{T}(d)$ is non-empty and so there is a partition associated with d, but it can be trivial. We have $|U \cap Y_j'| \geq M_{i+1}$ for at least one set Y_j' such that $j \geq 1$ and Y_j' is in the partition associated with d. Thus d is an option for v_{i+1} and so, since the sequence is maximal $d = v_k$ for some $k < i$. We have $i \leq h-2$ by Claim 17. Consequently $v_i \in V^* \cup D^* \cup Z$. ◀

Finally we show that vertices from $D_2(V_i')$ cover all but a constant number of vertices in V_i''.

▶ **Claim 19.** *There is a constant* $L = L(K, C, l, h)$ *such that* $|V_i'' \setminus \bigcup_{x \in D_2(V_i')} N_G(x)| \leq L$.

Proof. If $v_0 v_1 \ldots v_j$ is maximal and U is a fragment in $\mathcal{P}^W_{v_0 v_1 \ldots v_j}$, then $U \subseteq N(v_j)$ and $v_j \in D_2(V_i')$. For any fragment U and any sequence $v_0 v_1 \ldots v_k$ which is not maximal, U is partitioned using the process above, the sequence is extended, and the process terminates with a maximal sequence which has at most $h - 1$ vertices. For every fragment U obtained in the process above, we have

$$|\bigcup \mathcal{P}^U_{v_0 \ldots v_k}| \geq |U| - kCq - 2^{(K+1)C} M_k$$

and the number of all possible fragments is constant. In addition, the number of vertices in the union of the exceptional sets is constant. Consequently $|V_i'' \setminus \bigcup_{x \in D_2(V_i')} N(x)|$ is a constant. ◀

We can now describe the second phase of the algorithm.

Phase 2

Input: $G = (V, V', E)$ and $D_1 \subseteq V$

1. For every marked vertex $v \in V$ such that $d_G(v) \geq q$, construct $\mathcal{T}(v)$ consisting of all (α, q, l, K)-pseudo-covers of size at most K and the partition $V_0', V_1', \ldots, V_s'$ associated with v in G.
2. Let $V := V \setminus D_1$ and let $V'' := V' \setminus \bigcup_{v \in D_1} N(v)$.
3. For every v and every set V_i' in the partition associated with v let $V_i'' := V_i' \cap V''$. If $|V_i''| \geq M_0$, compute $D_2(V_i')$ and add all vertices from $D_2(V_i')$ to D_2.
4. For every $v' \in V'' \setminus \bigcup_{v \in D_2} N(v)$ add its mate $v \in V$ to D_3.
5. Return $D := D_1 \cup D_2 \cup D_3$.

Let Algorithm Dominating Set consist of Phase 1 followed by Phase 2. We will note a few facts about Algorithm Dominating Set which will complete our proof of Theorem 1.

First note the following:

▶ **Fact 20.** Algorithm Dominating Set *runs in* $O(1)$ *rounds in the Local model.*

Proof. Phase 1 runs in $O(1)$ rounds because computing the virtual graph G and $\mathcal{T}(v)$, in parallel for every v, requires only knowledge of vertices within distance two of v. Phase 2 runs in $O(1)$ rounds. Finding the partition is done locally by every vertex v and vertices in all sets considered in Phase 2 for v are within distance $O(1)$ of v. ◀

In addition, because of the 4th step of Phase 2, set D returned by Algorithm DominatingSet dominates V' in G and consequently V in H. Thus we have the following fact.

▶ **Fact 21.** *Set D returned by* Algorithm DominatingSet *is a dominating set in H.*

To finish our analysis, we will show that Algorithm DominatingSet returns a set of size $O(\gamma(H))$.

▶ **Fact 22.** *Let $k \in Z$ and let H be a graph with no TK_k. There is a constant $s = s(k)$ such that for the set D returned by the algorithm* Algorithm Dominating Set, $|D| \leq s\gamma(H)$.

Proof. Let $D^* \subseteq V$ be such that $|D^*| = \gamma'(G)$. From Lemma 16, $|D_1| = O(\gamma(H))$. From Claim 18, we have $D_2(V_i') \subseteq V^* \cup D^* \cup Z$ for every vertex v and every set V_i' considered in step 3. Thus $D_2 \subseteq V^* \cup D^* \cup Z$ and since $|Z| = O(\gamma(H))$, we have $|D_2| = O(\gamma(H))$. Let $X := G[V \setminus (D_1 \cup D_2), V'' \setminus \bigcup_{v \in D_2} N(v)]$ and let $v \in V \setminus (D_1 \cup D_2)$. By Claim 19, for every set V_i' with $i \geq 1$ in the partition associated with v, only a constant number of vertices L are not dominated by vertices in D_2. Since, by Fact 8 the number of sets V_i' is at most $2^{(K+1)C}$ and by Fact 9, $|V_0'| \leq Cq$, we have $d_X(v)$ bounded by some constant p for every $v \in V \setminus (D_1 \cup D_2)$. By Fact 15, $\gamma'(X) \geq |V'' \setminus \bigcup_{v \in D_2} N(v)|/p$ and at the same time vertices in $V'' \setminus \bigcup_{v \in D_2} N(v)$ can only be dominated by vertices in $V \setminus (D_1 \cup D_2)$ and so $\gamma'(X) \leq |D^*|$. Consequently, $|D_3| = |V'' \setminus \bigcup_{v \in D_2} N(v)| = O(\gamma'(X)) = O(|D^*|)$. ◀

Proof of Theorem 1. Combining Fact 20, Fact 21 and Fact 22 shows that given a graph H with no TK_h, Algorithm DominatingSet finds in L_h rounds a dominating set D such that $|D| \leq C_h \gamma(H)$ for some constants L_h and C_h which depend on h only.

As noted in the introduction, Theorem 1 in connection with methods developed in [2] (Theorem 3.4) immediately imply Theorem 2.

References

1 E. Szymańska A. Czygrinow, M. Hańćkowiak. Fast Distributed Algorithm for the Maximum Matching Problem in Bounded Arboricity Graphs. *Proc. of 20th International Symposium, ISAAC, LNCS 5878, pages 668–678*, 2009.

2 W. Wawrzyniak A. Czygrinow, M. Hańćkowiak. Fast distributed approximations in planar graphs. *International Symposium on Distributed Computing, DISC, Arcachon, France, September 2008, LNCS 5218, pages 78–92*, 2008.

3 R. Wattenhofer C. Lenzen. Minimum Dominating Set Approximation in Graphs of Bounded Arboricity. *International Symposium on Distributed Computing, DISC 2010, pages 510–524*, 2010.

4 R. Wattenhofer C. Lenzen, Y. A. Oswald. Distributed minimum dominating set approximations in restricted families of graphs. *Distrib. Comput., 26 (2), pages 119–137*, 2013.

5 S. Gutner N. Alon. Kernels for the Dominating Set Problem on Graphs with an Excluded Minor. *Electronic Colloquium on Computational Complexity, Report No. 66*, 2008.

6 S. Siebertz S. A. Amiri, S. Schmid. A Local Constant Factor MDS Approximation for Bounded Genus Graphs. *PODC 2016, Proceedings of the ACM Symp. on Principles of Distributed Computing, 227–233*, 2016.