

Leveraging Linial's Locality Limit

Christoph Lenzen and Roger Wattenhofer

Computer Engineering and Networks Laboratory (TIK)
ETH Zurich, 8092 Zurich, Switzerland
{lenzen,wattenhofer}@tik.ee.ethz.ch
www.dcg.ethz.ch

Abstract. In this paper we extend the lower bound technique by Linial for local coloring and maximal independent sets. We show that constant approximations to maximum independent sets on a ring require at least log-star time. More generally, the product of approximation quality and running time cannot be less than log-star. Using a generalized ring topology, we gain identical lower bounds for approximations to minimum dominating sets. Since our generalized ring topology is contained in a number of geometric graphs such as the unit disk graph, our bounds directly apply as lower bounds for quite a few algorithmic problems in wireless networking.

Having in mind these and other results about local approximations of maximum independent sets and minimum dominating sets, one might think that the former are always at least as difficult to obtain as the latter. Conversely, we show that graphs exist, where a maximum independent set can be determined without any communication, while finding even an approximation to a minimum dominating set is as hard as in general graphs.

1 Introduction

The recent hype about multi-hop wireless networks such as ad hoc, mesh, or sensor networks has sparked an unprecedented interest in distributed network algorithms, from inside the distributed computing community, and probably even more from outside. In the last decade reams of new distributed network algorithms have been proposed. One common theme of these algorithms is *locality*: As large networks demand fast and failure resistant algorithms, nodes should be able to make decisions solely by communicating to neighboring nodes a bounded number of times. The main challenge is to design local algorithms which can provide global guarantees. In the center of attention are classic graph optimization problems such as minimum dominating sets (MDS) and connected dominating sets, as they provide, e.g., energy-efficient backbone solutions for a variety of applications.

This abundance of distributed network algorithms is not matched by an equally rich knowledge about lower bounds and impossibility results. Indeed, on the lower bound side of locality research there are to the best of our knowledge only two results. One is a technique by Kuhn et al. [13] which proved that

many classic graph optimization problems including vertex cover, matching, or MDS cannot be polylogarithmically approximated in less than $\sqrt{\log n / \log \log n}$ time. However, the proof requires quite a peculiar “fractal” graph family barely occurring in real world problems, and certainly not in the world of wireless networks. The other is a nifty lower bound by Linial [12] who proved that computing a 3-coloring or a maximal independent set (MIS) even on a ring topology cannot be done in constant time. Indeed, using an indistinguishability argument, Linial shows that at least $\Omega(\log^* n)$ time is necessary on a ring with n nodes. Linial's lower bound is extremely weak compared to the one due to Kuhn et al., conversely the ring is a topology that may appear in virtually any network. In fact Linial's limit also holds on a simple list, hence prohibiting constant time solutions to the above problems on almost any graph families of practical use.

However, to the theorist's annoyance, Linial's bound only holds for 3-coloring and maximal independent set, leaving important problems as MDS or maximum independent set (MaxIS) approximations aside. In fact, the MDS approximation problem on Linial's ring topology offers a trivial solution: Simply take every node and you have a 3-approximation!

The first issue we address in this paper is whether one can extend Linial's lower bound towards approximation covering and packing problems such as MDS or MaxIS, such that the lower bound still holds in natural geometric graphs existing in wireless multi-hop networks. The weakest geometric graph model is the unit disk graph (UDG). Our lower bounds hold in UDGs, and henceforth in all generalizations thereof, e.g. quasi unit disk graphs, unit ball graphs, and growth-bounded graphs.

We will prove that constant approximations of the MDS problem on a quite simple family of UDG's cannot be obtained in $o(\log^* n)$ time. To the best of our knowledge, this is the first nontrivial lower bound for this problem holding in UDG's. More generally, if we allow for an $O(f(n))$ -approximation of the MDS problem in $O(g(n))$ time, the product $f(n) \cdot g(n)$ cannot be in $o(\log^* n)$. Strengthening Linial's results with respect to MIS, the same bounds are proved for the MaxIS problem on the ring.¹ Like Linial's limit our results apply to a very general computational model, where nodes can gather all information about nodes that are at most k hops away in k rounds, and may perform arbitrary local computations. In addition, all bounds still apply when allowing non-uniform algorithms, i.e., nodes to be aware of the size of the graph.

Finally we answer the question if the local complexities of both problems are in principle related to each other. In growth-bounded graphs the MIS problem is at least as hard as approximating a MDS, as a MIS is always a constant factor MDS approximation. The same order holds for all other known results. However, as is shown in the last part of the paper, this is not true in general. We will utilize a simple construction to generate a graph family for which an exact solution to the MaxIS problem is trivial, but MDS and approximations thereof stay as hard as in general graphs.

¹ Note that—as for 3-coloring—this result trivially generalizes to simple lists, since in this case most of the nodes observe the same topology as on a ring.

2 Related Work

Local algorithms have been a recurring research theme since the 80's [4, 9, 10, 11, 12]. Lately many contributions in this area have been motivated by demands of wireless ad hoc and sensor networks. The classical minimum dominating set and maximum independent set type problems are subject to quite a few basic protocols in such distributed systems. Energy consumption and communication capacities are directly affected by the quality of the given solutions. As communication ranges are limited in the systems in consideration, often the family of bounded growth graphs (or subfamilies thereof, e.g. unit disk graphs) are examined. Under this assumption the problems seem to be closely related, as a MIS becomes both a constant MDS and MaxIS approximation.

As a first highlight Luby [10] managed to compute a MIS in $O(\log n)$ time, where n is the number of nodes. The algorithm works on general graphs, however, in general graphs a MIS is neither a MDS nor a MaxIS approximation. It took until the beginning of the current century for the first distributed MDS algorithm non-trivial in both time and approximation to be published [7]. It yields a $O(\log \Delta)$ approximation in $O(\log n \log \Delta)$ time, where n is the number of nodes and Δ is the largest node degree. Kuhn et al. followed with the first constant time algorithm providing a non-trivial approximation ratio [6]. This result has been improved [5] to the currently best result for general graphs: A MDS can be approximated up to a factor of $O(\Delta^{1/\sqrt{k}} \log \Delta)$ in $O(k)$ time.

For a long time the only known lower bound for local algorithms had been Linial's $\Omega(\log^* n)$ bound on 3-coloring and MIS on the ring. Later Kuhn et al. [13] opposed the positive results by showing that in general graphs local algorithms cannot compute a polylogarithmic approximation of several optimization problems, including MDS, in less than $\Omega(\sqrt{\log n / \log \log n})$ time. Independent of and concurrent to our own results, Czygrinow et al. [15] proved the same lower bound on MaxIS approximations we show, but using a different argument. In the same work they present a randomized algorithm achieving an $(1 - \varepsilon)$ -approximation in $O(1/\varepsilon)$ time for any $\varepsilon > 0$, showing that in contrast to Linial's bound randomization does help.

Since the graphs used in the lower bound proof in [13] are complex and most unlikely to occur in practice, researchers started studying geometric graph classes like unit disk graphs (UDG's) or bounded growth graphs, which are regarded as abstractions of realistic wireless network topologies. Close-to-optimum deterministic respectively randomized MIS/MDS/MaxIS algorithms were presented by [3] and [8]. Recently, Schneider et al. [1] devised an algorithm computing a MIS on bounded growth graphs within $O(\log^* n)$ time. Our lower bounds show this bound to be tight also with respect to MDS or MaxIS approximations in bounded growth graphs, as a MIS yields constant approximations to both in this graph family. In other words, locally approximating a MDS or MaxIS in bounded growth graphs is not simpler than the special case of ascertaining a MIS.

Restricting the scope further, one can study UDG's with the nodes given global position information. In this setting even Linial's lower bound can be

beaten (e.g., in [14] a PTAS for the MDS problem is given), as the positions can be used to partition the problem into efficiently solvable local instances. Thus, the main difficulty when approximating MDS or MaxIS in UDG's is to break the symmetry of the problem, which is reflected in Linial's Limit.

3 Model and Notation

We model a network as a simple undirected graph $G = (V, E)$, where nodes represent processors and edges represent bidirectional communication links. Basically we use Linial's classic synchronous message passing model, where in one communication round each node of the network graph can send a message to each of its direct neighbors. We allow those messages to be of arbitrary size. However, at the beginning a node $v \in V$ is only equipped with information about its communication channels and a unique identifier of $O(\log n)$ size, which for simplicity we will refer to as v as well. Thus, a node can gather knowledge about node identifiers and edges between nodes at most k hops away in k communication rounds. Each processor may perform arbitrary local computations. Thus in this model an algorithm running in at most k rounds can be expressed as a function of the topology and identifiers of the (inclusive) k -neighborhood $\mathcal{N}_k^+(v) := \{w \in V \mid w \text{ is in at most } k \text{ hops distance of } v\}$ of each node v to a result $c(v)$. For an algorithm to be correct, it is required that combining the choices $c(v)$ of all $v \in V$ yields a feasible global solution of the considered problem.

We modify this standard model by dropping the assumption of uniformity, i.e., we allow nodes to know the size $n := |V|$ of the graph. Though we need this in our proofs of the claimed lower bounds, we acquire even stronger results. We will solely consider symmetric graphs, in the sense that an embedding exists where for any two nodes $v, w \in V$ and any $k \in \mathbb{N}$, the k -neighborhood $\mathcal{N}_k^+(v)$ is identical up to translation and rotation to $\mathcal{N}_k^+(w)$. Hence an Algorithm \mathcal{A} running in at most k rounds on a node $v \in V$ will be a function from $\mathcal{N}_k^+(v)$, its topology, and n to the set of possible decisions $c(v)$, independent of whether nodes can gather any *local* geometric information. This implies that our bounds also hold, e.g., when we assume an Euclidean embedding of the graph where nodes can determine the exact distance an edge bridges. As discussed in the related work section, with *global* position information better solutions become possible.

Definition 1 (Local f -approximations of MaxIS). *Given a graph $G = (V, E)$, an independent set (IS) of G is a set $I \subseteq V$ such that for all $v, w \in I$ we have $\{v, w\} \notin E$. A maximal independent set (MIS) is an independent set M so that no set $S \supset M$ can be an IS. A maximum independent set (MaxIS) is an IS of maximum cardinality. Let f be a function from \mathbb{N} to $[1, \infty) \subset \mathbb{R}$. A local f -approximation algorithm for the MaxIS problem computes for each node $v \in V$ a choice $c(v) \in \{0, 1\}$ such that $I := \{v \in V \mid c(v) = 1\}$ is an IS and for any graph G the inequality $f(n)|I| \geq |M|$ holds, where M is an arbitrary MIS of G .*

Definition 2 (Local f -approximations of MDS). Given a graph $G = (V, E)$, a dominating set (DS) of G is a set $D \subseteq V$ such that for each $v \in V \setminus D$ a $d \in D$ exists with $\{v, d\} \in E$. A minimum dominating set (MDS) is a DS of minimum cardinality. Let f be a function from \mathbb{N} to $[1, \infty) \subset \mathbb{R}$. A local f -approximation algorithm for the MDS problem computes for each node $v \in V$ a choice $c(v) \in \{0, 1\}$ such that $D := \{v \in V \mid c(v) = 1\}$ is a DS and for any graph G the inequality $|D| \leq f(n)|M|$ holds, where M is an arbitrary MDS of G .

Definition 3 (Local 3-coloring). A valid 3-coloring of a graph $G = (V, E)$ is a function $c : V \rightarrow \{r, g, b\}$ such that $c(v) \neq c(w)$ for all $\{v, w\} \in E$.

Definition 4 (Unit Disk Graph (UDG)). A Unit Disk Graph (UDG) is a graph $UDG(\iota) = (V, E)$, defined by an injective function $\iota : V \rightarrow \mathbb{R}^2$, where $E = \{\{v, w\} \in V \times V \mid 0 < \|\iota(v) - \iota(w)\|_{\mathbb{R}^2} \leq 1\}$.

Definition 5 (R_n and R_n^k). Define the ring with n nodes as $R_n := (V_n, E_n)$, where $V_n := \{v_1, \dots, v_n\}$ and $E_n := \{\{v_1, v_2\}, \dots, \{v_{n-1}, v_n\}, \{v_n, v_1\}\}$. Thus R_n is simply a circle consisting of n nodes. Denote by $R_n^k := (V_n, E_n^k)$ the k -ring with n nodes, i.e., R_n extended by all edges $\{v_i, v_j\}$ with $v_j \in \mathcal{N}_k^+(v_i) \setminus \{v_i\}$ with respect to R_n (see Figure 1).

Proposition 6. R_n^k can be realized as UDG k is controlled by the scaling.

Proof. Place all nodes equidistantly on a circle of radius $\frac{1}{2} \left(\sin\left(\frac{k\pi}{n}\right)\right)^{-1}$, as illustrated by Figure 1.




Fig. 1. R_{16}^3 . Realized as UDG k is controlled by the scaling.

4 Proofs of the Bounds

For brevity, in the subsequent analysis the term *algorithm* will refer to deterministic, local algorithms in the sense of the model described in the previous section. We show the claimed lower bounds by means of a reduction of local 3-coloring of the ring. For this problem Linial [12] proved the following bound:

Theorem 7 (Lower bound for local 3-coloring of the ring). *There is no deterministic local algorithm 3-coloring the ring R_n requiring less than $\frac{1}{2}(\log^* n - 1)$ communication rounds.*

Proof. The proof in [9] applies, as it also holds when we assume the nodes to know the size of the network n .

Naor proved an analogous result for randomized algorithms [2]. We will need the following notion:

Definition 8 ($\sigma(n)$ -alternating algorithm). *Suppose \mathcal{A} is an algorithm operating on R_n which assigns each node $v_i \in V_n$ a value $c(v_i) \in \{0, 1\}$. We call \mathcal{A} $\sigma(n)$ -alternating, if the length k of any monochromatic sequence $c(v_i) = c(v_{i+1}) = \dots = c(v_{i+k})$, indices taken modulo n , is smaller than $\sigma(n)$.*

If a $\sigma(n)$ -alternating algorithm is given, one can easily obtain a 3-coloring of the ring R_n in $O(\sigma(n))$ time:

Lemma 9 (3-coloring the marked ring). *Given a $\sigma(n)$ -alternating algorithm \mathcal{A} running in $O(\sigma(n))$ rounds, a 3-coloring of the ring can be computed in $O(\sigma(n))$ rounds.*

Proof. Recall that we identify nodes with their identifier, thus we can compare two nodes $v, w \in R_n$. We define the following algorithm for 3-coloring the ring R_n nodewise for each node $v \in V_n$:

1. Run \mathcal{A} . Let $d(v) \in \{0, 1\}$ denote the result of this run.
2. Find a pair of neighboring nodes $\{w_1, w_2\}$ with $d(w_1) \neq d(w_2)$ which is closest to v . If $v \in \{w_1, w_2\}$, set $c(v) := b$, if $d(v) = 0$, and $c(v) := r$ otherwise. Else denote by δ the distance to the closer node in $\{w_1, w_2\}$, w.l.o.g. w_1 , and set $c(v) := c(w_1)$ if $\delta \in 2\mathbb{N}$ and $c(v) := c(w_2)$ else.
3. If v has a neighbor w with $c(v) = c(w)$ and $v > w$, set $c(v) := g$.
4. If v has a neighbor w with $c(v) = c(w) = g$ and $v > w$, set $c(v)$ to the color none of the neighbors of v has.
5. Return $c(v)$.

Clearly, the running time of this algorithm is in $O(\sigma(n))$, as by assumption not more than $\sigma(n)$ consecutive nodes take the same decision $d(v)$ when running \mathcal{A} .

We now show that it yields a valid 3-coloring of R_n . In step 2 at most one of the neighbors of any node $v \in V_n$ may take the same choice, as each node chooses different from one of its neighbors. In step 3 from each pair of neighbors with the same color one chooses g . Thus only neighbors both colored with g may remain. These will be resolved in step 4, as nodes to the right and left of a pair colored by g both must have a different color than g . \square




Fig. 2. An element of \mathcal{S}_n displayed as part of a labeling of the ring. A “?” indicates that the identifier or output of the corresponding node is unspecified respectively unknown. Independent of the identifiers left of v_1 and right of $v_{g(n)+2\sigma(n)}$ all nodes v_i from $v_{g(n)+1}$ to $v_{g(n)+\sigma(n)}$ will compute $c(v_i) = 0$.

To establish our lower bounds, we construct $\sigma(n)$ -alternating algorithms out of assumed approximation algorithms for MaxIS and MDS, respectively.

Lemma 10 (Modified MaxIS approximation). *Suppose an f -approximation algorithm \mathcal{A} for the MaxIS problem on the ring R_n running in at most $g(n) \geq 1$ rounds is given, where we have $f(n)g(n) \in o(\log^* n)$. Then an $o(\log^* n)$ -alternating algorithm \mathcal{A}' requiring $o(\log^* n)$ communication rounds exists.*

Proof. As stated in the last section, we identify nodes with their identifiers. Thus, the input of \mathcal{A} when executed on R_n is a sequence of identifiers (v_1, \dots, v_n) , where no identifier occurs twice. Recall that for a single node $v_i \in V_n$, $i \in \{1, \dots, n\}$, we can express the output $c(v_i)$ of \mathcal{A} as a function of n and the subsequence of identifiers $(v_{i-g(n)}, \dots, v_i, \dots, v_{i+g(n)})$, where indices are taken modulo n . Set $\sigma(n) := 10f(n)g(n)$ and define

$$\mathcal{S}_n := \left\{ (v_1, \dots, v_{\sigma(n)+2g(n)}) \mid \forall i \in \{g(n) + 1, \dots, \sigma(n) + g(n)\} : c(v_i) = 0 \right\} , \tag{1}$$

i.e., exactly the set of sequences preventing that \mathcal{A} is $\sigma(n)$ -alternating (see also Figure 2). Note that due to the preceding observations \mathcal{S}_n is well defined, although the choices of the leading and trailing $g(n)$ many nodes may depend on further identifiers.

For n fixed we construct a sequence of identifiers for R_n . Initially we choose an arbitrary subsequence $s \in \mathcal{S}_n$ and assign the identifiers of s to $v_1, \dots, v_{|s|}$. Now suppose we already assigned labels to the nodes v_1, \dots, v_j . If there exists a sequence $s \in \mathcal{S}_n$ that can be appended to v_1, \dots, v_j without duplicating an identifier, we do so. If no further sequence fits, we add $n - j$ arbitrary identifiers not yet present in v_1, \dots, v_j to complete the labeling (v_1, \dots, v_n) of R_n . Observe that each sequence from \mathcal{S}_n added implies that at least $\sigma(n)$ additional nodes will compute $c(v) = 0$ when \mathcal{A} is run on the constructed labeling.

Assume for contradiction, that for arbitrarily large n it is possible to label R_n as described in the preceding paragraph, with at least $n - \frac{n}{5f(n)}$ identifiers stemming from sequences out of \mathcal{S}_n . By construction at least

$$\frac{\sigma(n)n}{\sigma(n) + 2g(n)} - \frac{n}{5f(n)} \tag{2}$$

many nodes compute $c(v) = 0$ when \mathcal{A} is passed such a labeling as input. Thus we have

$$|I| \leq n - \left(\frac{\sigma(n)n}{\sigma(n) + 2g(n)} - \frac{n}{5f(n)} \right) \leq \frac{2n}{5f(n)}, \tag{3}$$

where $I = \{v \in V_n \mid c(v) = 1\}$ denotes the resulting independent set. Since the size of a MaxIS of R_n is $\lfloor \frac{n}{2} \rfloor$, this contradicts the assumption that \mathcal{A} is an f -approximation algorithm to the MaxIS problem on R_n .

Thus, an $n_0 \in \mathbb{N}$ must exist, such that for all $n \geq n_0$ we may choose a maximal set of disjoint sequences $\{s_1, \dots, s_{j_n}\} \subset S_n$ such that

$$\left| Id(n) \setminus \left(\bigcup_{i=1}^{j_n} s_i \right) \right| \geq \frac{n}{5f(n)}, \tag{4}$$

where $Id(n)$ is the set of admissible identifiers for nodes on R_n . In other words, at least $\frac{n}{5f(n)}$ identifiers remain which cannot form a further sequence from S_n . W.l.o.g. we may restrict $Id(n)$ such that $|Id(n)| = n$, as \mathcal{A} must yield correct results for any admissible set of identifiers. Hence, by setting $n' := \max\{n_0, 5f(n)n\}$, we can define an injective relabeling function $r_n : Id(n) \rightarrow Id(n')$ such that no sequence $s \in S_{n'}$ is completely contained in the image of r_n .

The algorithm \mathcal{A}' claimed to exist now consists of redefining all identifiers by r_n and simulating a run of \mathcal{A} on the modified instance, where instead of n the algorithm is given n' as the number of nodes. As $g(n) \leq g(n)f(n) \in o(\log^* n)$, the running time $g(n')$ of \mathcal{A}' is certainly in $o(\log^* n)$ as well. Since \mathcal{A} computes an IS, no two consecutive nodes are assigned $c(v) = 1$.² By definition no sequence from $S_{n'}$ is contained completely in the image of r_n , hence at most $\sigma(n') - 1 \in O(f(n')g(n')) \subset o(\log^* n') = o(\log^* n)$ consecutive nodes compute $c(v) = 0$. Thus \mathcal{A}' is $o(\log^* n)$ -alternating as desired. \square

We will need a similar result for the MDS approximation problem. In a ring topology choosing every node is a trivial, yet constant MDS approximation. Hence we will resort to the slightly more complex topology of R_n^k , which still is present in UDG’s.

Lemma 11 (Modified MDS approximation). *Assume an f -approximation algorithm \mathcal{A} for the MDS problem on UDG’s running in at most $g(n) \geq 1$ rounds is given, where $f(n)g(n) \in o(\log^* n)$. Then an $o(\log^* n)$ -alternating algorithm \mathcal{A}' requiring $o(\log^* n)$ communication rounds exists.*

Proof. We will extend the proof of Lemma 10. In a simple ring topology choosing all nodes is a constant MDS approximation. This is not true in R_n^k . Set $\sigma_k(n) := \max\{f(n), k\}g(n)$ and define

$$S_n^k := \left\{ (v_1, \dots, v_{\sigma_k(n)+2kg(n)}) \right. \\ \left. \mid \forall i \in \{kg(n) + 1, \dots, \sigma_k(n) + kg(n)\} : c(v_i) = 1 \text{ on } R_n^k \right\}, \tag{5}$$

² Independence is a local property, which is only affected by the input of \mathcal{A} at a node v and its neighbors, i.e., the identifiers of the $g(n) + 1$ -neighborhood of v , and n . Since any subsequence of identifiers $(v_{i-g(n)-1}, \dots, v_i, \dots, v_{i+g(n)+1}) \subset Id(n')$ may occur on $R_{n'}$, the output of \mathcal{A}' must still form an IS.

i.e., the set of sequences of identifiers yielding $\sigma_k(n)$ consecutive nodes taking the decision $c(v) = 1$ when \mathcal{A} is executed on R_n^k , where the choices of the leading and trailing $kg(n)$ many nodes may also depend on identifiers not in the considered sequence. As the decision of any node v depends only on identifiers of nodes in $\mathcal{N}_{kg(n)}^+(v)$, \mathcal{S}_n^k is well defined.

We make a case decision. The first case is that a $k_0 \in \mathbb{N}$ exists allowing a similar relabeling procedure as in Lemma 10. More precisely, $k_0, n_0 \in \mathbb{N}$ exist, such that for $n \geq n_0$ at most $\frac{n}{2}$ identifiers can simultaneously participate in disjoint sequences from $\mathcal{S}_n^{k_0}$ in a valid labeling of $R_n^{k_0}$. Thus, by setting $n' := \max\{n_0, 2n\}$, we can define \mathcal{A}' to simulate a run of \mathcal{A} on $R_{n'}^{k_0}$ and return the computed result. Each simulated round of \mathcal{A} will require k_0 communication rounds, thus the running time of \mathcal{A}' is bounded by $k_0g(n') \in o(\log^* n)$. At most $2k_0$ consecutive nodes will compute $c(v) = 0$, as \mathcal{A} determines a DS, and by definition of $\mathcal{S}_{n'}^{k_0}$ at most $\sigma_{k_0}(n') - 1 \in O(f(n')g(n')) \subset o(\log^* n') = o(\log^* n)$ consecutive nodes take the decision $c(v) = 1$.

The second case is that no pair $k_0, n_0 \in \mathbb{N}$ as assumed in the first case exists. Similar to the proof of Lemma 10, we construct a labeling of R_n^k with at least $\frac{n}{2}$ many identifiers stemming from sequences in \mathcal{S}_n^k . Running \mathcal{A} on this instance will yield at least

$$\frac{\sigma_k(n)n}{2(\sigma_k(n) + 2kg(n))} \geq \frac{n}{6} \in \Omega(n) \tag{6}$$

many nodes choosing $c(v) = 1$. On the other hand, varying k , we get minimum dominating sets with $O(\frac{n}{k})$ many nodes. Define n_k to be the minimum n , such that it is possible to construct labelings of R_n^k with $\frac{n}{2}$ identifiers from sequences in \mathcal{S}_n^k . Since \mathcal{A} is an f -approximation algorithm to the MDS problem on R_n^k , we conclude

$$f(n_k) \in \Omega(k) . \tag{7}$$

We choose $k(n)$ minimum such that $n' := 2n < n_{k(n)}$, allowing to define an injective relabeling function $r_n : Id(n) \rightarrow Id^{k(n)}(n')$, such that no element of $\mathcal{S}_{n'}^{k(n)}$ lies completely in the image of r_n . Here $Id(n)$ and $Id^k(n)$ denote the sets of admissible identifiers of R_n and R_n^k , respectively, where w.l.o.g. we assume $|Id(n)| = |Id^k(n)| = n$. We define \mathcal{A}' to be the algorithm operating on R_n , but returning the result of a simulated run of \mathcal{A} on $R_{n'}^{k(n)}$, where we relabel all nodes $v \in R_n$ by $r_n(v)$. By definition of $k(n)$ we have $n_{k(n)-1} \leq n'$. Together with (7) this yields

$$k(n) = (k(n) - 1) + 1 \in O(f(n_{k(n)-1}) + 1) = O(f(n')) = O(f(n)) , \tag{8}$$

since f grows asymptotically sublinear. Hence we can estimate the running time of \mathcal{A}' by $k(n)g(n') \in O(f(n)g(n))$, using that g grows asymptotically sublinear as well.

Since the simulated run of \mathcal{A} yields a dominating set, at worst $2k(n) \in O(f(n)) \subseteq O(f(n)g(n))$ many consecutive nodes may compute $c(v) = 0$. By the definitions of \mathcal{S}_n^k and r_n at most $s_{k(n)}(n') - 1 < \max\{f(n'), k(n)\}g(n') \in O(f(n)g(n))$ consecutive nodes may take the decision $c(v) = 1$. Thus \mathcal{A}' is $o(\log^* n)$ -alternating, as claimed. \square

The two main theorems follow immediately from the preceding statements.

Theorem 12 (Lower bound on MaxIS approximations). *No f -approximation algorithm to the MaxIS problem on the ring R_n running in at most $g(n) \geq 1$ communication rounds with $f(n)g(n) \in o(\log^* n)$ exists.*

Proof. Assuming the contrary, we can combine Lemma 10 and Lemma 9 to construct an algorithm contradicting Theorem 7.

Theorem 13 (Lower bound on MDS approximations on UDG's). *No f -approximation algorithm to the MDS problem on UDG's running in at most $g(n) \geq 1$ time with $f(n)g(n) \in o(\log^* n)$ exists.*

Proof. Assuming the contrary, we can combine Lemma 11 and Lemma 9 to construct an algorithm contradicting Theorem 7.

Note that $g(n) \geq 1$ is just a formal restriction. If $g(n) = 0$ for infinitely many n , the approximation ratio f must be trivial, i.e., $f(n) \notin o(n)$.

5 MaxIS Graphs

In this section we address the question whether the difficulties of MaxIS and MDS approximations are related in general. As shown in Theorem 12, on a ring topology one cannot compute a constant MaxIS approximation in constant time. Conversely, for MDS this is trivially possible by taking all nodes to be in the DS. On UDG's³ a MaxIS, or even any maximal independent set, is a constant MDS approximation, since the number of independent neighbors of a node is bounded by a constant. Schneider et al. [1] showed how a maximal independent set can be computed on UDG's in $O(\log^* n)$ time. As direct consequence of Theorem 7 this bound is tight, as a maximal independent set on a ring allows for a 3-coloring in a single round. Moreover, Theorem 13 shows this bound to be tight also with respect to MDS approximations on UDG's.

In the light of these results one might expect that approximating MaxIS is always at least as difficult as approximating MDS. On the contrary, we will now present an example showing that finding a MaxIS can be trivial (see Lemma 15), while computing a MDS remains as hard as in general graphs (see Lemma 16). Towards this end, we construct a family of graphs for which an exact solution for the MaxIS problem can be given without any communication and without knowing the problem size n . Conversely, MDS approximation on general graphs reduces to MDS approximation on this graph family. Kuhn et al. [13] showed that this problem cannot be approximated well in less than $\Omega\left(\sqrt{\log n / \log \log n}\right)$ time. The graph class examined is constructed for this special purpose, being of no practical relevance. Indeed, the solution to the MaxIS problem will already be encoded in the local topology of the graph. The main impact is that the (local) complexities of MDS and MaxIS type problems depend strongly on the considered types of graphs.

³ More generally, on bounded growth graphs.

Definition 14 (MaxIS Graphs). Given an arbitrary graph $G = (V, E)$, we construct a new graph $H = (V_H, E_H)$ from G . We define $V_H := V_1 \cup V_2 \cup V_3 \cup V_4$, where the V_i , $i \in \{1, 2, 3, 4\}$, are four disjoint copies of V . Denote by $v_i \in V_i$ the nodes that are copies of a node $v \in V$. The edge set connects all nodes $v_1 \in V_1$ and $v_2 \in V_2$ to all w_i that are copies of some $w \in \mathcal{N}_1^+(v)$, i.e.,

$$E_H := \{ \{v_i, w_j\} \mid w \in \mathcal{N}_1^+(v), i \in \{1, 2\}, j \in \{1, 2, 3, 4\}, v_i \neq w_j \} . \quad (9)$$

Thus, the subgraphs induced by V_1 and V_2 are copies of G , and each $v_3 \in V_3$ (and $v_4 \in V_4$, respectively) is connected exactly to both copies of $\mathcal{N}_1^+(v)$ in V_1

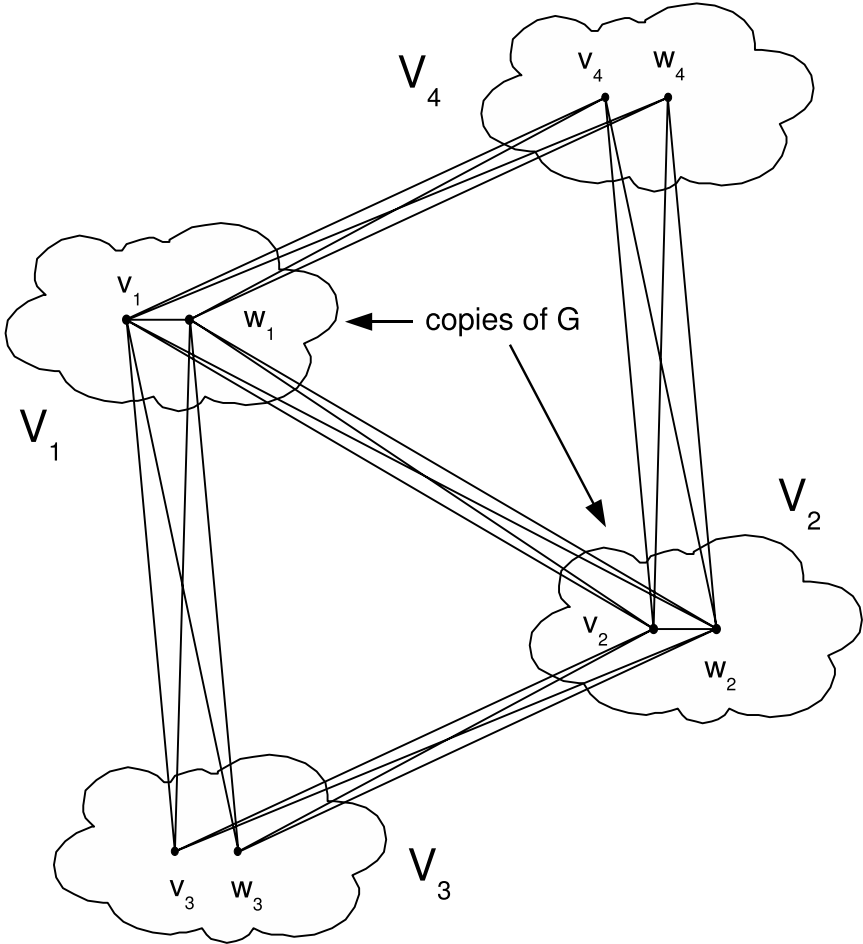


Fig. 3. Overview of the structure of a MaxIS Graph H constructed out of some graph $G = (V, E)$. The displayed vertices and edges form the subgraph S of H induced by the copies of two nodes $v, w \in V$, where $\{v, w\} \in E$. This subgraph is the complete graph without any edges between two nodes both in $V_3 \cup V_4$. Thus the degree of a node in $V_1 \cup V_2$ with respect to S equals the one of a node in $V_1 \cup V_2$ plus three.

and V_2 . An illustration is given in Figure 3. Any graph that can be constructed in this way is a MaxIS Graph.

We will prove our statements by giving a simple criterion allowing to compute a MaxIS on this class of graphs in zero rounds, and a local reduction of the general MDS approximation problem to MaxIS Graphs.

Lemma 15 (Local computation of a MaxIS on MaxIS Graphs). *The set $\{v \in V \mid |\mathcal{N}_1^+(v)| \bmod 2 = 1\}$ is a MaxIS for any MaxIS Graph. It can be determined locally, without communication.*

Proof. We use the notation of Definition 14. The set of nodes $I := V_3 \cup V_4$ is independent by construction. A node $v_i \in V_i$, $i \in \{3, 4\}$, has $2|\mathcal{N}_1^+(v)|$ many neighbors in $V_1 \cup V_2$, hence $|\mathcal{N}_1^+(v_i)|$ is odd. On the other hand, for a node $v_i \in V_i$, $i \in \{1, 2\}$, we have $|\mathcal{N}_1^+(v_i)| = 4|\mathcal{N}_1^+(v)|$, which is even. Thus $I = \{v \in V \mid |\mathcal{N}_1^+(v)| \bmod 2 = 1\}$ holds. Moreover, since the sequence of nodes (v_1, v_3, v_2, v_4) forms a cycle, at most two of them may participate in an IS, thus I is a MaxIS.

As nodes know the number of their neighbors, each node can determine whether it is in I or not without any communication. □

Lemma 16 (Reduction of MDS to MDS on MaxIS Graphs). *We use the notation of Definition 14. Given an f -approximation algorithm \mathcal{A} to the MDS problem on MaxIS Graphs running in $g(n)$ time, we can define the following algorithm \mathcal{A}' operating on an arbitrary graph G :*

1. Simulate a run of \mathcal{A} on the MaxIS Graph H constructed from G .
2. Return for each node $v \in V$ $c(v) = 1$ if \mathcal{A} computed $c(v_i) = 1$ for some $i \in \{1, 2, 3, 4\}$, and $c(v) = 0$ else.

Algorithm \mathcal{A}' is an $f(4n)$ -approximation algorithm to the MDS problem running in $g(4n)$ rounds. Up to constants it is as efficient as the original one, i.e., it is an $O(f(n))$ approximation running in $O(g(n))$ time.

Proof. Since \mathcal{A} computes a MDS of H , Algorithm \mathcal{A}' will return a MDS of G : If $v_1 \in V_1$ is covered by $w_i \in V_i$ for some $i \in \{1, 2, 3, 4\}$, w will cover v in G . Thus \mathcal{A}' works correctly.

If M is a MDS of G , the copy $M_1 := \{m_1 \mid m \in M\}$ is a MDS of H . Conversely, any node $v_1 \in V_1$ can only be covered by copies of nodes $w \in \mathcal{N}_1^+(v)$, hence the sizes of minimum dominating sets of G and H coincide. Thus, if \mathcal{A} is an $f(n)$ -approximation algorithm to the MDS problem on MaxIS Graphs, \mathcal{A}' will be an $f(4n)$ -approximation algorithm to the MDS problem on general graphs. Since any MDS approximation algorithm will trivially reach at least an approximation ratio of n , \mathcal{A}' is an $O(f(n))$ -approximation.

Having the nodes of G simulate a run of \mathcal{A} on H does not require any additional communication rounds. As we have no restrictions to message sizes, we can simply append the information which edge in H is used to all communications, while each node $v \in V$ simulates v_i , $i \in \{1, 2, 3, 4\}$. This is a simple task, as the neighbourhood of each v_i is determined solely by $\mathcal{N}_1^+(v)$, and we have no restrictions to local

computations. Hence, if \mathcal{A} runs in $g(n)$ time, \mathcal{A}' will require at most $g(4n)$ time to complete. As the diameter of any graph is bounded by the number of nodes n , any algorithm in our computational model can be realized using at most n communication rounds. Thus \mathcal{A}' needs at worst $O(g(n))$ rounds to complete. \square

Finally we conclude that the local complexities of the MaxIS and MDS approximation problems are incomparable.

Theorem 17. *Assume an f_1 -approximation algorithm to the MaxIS problem and an f_2 -approximation algorithm to the MDS problem, both on a family of graphs \mathcal{F} , are given. Denote by $g_1(n) \geq 1$ and $g_2(n) \geq 1$ bounds for their running times. Furthermore assume the products $p_1(n) := f_1(n) \cdot g_1(n)$ and $p_2(n) := f_2(n) \cdot g_2(n)$ are minimum, i.e., the algorithms are optimum in this sense. Then neither $p_1 \in O(p_2)$ nor $p_2 \in O(p_1)$ holds independent of \mathcal{F} .*

Proof. The lower bound of Kuhn et al. [13] and Lemma 16 show that on MaxIS Graphs we have $p_2 \in \Omega\left(\sqrt{\log n / \log \log n}\right)$, while Lemma 15 yields $p_1 \in O(1)$. Conversely, for a ring topology we trivially have $p_2 \in O(1)$, while Theorem 12 gives $p_1 \notin o(\log^* n)$, implying $p_1 \notin O(1) = O(p_2)$. \square

6 Conclusion

In this paper we extended Linial's lower bound to the well-known MDS and MaxIS problems on UDG's. The product between running time and approximation quality of any deterministic local algorithm for these problems cannot be in $o(\log^* n)$.

In a couple of graph classes, especially geometric graphs, a MIS is a special case of the MDS approximation problem. Consequently one might believe that coming up with a distributed algorithm for MIS is harder, or at least not simpler, than for MDS. In the second part of the paper we showed that this is not always true. The constructed MaxIS graphs demonstrate that the two problems are generally incomparable: In this graph class a MaxIS can be "computed" locally without any communication, while any MDS approximation algorithm on MaxIS graphs could be used to solve the problem on general graphs. Adding one more tessera to the picture, the extension of Linial's lower bound to the MaxIS approximation problem holds on virtually any graph class.

We hope that our findings will help to get a better understanding of distributed algorithms, eventually permitting a classification of local problems reflecting their complexity.

References

1. Schneider, J., Wattenhofer, R.: A Log-Star Distributed Maximal Independent Set Algorithm for Growth-Bounded Graphs. In: Proc. Twenty-Seventh Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (2008)
2. Naor, M.: A Lower Bound on Probabilistic Algorithms for Distributive Ring Coloring. SIAM Journal on Discrete Mathematics 4(3), 409–412 (1991)

3. Kuhn, F., Moscibroda, T., Wattenhofer, R.: On the Locality of Bounded Growth. In: Proc. 24th ACM Symposium on the Principles of Distributed Computing (2005)
4. Cole, R., Vishkin, U.: Deterministic Coin Tossing with Applications to Optimal Parallel List Ranking. *Information and Control* 70(1), 32–53 (1986)
5. Kuhn, F., Moscibroda, T., Wattenhofer, R.: The Price of Being Near-Sighted. In: Proc. 17th ACM-SIAM Symposium on Discrete Algorithms (2006)
6. Kuhn, F., Wattenhofer, R.: Constant-Time Distributed Dominating Set Approximation. In: Proc. 22nd ACM Symposium on the Principles of Distributed Computing (2003)
7. Jia, L., Rajaraman, R., Suel, T.: An Efficient Distributed Algorithm for Constructing Small Dominating Sets. *Distributed Computing* 15(4), 193–205 (2002)
8. Gfeller, B., Vicari, E.: A Randomized Distributed Algorithm for the Maximal Independent Set Problem in Growth-Bounded Graphs. In: Proc. 26th annual ACM symposium on Principles of distributed computing (2007)
9. Peleg, D.: *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2000)
10. Luby, M.: A Simple Parallel Algorithm for the Maximal Independent Set Problem. *SIAM J. Comput.* 15(4), 1036–1055 (1986)
11. Naor, M., Stockmeyer, L.: What Can Be Computed Locally? *SIAM J. Comput.* 24(6), 1259–1277 (1995)
12. Linial, N.: Locality in Distributed Graph Algorithms. *SIAM J. Comput.* 21(1), 193–201 (1992)
13. Kuhn, F., Moscibroda, T., Wattenhofer, R.: What Cannot Be Computed Locally. In: Proc. 23rd annual ACM symposium on Principles of distributed computing (2004)
14. Wiese, A., Kranakis, E.: Local PTAS for Independent Set and Vertex Cover in Location Aware Unit Disk Graphs. In: Proc. 4th IEEE/ACM International Conference on Distributed Computing in Sensor Systems (2008)
15. Czygrinow, A., Hańćkowiak, M., Wawrzyniak, W.: Fast distributed approximations in planar graphs. In: Proc. 22nd International Symposium on Distributed Computing (2008)