

Meta-Theorems for Cuttable Distributed Problems

MARTHE BONAMY, LaBRI, University of Bordeaux, CNRS, France

AVINANDAN DAS, Aalto University, Finland

CYRIL GAVOILLE, LaBRI, University of Bordeaux, CNRS, France

TIMOTHÉ PICAVET, LaBRI, University of Bordeaux, CNRS, France

JUKKA SUOMELA, Aalto University, Finland

ALEXANDRA WESOLEK, LaBRI, University of Bordeaux, CNRS, France

We prove that given any α -approximation LOCAL algorithm for MINIMUM DOMINATING SET (MDS) on planar graphs, we can construct an $f(g)$ -round $(3\alpha + 1)$ -approximation LOCAL algorithm for MDS on graphs embeddable in a given Euler genus- g surface. Heydt et al. [*European Journal of Combinatorics* (2025)] gave an algorithm with $\alpha = 11 + \varepsilon$, from which we derive a $(34 + \varepsilon)$ -approximation algorithm for graphs of genus g , therefore improving upon the current state of the art of $24g + O(1)$ due to Amiri et al. [*ACM Transactions on Algorithms* (2019)]. It also improves the approximation ratio of $91 + \varepsilon$ due to Czygrinow et al. [*Theoretical Computer Science* (2019)] in the particular case of orientable surfaces.

We generalize this result into two directions: (1) by considering other graph problems studied in Distributed Computing such as MINIMUM k -TUPLE DOMINATING SET, for which constant-round approximation algorithms were known for planar graphs, but not for graphs of bounded genus; and (2) by considering graph classes beyond bounded genus graphs, called *locally nice*, and relying on the asymptotic dimension of the class. We prove these results by a series of meta-theorems about *cuttable* minimization problems with constant-round approximation LOCAL algorithms. Roughly speaking, in cuttable problems, one can systematically extract small subgraphs whose solutions are in proportion to the global solution restricted to the neighbourhood of the subgraph.

CCS Concepts: • **Theory of computation** → **Distributed computing models**; **Graph algorithms analysis**.

Additional Key Words and Phrases: distributed algorithm, meta-theorem, LOCAL model, dominating set

1 Introduction

MINIMUM DOMINATING SET (MDS) is a famous minimization problem on graphs, known to be NP-complete even in cubic planar graphs [GJ79, KYK80]. The goal is to find a smallest subset of vertices of the input graph that intersects all radius-1 balls of the graph.

In this paper, we focus our attention on distributed algorithms that can approximate MDS on the graph underlying the topology of the network. Due to the covering property of a dominating set, the problem and its variants (like CONNECTED DOMINATING SET [WAF02]) get increasingly more attention in Networking, in particular for mobile and ad-hoc networks. Not only are mobile networks important for point-to-point communications, but specialized ad-hoc networks, such as sensor networks, are important for environmental monitoring tasks. We refer to [KW05, vRWZ09] for extended discussions about the motivations of MDS for routing purposes of such networks.

In the LOCAL model, approximating MDS up to a constant factor in general n -vertex graphs is known to require $\Omega(\sqrt{\log n / \log \log n})$ rounds [KMW16, CL21]. On the positive side, it is possible to $(1 + \varepsilon)$ -approximate MDS for any graph in $\text{poly}(\varepsilon^{-1} \log n)$ rounds by combining the techniques of [GKM17] and of [RG20, Corollary 3.11].

For specific graphs, better round complexities can be achieved; cf. [LPW13, Suo13] for a large collection of results (including unit-disk graphs, and planar graphs as a basis of the Gabriel graph

model widely used in ad-hoc networks [WY07]). For instance, $O(\log^* n)$ rounds suffice for planar graphs [CHW08], or more generally for K_t -minor-free graphs [CHW18] and for sub-logarithmic expansion graphs¹ [ASS19], and $o(\log^* n)$ rounds are not sufficient to get a $(1 + \varepsilon)$ -approximation of MDS on cycles [CHW08] or unit-disk graphs [LW08]. So, achieving constant-round algorithms must be at the price of relaxing the $(1 + \varepsilon)$ approximation ratio.

For planar graphs, the quest for constant-approximation and constant-round algorithms seems to start with [LOW08][Len11, Chp. 13][LPW13], with a 130-approximation. Since then, a long line of research has been aimed at improving this ratio. In short, the best approximation ratio to date is $11 + \varepsilon$, due to [HKOdM⁺25]. We refer to the nice survey of [HKOdM⁺25] for planar graphs and sub-families, including lower bounds.

In the meantime, the quest for constant-approximation ratio and constant-round distributed algorithms has been proposed for larger classes of graphs. However, very few examples are known of graph classes $\mathcal{C}(p)$, depending on some fixed parameter p , where MDS can be solved by a LOCAL algorithm with round complexity $f(p)$ and truly-constant approximation ratio (independent of p).

Larger graph classes that make good candidates for extending planar graphs are H -minor-free graphs, with some graph H not limited to K_5 or $K_{3,3}$. For K_p -minor-free graphs, we only know of an exponential approximation ratio in p [KSV21, HKOdM⁺25], whereas a linear approximation ratio could be possible. For $K_{3,p}$ -minor-free graphs [HKOdM⁺25], the dependency is much better (linear in p), but still not truly constant. Very recently, it was shown [BGPW25] that, for $K_{2,p}$ -minor-free graphs, the approximation ratio is 50 regardless of the value of p .

On the other end of the spectrum, the narrowest way to meaningfully extend planar graphs is through embeddable graphs² on surfaces³ of Euler genus g . Such surfaces can be obtained from a sphere by adding $g/2 \geq 0$ handles⁴ (if they are orientable) or by adding $g \geq 1$ cross-caps⁵ (if they are nonorientable). The *Euler genus* of a graph G is the minimum number g such that G embeds on a surface of Euler genus g . In this context, [ASS19, Theorem 3.11] proposed an $(24g + O(1))$ -approximation algorithm with constant-round complexity (actually linear in g). For graphs of orientable genus g , i.e., embeddable on an orientable surface of genus g , [CHWW19] designed a constant-round algorithm that returns a dominating set of size at most $91M + 76g - 66$, where M is the optimal size. By adding a simple extra brute-force step in their algorithm, we can convert it into a $(91 + \varepsilon)$ -approximation (cf. Proposition D.2), so with a ratio independent of g . However, the result of [CHWW19] does not transfer to graphs of bounded Euler genus, as for any g , there are graphs embeddable on the projective plane⁶ (so of Euler genus 1) that cannot be embedded on any orientable surface of genus g (so of orientable genus at least $g + 1$). So, for these graphs, the approximation ratio will be $91 + \varepsilon$, but with the number of rounds possibly depending on the number of vertices⁶.

Graphs of Euler genus g are included in $K_{3,2g+3}$ -minor-free graphs, because the Euler genus of $K_{3,2g+3}$ is at least $g + 1$ (cf. [MT01, Theorem 4.4.7]) and the class of Euler genus- g graphs is closed under taking minors. We observe that the approximation ratio for $K_{3,p}$ -minor-free graphs, for $p = 2g + 3$, due to [HKOdM⁺25, Theorem 2.2] provides an approximation ratio of $O(\sqrt{g})$. Indeed,

¹Roughly speaking, the expansion of a graph G is a function f that, for any r , bounds the maximum edge density of a depth- r minor of G by $f(r)$, where a depth- r minor is a minor of G where each branch set has radius at most r .

²I.e., that can be drawn on the surface without edge crossing, like planar graphs for the sphere.

³That are compact, connected 2-dimensional manifold without boundary.

⁴By adding a handle, we mean that two disjoint disks of the sphere are replaced by a cylinder.

⁵By adding a cross-cap, we mean that a disk of the sphere is replaced by a Möbius strip.

⁶For instance, a cycle with $n = 2g + 6$ vertices where opposite vertices are connected by an extra edge as Euler genus 1 but orientable genus at least $g + 1 = n/2 - 2$, cf. [ABY63].

the ratio of their algorithm is precisely $(2 + \varepsilon)(2\nabla_1 + 1)$, where ∇_1 is the maximal edge density of a depth-1 minor of any graph of the class. As we can check (cf. Proposition D.3) that $\nabla_1 \leq \sqrt{3g/2} + 3$, the approximation ratio is at most $4\sqrt{3g/2} + 14 + \varepsilon$. As it may occur that $\nabla_1 \geq \sqrt{3g/2} - O(1)$ for some Euler genus- g graph (cf. Proposition D.3), the best known approximation ratio for MDS in Euler genus- g graphs is $\Theta(\sqrt{g})$, and $91 + \varepsilon$ for orientable genus- g graphs.

It is important to note that no LOCAL algorithms for Euler genus- g graphs can achieve simultaneously constant round complexity and constant approximation ratio. Indeed, by taking $n = \Omega(\sqrt{g})$, since there exists n -vertex graphs of Euler genus $\Theta(n^2)$, the $\Omega(\sqrt{\log n / \log \log n})$ lower bound of [KMW16, CL21] yields a lower bound of $\Omega(\sqrt{\log g / \log \log g})$.

Our Contributions. In the remainder of the paper, by “LOCAL algorithm”, we mean deterministic distributed algorithm in the classical LOCAL model.

- We propose a $(34 + \varepsilon)$ -approximation LOCAL algorithm for MDS in Euler genus- g graphs whose round complexity depends only on g (cf. Corollary 4.3). This improves approximation ratios in both the orientable (previously $91 + \varepsilon$) and non-orientable (previously $24g + O(1)$) genus case. The analysis of the algorithm uses the fact that such graphs have asymptotic dimension at most two.
- This result follows directly from more general meta-theorems. More generally, we prove that any constant-round α -approximation for MDS on planar graphs can be turned into a $f(g)$ -round $(3\alpha + 1)$ -approximation algorithm for MDS on Euler genus- g graphs, for some function f .
- We extend this result in two ways. First, we show that our meta-theorems hold for all local minimization problems that are well-behaved (including MDS and its variants). We provide a simple graph-theoretic property on the problem that is sufficient for our meta-algorithms to apply. This significantly simplifies proofs and circumvents having to reason about various algorithms.
- We apply our results on problems such as MINIMUM k -TUPLE DOMINATING SET (cf. Corollary 4.5), obtaining constant-ratio approximations on problems where none were previously known. We also obtain a ratio of $O(\sqrt{g})$ for MINIMUM CONNECTED DOMINATING SET in Euler genus- g graphs.
- Second, we prove that our meta-theorems hold on much broader “locally nice” graph classes, such as locally bounded-genus graphs, or locally bounded-genus graphs with a bounded number of apex vertices (whose unconstrained neighborhoods can break the surface embeddability of the graph).
- We prove three versions of our meta-theorem (see Meta-Theorem I, Meta-Theorem II and Meta-Theorem III), making it easily applicable and adapted to different situations.

Unlike similar notions, such as the classical network decomposition [AGLP89], no preliminary construction is required for any of our algorithms. Obviously, they depend on the parameters of the graph classes on which they are applied, like the Euler genus g and/or the asymptotic dimension of the class with its control function.

Organization. Basic notations and asymptotic dimension are introduced in Section 2. Section 3 derives a 616-approximation for MDS on bounded-genus graphs, gently introducing the ideas and motivations of the later meta-theorems. Finally, Section 4 presents our different meta-theorems, and their applications to different problems of the literature.

2 Preliminaries

Graphs and Domination. A subset $X \subseteq V(G)$ *dominates* S if $S \subseteq N[X]$. We denote by $\text{MDS}(G, S)$ the minimum size of a subset of $V(G)$ dominating S . Note that $\text{MDS}(G, V(G))$, denoted by $\text{MDS}(G)$ for short, is the size of a minimum dominating set for G .

Asymptotic Dimension. In this section, we focus on defining asymptotic dimension⁷ (a non-trivial task, as it happens) and explaining how to exploit it, in the hope that others may be able to exploit it in turn.

A graph G with a spanning subgraph H is *k -colorable δ -bounded in H* if each vertex of G can be assigned a color from $\{0, \dots, k-1\}$ so that the distance in H between any two vertices taken in a monochromatic path of G is at most δ . In other words, all monochromatic connected components of G must have a weak diameter in H at most δ . A proper k -coloring of G is nothing else than a k -coloring 0-bounded in H . If H has diameter D , then it is 1-colorable D -bounded in H .

The asymptotic dimension of a graph class \mathcal{C} is a non-negative integer denoted by $\text{asdim}(\mathcal{C})$, and related to the k -coloring boundedness of the r -th power of G . It is a notion closely related to *sparse partitions*⁸ and *weak sparse covers* (see [BBE⁺24, §1.8] for precise connection between these notions). Recall that the r -th power of G , denoted by G^r , is the graph obtained from $V(G)$ by adding edges between pairs of vertices at distance at most r in G .

Rather than giving its standard definition, we prefer the following equivalent one⁹:

Proposition 2.1 ([BBE⁺24, Proposition 1.17]). *For every graph class \mathcal{C} , $\text{asdim}(\mathcal{C}) \leq d$ if and only if there exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$, called control function, such that for every $G \in \mathcal{C}$ and integer $r \geq 1$, G^r is $(d+1)$ -colorable $f(r)$ -bounded in G .*

Any associated control function f that makes $\text{asdim}(\mathcal{C}) = d$ is also called a d -dimensional control function of \mathcal{C} . The Assouad-Nagata dimension of \mathcal{C} has a similar definition, except that the control function must be linear. A crude example is that paths have Assouad-Nagata dimension 1, because, for each $r \geq 1$, it suffices to alternately color subsets of r consecutive vertices with two colors, see Fig. 1.




Fig. 1. An example that shows that P^r is 2-colorable $(r-1)$ -bounded in P for $r=3$ (on the left), and an example of a 3-coloring of grids for $r=2$ and $f(r)=2(r-1)=2$ (on the right).

⁷Asymptotic dimension was originally defined on metric spaces by Gromov [Gro93] in the field of geometric group theory.

⁸However, those notions follow a different philosophy. In the study of asymptotic dimension, we want to minimize the dimension. In the study of sparse partitions and covers, the main goal is to minimize a function of the dimension and the diameter of the sets.

⁹In the original Proposition 1.17 of [BBE⁺24], G^r is $(d+1)$ -colorable $f(r)$ -bounded in G^r , instead of in G . By following this original definition, the function f is not a control function, whereas in Proposition 2.1 it is.

More generally, d -dimensional grids have a d -dimensional control function $f(r) = d \cdot (r - 1)$.

Trees, and more generally graph classes of bounded treewidth (resp. layered treewidth) have asymptotic dimension 1 (resp. 2). Planar graphs, and more generally the classes of H -minor-free graphs (for any fixed H) have asymptotic dimension 2 as shown by [BBE⁺24]: the dependency in H only shows in the control function f . Dense graph classes may also have small asymptotic dimension. For example, it is sufficient that there is a quasi-isometry into a class having small asymptotic dimension¹⁰. The dense class of chordal graphs, being quasi-isometric to the class of trees, has asymptotic dimension 1.

Intuitively, each connected component of a color class in G^r can be viewed as a “simple” component of the graph, that we already understand well. In the setting of LOCAL algorithms with constant round complexity $r/2$, each vertex can observe only its distance- $(r/2)$ neighborhood. If the underlying graph class has asymptotic dimension d , this neighborhood intersects at most $d + 1$ such simple parts. As a result, we can derive strong approximation guarantees, since each vertex’s view of the graph remains constrained and conceptually simple. We observe that all our algorithms make use of the control function f , but not for any particular pre-computing coloring or partitioning, so that the asymptotic dimension d is actually used only in analysis of the approximation ratio.

3 MDS on Bounded Euler Genus Graphs

This section explains how asymptotic dimension helps us devise approximation algorithms for MDS on Euler-genus- g graphs. It also serves as an introduction to the method of cutting and motivates the later meta-theorems. The following algorithm for planar graphs will serve as a black-box building block for the more involved algorithm of Theorem 3.2.

Proposition 3.1. *There is a 205-approximation LOCAL algorithm A for MINIMUM DOMINATING SET in planar graphs with round complexity 6.*

To present the algorithm (cf. Algorithm 1) we first need to define the “best” minimum dominating sets of a labelled graph H . Among all dominating sets, we first discard those which contain some v for which there exists w with $N[v] \subsetneq N[w]$, and among the remaining ones, we declare “best” the one lexicographically smallest considering the labels of the vertices.

Algorithm 1 The algorithm for dominating set on planar graphs.

Require: A planar graph G .

Ensure: A dominating set D of G such that $|D| \leq 205 \cdot \text{MDS}(G)$.

- 1: $D \leftarrow \emptyset$
 - 2: Each vertex u computes $G[N^5[u]]$ and all minimum dominating sets of $N^4[u]$ in G . Among those sets, u selects the “best” one as D_u .
 - 3: Each vertex u picks the vertex v_u of smallest label in $D_u \cap N[u]$, and adds it to D .
-

By construction, Algorithm 1 outputs a dominating set D within 6 rounds. In fact, Algorithm 1 is very natural in the sense that in order to compute a global dominating set we use local dominating sets. The proof on the approximation factor is given in the Appendix C.1. We extend Algorithm 1 to genus g graphs by dealing separately with local subgraphs of G which are planar and which are not planar. For a graph G , the T -error set of G w.r.t. planarity is the vertex set defined by $X = \{u \in V(G) \mid G[N^T[u]] \text{ is not planar}\}$.

¹⁰It is conjectured in [BBE⁺24] that, even more generally, any graph class forbidding some graph as a *fat* minor should also have asymptotic dimension at most 2.

THEOREM 3.2. *There exists a function C such that for any g , there is a 616-approximation LOCAL algorithm for MINIMUM DOMINATING SET in Euler genus- g graphs with round complexity $C(g)$.*

PROOF. Suppose G has Euler genus g and let f be a 2-dimensional control function of the class of Euler genus- g graphs. A 2-hop in $X \subset V(G)$ is a sequence of vertices x_1, \dots, x_k in X such that consecutive vertices are at distance at most two in G .

Algorithm 2 The algorithm for dominating set on bounded genus graphs.

Require: A graph G of genus $\leq g$.

Ensure: A dominating set D of G such that $|D| \leq 616 \cdot \text{MDS}(G)$.

- 1: If $N^{f(10)+5}[u]$ is planar, u runs Algorithm 1 on the corresponding locally planar subgraph. Let X be the set of undominated vertices after step 1.
 - 2: If $u \in X$, u computes lexicographically smallest minimum set $D_{X'}$ in G which dominates the set $X' \subseteq X$ of vertices reachable from u with a 2-hop in X . It adds $N[u] \cap D_{X'}$ to D .
-

Claim 3.3. *The round complexity is some constant $C(g)$ depending only on g .*

PROOF OF CLAIM. It suffices to show that there is no 2-hop of weak diameter $\geq 2(f(10) + 6)g$ of vertices in X . Suppose such a 2-hop on vertices of X exists. Take two vertices x and y on the 2-hop which are at least $2(f(10) + 6)g$ apart. Note that for any $i < 2(f(10) + 6)g$ there exists a vertex on the 2-hop at distance i or $i + 1$ from x . Hence there exist $u_1, \dots, u_{g+1} \in X$ where u_i is at distance $2(f(10) + 6)(i - 1)$ or $2(f(10) + 6)(i - 1) + 1$ from x . Therefore, the sets $N^{f(10)+5}[u_1], \dots, N^{f(10)+5}[u_{g+1}]$ are pairwise disjoint and they are non-planar, so the genus of the graph is at least $g + 1$ by additivity (cf. [MT01, Theorem 4.4.3]). As the genus is at most g , this is a contradiction. \diamond

The coloring of the asymptotic dimension gives us color classes C_1, C_2, C_3 , such that:

- $C_i = B_1^i \cup \dots \cup B_{j_i}^i$ with¹¹ $\text{diam}_G(B_j^i) \leq f(10)$ for all $1 \leq j \leq j_i$, and
- for any $j \neq \ell$, we have $B_j^i \neq B_\ell^i$. More precisely, B_j^i and B_ℓ^i are at distance at least 11.

The Cutting. Let Y_G be a minimum dominating set of G . Let $G' = G[N[Y']]$ where $Y' = Y_G \cap N^5[B_j^i]$. Note that G' contains $N^4[B_j^i]$ and $|\text{MDS}(G')| \leq |Y_G \cap N^5[B_j^i]|$.

The Analysis Outside of the Error Set. Note that G' either contains only vertices from X or contains at least one vertex which is not in X and is therefore planar. Hence, vertices in $B_j^i \setminus X$ choose at most $205 \cdot |\text{MDS}(G')| \leq 205 \cdot |Y_G \cap N^5[B_j^i]|$ vertices as their neighbourhood of distance 5 in G and G' is the same. Vertices not in X choose at most $205|Y_G|$ vertices in each color class. Across all color classes, this adds up to at most $3 \cdot 205|Y_G|$.

The Analysis With Non-Empty Error Set. Let X_1, \dots, X_k be the 2-hop components. Note that for two distinct 2-hop components X_i and X_j , the neighbourhoods $N[X_i]$ and $N[X_j]$ do not intersect. As $|D_{X_i}| \leq |Y_G \cap N[X_i]|$, the vertices in X choose at most $|Y_G|$ vertices.

In total, at most $3 \cdot 205|Y_G| + |Y_G| = 616 \cdot |\text{MDS}(G)|$ vertices are added to D : the algorithm is a 616-approximation. \square

¹¹We denote by $\text{diam}_G(X)$ the largest distance in G between any two vertices of X .

In fact, if Algorithm 1 gives an approximation ratio α , then the above algorithm for genus g graphs is a $(3\alpha + 1)$ -approximation algorithm. One might now ask whether any algorithm for MDS on planar graphs can be turned into one for genus g graphs. We show that this is indeed possible. In order to do so, we need to show that the cutting step and its analysis work for any MDS algorithm for planar graphs. We delve into this deeper in the next section.

4 Meta-Theorems

In this section, we present our meta-theorems, which take as input a LOCAL approximation algorithm for a graph class \mathcal{C} and returns a LOCAL approximation algorithm for a broader class \mathcal{D} that locally looks like graphs in \mathcal{C} , possibly with some local exceptions.

Meta-theorems using black-box LOCAL algorithms might be difficult to state as there are some subtle issues with vertex identifiers (see for instance [FKP13, FGKS13, GHS13]). The LOCAL model assumes the input graph G has unique $O(\log n)$ -bit vertex labels, i.e., with identifiers polynomial in G . In particular, running a LOCAL algorithm A on only a small part H of G might fail because: (1) vertex labels in H may not be polynomial in H anymore, and (2) algorithm A may require polynomial-size identifiers to work properly. Thus, Meta-Theorem I assumes that the black-box algorithm A does not require polynomial identifiers. This captures all order-invariant algorithms¹². In contrast, Meta-Theorem II and Meta-Theorem III allow A to use polynomial identifiers, but impose restrictions on the optimization problem. Due to space constraints, Meta-Theorem II & III are presented in Section A.

From Locally- \mathcal{C} to Locally Nice Graphs. Given a graph class \mathcal{C} , we say that a graph class \mathcal{D} is T -locally- \mathcal{C} if for all $G \in \mathcal{D}$ and $u \in V(G)$, $G[N^T[u]] \in \mathcal{C}$.

This definition can already capture many graph classes, such as graphs of girth at least $2k + 2$: those are just graphs that are k -locally- \mathcal{T} , where \mathcal{T} is the class of trees. Similarly, graphs embedded on a surface with edge-width¹³ at least $2k + 2$ are k -locally- \mathcal{P} , where \mathcal{P} is the class of planar graphs. However, our meta-theorem captures broader classes of graphs as it can support local exceptions.

Given a graph class \mathcal{D} and $G \in \mathcal{D}$, the T -error set of G w.r.t. \mathcal{C} is the vertex set defined by $X = \{u \in V(G) \mid G[N^T[u]] \notin \mathcal{C}\}$. In other words, this is the set of vertices that are witnesses for G not being T -locally- \mathcal{C} . Clearly, if G has no T -errors, that is $X = \emptyset$, then G is T -locally- \mathcal{C} . Given a ρ -local¹⁴ problem Π , the class \mathcal{D} is T -locally δ -nice w.r.t. \mathcal{C} and Π , if the maximum weak diameter in any $G \in \mathcal{D}$ of any connected component of $G[N^{2\rho}[X]]$ is at most δ , where X is the T -errors of G w.r.t. \mathcal{C} . This means that the set of T -errors can be covered by far-away bounded-radius balls of G . By Proposition C.11, bounded Euler genus graphs form a locally nice class of graphs w.r.t. \mathcal{P} .

Uniform Approximation. Our meta-theorems are inspired by the intriguing property [BGPW25, Proposition 3.1] that any approximate algorithm A for MDS is going to perform well in a target graph class \mathcal{D} that is locally- \mathcal{C} , as long as \mathcal{D} has small asymptotic dimension, and A makes "good local choices" in \mathcal{C} . More precisely, a LOCAL algorithm A is a k -uniform α -approximation in \mathcal{C} , if for all $G \in \mathcal{C}$ and $S \subseteq V(G)$, $|A(G) \cap S| \leq \alpha \cdot \text{MDS}(G, N^k[S])$, where $A(G)$ is the set returned by A when run on G , and $\text{MDS}(G, X)$ is the minimum size of a subset of $V(G)$ dominating vertices of X in G . Unless explicitly said, we will assume for convenience that $k \geq 1$, as otherwise only very

¹²A model where vertices have unique labels, but the output of A is invariant under relabelings preserving the label order [GHS13].

¹³The length of shortest noncontractible cycle in the embedded graph.

¹⁴A problem whose solution can be verified by checking all neighborhoods of radius ρ . See next page for a formal definition.

restricted graph classes \mathcal{C} can support 0-uniform constant-approximation algorithms for MDS (cf. Proposition D.1).

We can now state properly the fact that any k -uniform approximate algorithm that performs well for \mathcal{C} , performs well for a locally- \mathcal{C} class of bounded asymptotic dimension. It can be shown that the bounded asymptotic dimension assumption cannot be removed, even for MDS¹⁵.

Proposition 4.1 ([BGPW25, Proposition 3.1]). *Let A be a k -uniform α -approximation LOCAL algorithm for MINIMUM DOMINATING SET in a hereditary¹⁶ class of graphs \mathcal{C} with round complexity $r \geq 1$. Let \mathcal{D} be a graph class with d -dimensional control function f , and that is¹⁷ $(f(2k+2) + r)$ -locally- \mathcal{C} . Then A is also an $\alpha(d+1)$ -approximation algorithm on \mathcal{D} .*

We extend Proposition 4.1 in multiple ways. First, we allow \mathcal{D} to be locally nice, and secondly, we prove that this also applies to problems other than MDS. This means that, for example, any approximate algorithm A solving a well-behaved problem that performs well on planar graphs can be converted into an algorithm B that performs well on bounded genus graphs, as long as A is k -uniform on \mathcal{C} , and \mathcal{D} has small asymptotic dimension. Before giving the full statement of Meta-Theorem I, we define the class of problems we are interested in.

LOCAL Optimization Problems. As defined in [GHS13], a *simple graph problem* Π is an optimization problem in which a feasible solution¹⁸ is a subset of vertices or edges, and the goal is to optimize the size of a feasible solution. A simple graph problem Π is ρ -local if there is a LOCAL algorithm R with round complexity ρ that can recognize a feasible solution¹⁹. Algorithm R recognizes S as a feasible solution if and only if, whenever applied on G with S given, R returns true for all vertices of G . See Section B.1 for more formal definitions. We are mainly interested in small locality problems, that is when ρ is constant, as clearly all simple graph problems are ρ -local for ρ large enough. A lot of well-known problems such as optimization variants of LCL problems [NS95]. This includes MINIMUM VERTEX COVER, MINIMUM DOMINATING SET, MAXIMUM INDEPENDENT SET, MAXIMAL MATCHING, and MINIMUM k -TUPLE DOMINATING SET. The distance- d versions of these problems²⁰ are also d -local problems. However, MINIMUM CONNECTED DOMINATING SET is $\Omega(n)$ -local²¹. Our framework includes problems such as RED-BLUE DOMINATING SET that can take vertex or edge input labels on the graph, and problems defined on arbitrarily large degree graphs.

In this work, we focus on subsets of vertices, but the techniques presented here extend to subsets of edges. In our setting, we restrict our attention to cases where the vertices do not have access to n , the size of the graph. For some problems, such as MINIMUM DOMINATING SET, this restriction can

¹⁵This can be achieved by observing that a simple constant-round 3-approximation LOCAL algorithm exists for MDS in trees, but (high) girth- g graphs do not admit a constant-round $o(g)$ -factor approximation. The latter follows by combining Ramsey-theoretic arguments with the lower bound proof due to [GHS13] based on $2k$ -regular graphs of girth g , where vertices can be linearly ordered so that a positive fraction of them have pairwise isomorphic radius- r neighbourhoods, for any r, k, g .

¹⁶A class of graphs closed under vertex deletion.

¹⁷The original statement claimed erroneously that \mathcal{D} being t -locally- \mathcal{C} for $t = f(2k+1)$ is enough, instead of the correct $t = f(2k+2) + r$. This has no impact, as in practice we use Proposition 4.1 in Corollary 4.4 only for $k \leq r = O(1)$.

¹⁸Here, a feasible solution is any subset of vertices or edges that satisfies the constraints of the graph problem. Said differently, an admissible candidate solution.

¹⁹We do not require that the algorithm checks the optimality of the solution. If we require this, most interesting problems would require $\rho = \Omega(n)$ to be ρ -local.

²⁰E.g. for MINIMUM DOMINATING SET, every vertex has a dominating vertex in its distance- d neighborhood.

²¹In a n -vertex cycle, no LOCAL algorithm R of round complexity $< n/2$ can distinguish the situation where a feasible solution is $S = V(G) \setminus \{u\}$ from the situation where $S = V(G) \setminus \{u, v\}$ with two opposite vertices u, v .

be relaxed²². However, we omit this case for brevity, since we are not aware of any constant-round algorithm for a ρ -local problem that requires n . As, in the LOCAL model with constant number of rounds, the literature mostly studies minimization problems, we also focus on those. Observe that many classically interesting problems like MAXIMUM INDEPENDENT SET and MAXIMUM MATCHING do not admit constant-time constant-ratio approximations (see [BH25] for references), even on the restricted graph class of paths²³.

Given a ρ -local problem Π , a graph G and $X \subseteq V(G)$, we define $\text{OPT}_\Pi(G, X)$ as the optimal size of a subset of vertices of G that is feasible on X (and not necessarily for vertices outside X). We naturally extend the definition of k -uniform algorithm as follows: a LOCAL algorithm A is a k -uniform α -approximation for Π in a class of graphs \mathcal{C} , if for all $G \in \mathcal{C}$ and $S \subseteq V(G)$, $|A(G) \cap S| \leq \alpha \cdot \text{OPT}_\Pi(G, N^k[S])$.

We can now state our strengthening of Proposition 4.1. Note that we further require (when the set of errors is not empty) that our problem is *additive*, that is, any superset of a solution feasible on X is still feasible on X . For instance, MDS is additive because adding vertices to any dominating set maintains a dominating set.

THEOREM 4.2 (META-THEOREM I). *Let A be a k -uniform α -approximation LOCAL algorithm that does not require polynomial identifiers, for an additive ρ -local problem Π in a class of graphs \mathcal{C} with round complexity r , where \mathcal{C} is hereditary and stable by disjoint union. Let \mathcal{D} be a graph class of asymptotic dimension d with d -dimensional control function f that is T -locally δ -nice w.r.t. \mathcal{C} and Π , where $T = f(2k + 2\rho) + \max\{k + \rho, r\}$. Then, there exists a $\max\{k, \rho\}$ -uniform $(\alpha(d + 1) + 1)$ -approximation algorithm B on \mathcal{D} with round complexity $T + \delta + 2$. Furthermore, if \mathcal{D} is T -locally \mathcal{C} , the approximation ratio is $\alpha(d + 1)$ and the additivity condition can be dropped.*

Cutability. It may be difficult to prove that some algorithm is k -uniform. Instead, we provide a simple sufficient criterion for deciding if *any* constant-round constant-ratio approximation algorithm for Π is k -uniform.

A ρ -local problem Π is *cuttable* for a hereditary graph class \mathcal{C} with parameter $\beta > 0$ if for all $G \in \mathcal{C}$ and $X \subseteq V(G)$, there exists some $Y \supseteq X$ such that $\text{OPT}_\Pi(G[Y]) \leq \beta \cdot \text{OPT}_\Pi(G, X)$. Informally, this asks that X can be augmented into a set Y (if not the whole vertex set) so that the graph induced by this bigger set Y has an optimal solution somewhat smaller than the optimal feasible solution for X .

For instance, MDS is cuttable with parameter $\beta = 1$ (see Claim C.3), as is the distance- d version of MDS (see Proposition C.6). We prove in Section C.6 the fundamental property that every constant-round constant-ratio approximation for a cuttable problem is also k -uniform (as long as the algorithm does not require polynomial identifiers). Therefore, MDS admits uniform approximation algorithms. This means we can actually use Meta-Theorem I for any cuttable problem without the k -uniform assumption. Actually, for MDS, we can show that this property holds even if the algorithm requires polynomial identifiers (cf. Proposition C.2).

²²In full generality, this applies to all problems for which, for every n , there exists an n -vertex graph within the input graph class that has bounded diameter and admits a constant-size solution. For MDS on planar graphs, we can take a depth-1 tree.

²³The rough intuition for this is that on paths, any constant-time algorithm run on a vertex far away from the end of the path sees the same neighborhood, and therefore will output the same result on every vertex or edge. The only way to comply with the problem description, unlike minimization problem like MDS, is to not take any vertex in the independent set (or any edge in the maximum matching), at least in the middle of the path. This leads to an arbitrarily bad approximation ratio.

Recall that Euler genus- g graphs have asymptotic dimension at most 2, and are locally nice (cf. Proposition C.11). Thus, Meta-Theorem I implies that any constant-round α -approximation for MDS gives a $C(g)$ -round $(3\alpha + 1)$ -approximation for graphs of Euler genus g .

The best known ratio for MDS in planar graphs is $\alpha = 11 + \varepsilon$ due to Heydt et al. [HKOdM⁺25]. Therefore,

Corollary 4.3. *For any $\varepsilon > 0$, for any g , there is a $k(\varepsilon)$ -uniform $(34 + \varepsilon)$ -approximation LOCAL algorithm for MINIMUM DOMINATING SET in Euler genus- g graphs with round complexity $C(\varepsilon, g)$ for some functions C and k .*

We can even go one step further by applying Meta-Theorem I to the uniform approximation LOCAL algorithm derived from Corollary 4.3, and the class \mathcal{C} of Euler genus- g graphs. We immediately obtain the following, where 34.01 can be thought of as $34 + \varepsilon$.

Corollary 4.4. *Let \mathcal{C} be the class of Euler genus- g plus graphs, and \mathcal{D} be a graph class of asymptotic dimension d that is r -locally- \mathcal{C} for some r large enough. Then, there is a $34.01(d + 1)$ -approximation LOCAL algorithm for MINIMUM DOMINATING SET in \mathcal{D} with round complexity $C(d, g, r)$ for some function C .*

Actually, Corollaries 4.3 and 4.4 hold for the larger class of Euler genus- g plus a -apex graphs, i.e. graphs obtained from a graph of Euler genus g where at most a vertices are allowed to have an arbitrary neighborhood. These graphs exclude $K_{3+a, 2g+3+a}$ as minor (as, without apices, the graphs exclude $K_{3, 2g+3}$), and thus have asymptotic dimension 2 as well. They are also locally nice, since a apices cannot create short-cuts longer than $2a$. So, the approximation ratio is preserved, whereas the round complexity depends on a .

Moreover, our results can be applied to get approximation algorithms on graphs of bounded Euler genus for *all* other minimization problems in the Distributed Computing literature that admit a constant-round, constant-ratio approximation on planar graphs (for a comprehensive list of such problems, see the survey of Feuilloley [Feu23]).

For instance, let us focus on the MINIMUM k -TUPLE DOMINATING SET problem (studied extensively in [WAF02]), a generalization of MINIMUM DOMINATING SET. In this variant, the goal is to find the smallest subset of vertices such that every vertex outside the set is adjacent to at least k vertices within it. The classical MDS problem corresponds to the special case $k = 1$. For planar graphs, this problem admits a constant-round $k/(k - 2)$ -approximation algorithm [CHS⁺14] for $k > 2$, and a 6-approximation [CHS⁺17] for $k = 2$. Both algorithms do not require polynomial identifiers.

Using the fact that MINIMUM k -TUPLE DOMINATING SET is cuttable and additive (cf. Proposition C.5), that Euler genus- g graphs are locally nice (cf. Proposition C.11), and by applying Proposition C.7 and Meta-Theorem I, we get the following:

Corollary 4.5. *For every $k \geq 2$, there is a $C_1(k)$ -uniform α -approximation LOCAL algorithm for MINIMUM k -TUPLE DOMINATING SET in Euler genus- g graphs with round complexity $C_2(g, k)$, for some functions C_1, C_2 , where $\alpha = 19$ if $k = 2$ and $\alpha \leq 3k/(k - 2) + 1$ otherwise.*

For MINIMUM CONNECTED DOMINATING SET [WAF02], which is an example of non-local problem, we can obtain a constant-round approximation with ratio $O(\sqrt{g})$ in Euler genus- g graphs by applying the "connected" reduction from MDS in bounded expansion graphs due to [AAOdMRS18, Theorem 5.8]. This gives an approximation ratio of $2\alpha \nabla_1 \leq 68\sqrt{3g/2} + O(1)$, where $\alpha = 34 + \varepsilon$ is the best approximation ratio for MDS in Euler genus- g graphs (Corollary 4.3), and $\nabla_1 \leq \sqrt{3g/2} + 3$

(Proposition D.3) is the expansion of Euler genus- g graphs. We leave open the question of improving the ratio to a constant independent of the genus.

Acknowledgments

The first, third and fourth authors have been partially supported by the French ANR projects ENEDISC (ANR-24-CE48-7768) and TEMPOGRAL (ANR-22-CE48-0001). The second author was supported in part by the Research Council of Finland, Grants 363558 and 359104. The sixth author was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID:390685689).

References

- [AAOdMRS18] S. AKHOONDIAN AMIRI, P. OSSONA DE MENDEZ, R. RABINOVICH, AND S. SIEBERTZ, *Distributed domination on graph classes of bounded expansion*, in 30th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA), ACM Press, July 2018, pp. 143–151. doi: 10.1145/3210377.3210383.
- [ABY63] L. AUSLANDER, T. A. BROWN, AND J. W. YOUNGS, *The imbedding of graphs in manifolds*, Journal of Mathematics and Mechanics, 12 (1963), pp. 629–634. <http://www.jstor.org/stable/24900840>.
- [AGLP89] B. AWERBUCH, A. V. GOLDBERG, M. LUBY, AND S. A. PLOTKIN, *Network decomposition and locality in distributed computation*, in 30th Annual IEEE Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society Press, May 1989, pp. 364–369. doi: 10.1109/SFCS.1989.63504.
- [ASS19] S. A. AMIRI, S. SCHMID, AND S. SIEBERTZ, *Distributed dominating set approximations beyond planar graphs*, ACM Transactions on Algorithms, 15 (2019), pp. Article No. 39, pp. 1–18. doi: 10.1145/3093239.
- [BBE⁺24] M. BONAMY, N. BOUSQUET, L. ESPERET, C. GROENLAND, C.-H. LIU, F. PIROT, AND A. SCOTT, *Asymptotic dimension of minor-closed families and Assouad–Nagata dimension of surfaces*, Journal of the European Mathematical Society, 26 (2024), pp. 3739–3791. doi: 10.4171/JEMS/1341.
- [BBK⁺26] A. BALLIU, S. BRANDT, F. KUHN, D. OLIVETTI, T. PICALET, AND G. SCHMID, *The distributed complexity landscape on trees depends on the knowledge about the network size*, 2026. <https://arxiv.org/abs/2605.12787>.
- [BGPW25] M. BONAMY, C. GAVOILLE, T. PICALET, AND A. WESOLEK, *Local constant approximation for dominating set on graphs excluding large minors*, in 44th Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, June 2025. doi: 10.1145/3732772.3733531.
- [BH25] R. B. BOPPANA AND M. M. HALLDÓRSSON, *Approximating independent sets in constant distributed rounds*, in 32nd International Colloquium on Structural Information & Communication Complexity (SIROCCO), U. Schmid and R. Kuznets, eds., vol. 15671 of Lecture Notes in Computer Science, Springer, June 2025, pp. 144–158. doi: 10.1007/978-3-031-91736-3_9.
- [CHS⁺14] A. CZYGRINOW, M. HANĆKOWIAK, E. SZYMAŃSKA, W. WAWRZYŃIAK, AND M. WITKOWSKI, *Distributed local approximation of the minimum k -tuple dominating set in planar graphs*, in International Conference on Principles of Distributed Systems, Springer, 2014, pp. 49–59.
- [CHS⁺17] A. CZYGRINOW, M. HANĆKOWIAK, E. SZYMAŃSKA, W. WAWRZYŃIAK, AND M. WITKOWSKI, *Improved distributed local approximation algorithm for minimum 2-dominating set in planar graphs*, Theoretical Computer Science, 662 (2017), pp. 1–8. doi: <https://doi.org/10.1016/j.tcs.2016.12.001>, <https://www.sciencedirect.com/science/article/pii/S0304397516307071>.
- [CHW08] A. CZYGRINOW, M. HANĆKOWIAK, AND W. WAWRZYŃIAK, *Fast distributed approximations in planar graphs*, in 22nd International Symposium on Distributed Computing (DISC), vol. 5218 of Lecture Notes in Computer Science, Springer, September 2008, pp. 78–92. doi: 10.1007/978-3-540-87779-0_6.
- [CHW18] A. CZYGRINOW, M. HANĆKOWIAK, AND W. WAWRZYŃIAK, *Distributed approximation algorithms for the minimum dominating set in K_h -minor-free graphs*, in 29th International Symposium on Algorithms and Computation (ISAAC), vol. 123 of LIPIcs, December 2018, pp. 22:1–22:12. doi: 10.4230/LIPIcs.ISAAC.2018.22.
- [CHWW19] A. CZYGRINOW, M. HANĆKOWIAK, W. WAWRZYŃIAK, AND M. WITKOWSKI, *Distributed $CONGEST_{BC}$ constant approximation of MDS in bounded genus graphs*, Theoretical Computer Science, 757 (2019), pp. 1–10. doi: 10.1016/j.tcs.2018.07.008.
- [CL21] C. COUPETTE AND C. LENZEN, *A breezing proof of the KMW bound*, in Symposium on Simplicity in Algorithms (SOSA), ACM-SIAM, January 2021, pp. 184–195. doi: 10.1137/1.9781611976496.21.
- [Feu23] L. FEUILLOLEY, *Bibliography of distributed approximation beyond bounded degree*, Tech. Rep. 2001.08510v4 [cs.CG], arXiv, November 2023.
- [FGKS13] P. FRAIGNIAUD, M. GÖÖS, A. KORMAN, AND J. SUOMELA, *What can be decided locally without identifiers?*, in 32nd Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, July 2013, pp. 157–165. doi: 10.1145/2484239.2484264.

- [FKP13] P. FRAIGNIAUD, A. KORMAN, AND D. PELEG, *Towards a complexity theory for local distributed computing*, Journal of the ACM, 60 (2013), pp. Article No. 35, pp. 1–26. doi: 10.1145/2499228.
- [GHS13] M. GÖÖS, J. HIRVONEN, AND J. SUOMELA, *Lower bounds for local approximation*, Journal of the ACM, 60 (2013), pp. Article No. 39, pp. 1–23. doi: 10.1145/2528405.
- [GJ79] M. R. GAREY AND D. S. JOHNSON, *Computers and Intractability - A Guide to the Theory of NP-Completeness*, W.H. Freeman, 1979.
- [GKM17] M. GHAFFARI, F. KUHN, AND Y. MAUS, *On the complexity of local distributed graph problems*, in 49th Annual ACM Symposium on Theory of Computing (STOC), ACM Press, June 2017, pp. 784–797. doi: 10.1145/3055399.3055471.
- [Gro93] M. GROMOV, *Geometric Group Theory: Asymptotic invariants of infinite groups*, Cambridge University Press, 1993.
- [HKOdM⁺25] O. HEYDT, S. KUBLENZ, P. OSSONA DE MENDEZ, S. SIEBERTZ, AND A. VIGNY, *Distributed domination on sparse graph classes*, European Journal of Combinatorics, 123 (2025), p. 103773. doi: 10.1016/j.ejc.2023.103773.
- [HLS14] M. HILKE, C. LENZEN, AND J. SUOMELA, *Brief announcement: local approximability of minimum dominating set on planar graphs*, in 33rd Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, July 2014, pp. 344–346. doi: 10.1145/2611462.2611504.
- [KMW16] F. KUHN, T. MOSCIBRODA, AND R. WATTENHOFER, *Local computation: Lower and upper bounds*, Journal of the ACM, 63 (2016), pp. Article No. 17, pp. 1–44. doi: 10.1145/2742012.
- [KSV21] S. KUBLENZ, S. SIEBERTZ, AND A. VIGNY, *Constant round distributed domination on graph classes with bounded expansion*, in 28th International Colloquium on Structural Information & Communication Complexity (SIROCCO), T. Jurdziński and S. Schmid, eds., vol. 12810 of Lecture Notes in Computer Science, Springer, Cham, June 2021, pp. 334–351. doi: 10.1007/978-3-030-79527-6_19.
- [KW05] F. KUHN AND R. WATTENHOFER, *Constant-time distributed dominating set approximation*, Distributed Computing, 17 (2005), pp. 303–310. doi: 10.1007/s00446-004-0112-5.
- [KYK80] T. KIKUNO, N. YOSHIDA, AND Y. KAKUDA, *The NP-completeness of the dominating set problem in cubic planar graphs*, IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences, E63-E (1980), pp. 443–444.
- [Len11] C. LENZEN, *Synchronization and symmetry breaking in distributed systems*, PhD thesis, ETH Zurich, Switzerland, December 2011.
- [LOW08] C. LENZEN, Y. A. OSWALD, AND R. WATTENHOFER, *What can be approximated locally? – Case study: Dominating sets in planar graphs*, in 20th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA), ACM Press, June 2008, pp. 46–54. doi: 10.1145/1378533.1378540.
- [LPW13] C. LENZEN, Y.-A. PIGNOLET, AND R. WATTENHOFER, *Distributed minimum dominating set approximations in restricted families of graphs*, Distributed Computing, 26 (2013), pp. 119–137. doi: 10.1007/s00446-013-0186-z.
- [LW08] C. LENZEN AND R. WATTENHOFER, *Leveraging Linial’s locality limit*, in 22nd International Symposium on Distributed Computing (DISC), vol. 5218 of Lecture Notes in Computer Science, Springer, September 2008, pp. 394–407. doi: 10.1007/978-3-540-87779-0_27.
- [MT01] B. MOHAR AND C. THOMASSEN, *Graphs on Surfaces*, The Johns Hopkins University Press, 2001.
- [NS95] M. NAOR AND L. STOCKMEYER, *What can be computed locally*, SIAM Journal on Computing, 24 (1995), pp. 1259–1277. doi: 10.1137/S0097539793254571.
- [RG20] V. ROZHOŇ AND M. GHAFFARI, *Polylogarithmic-time deterministic network decomposition and distributed derandomization*, in 52nd Annual ACM Symposium on Theory of Computing (STOC), ACM Press, June 2020, pp. 350–363. doi: 10.1145/3357713.3384298.
- [Suo13] J. SUOMELA, *Survey of local algorithms*, ACM Computing Surveys, 45 (2013), pp. Article No. 24, pp. 1–40. doi: 10.1145/2431211.2431223.
- [vRWZ09] P. VON RICKENBACH, R. WATTENHOFER, AND A. ZOLLINGER, *Algorithmic models of interference in wireless ad hoc and sensor networks*, IEEE/ACM Transactions on Networking, 17 (2009), pp. 172–185. doi: 10.1109/TNET.2008.926506.
- [WAF02] P.-J. WAN, K. M. ALZOUBI, AND O. FRIEDER, *Distributed construction of connected dominating set in wireless ad hoc networks*, in 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM), vol. 3, June 2002, pp. 1597–1604. doi: 10.1109/INFOCOM.2002.1019411.
- [WY07] P.-J. WAN AND C.-W. YI, *On the longest edge of gabriel graphs in wireless ad hoc networks*, IEEE Transactions on Parallel and Distributed Systems, 18 (2007), pp. 111–125. doi: 10.1109/TPDS.2007.253285.

A When the black-box algorithm requires polynomial identifiers

As we have seen, Proposition C.7 applies to lots of optimization problems. However, it does not fulfil the polynomial identifiers requirement of the LOCAL model. We would like to have the same statement to apply in the LOCAL model with polynomial identifiers with the same assumptions. However, assigning new identifiers to vertices while maintaining control of the output of the algorithm is very challenging. Ramsey-type techniques (for example, for finding lower bounds on the approximation ratio on particular graph classes [HLS14]) can be useful for this, but only on very particular graphs: ones where each neighborhood is similar. In general graphs, to the best of our knowledge, we do not know how to adapt this technique. Thankfully, for locally- \mathcal{C} classes, where the set of errors is empty, we can weaken the statement of Proposition C.7 in a way that we can still apply a variation of Meta-Theorem I afterwards. The trick is to modify the coloring given by the asymptotic dimension so that all color classes are almost balanced in size (see Lemma C.12). With this technique, we just need to apply uniformity on sets that are large. Formally, LOCAL algorithm A is an ε -weakly k -uniform α -approximation for Π in a class of graphs \mathcal{C} , if for all $G \in \mathcal{C}$ and $S \subseteq V(G)$ such that $|S| \geq \lfloor \varepsilon |V(G)| \rfloor$, we have $|A(G) \cap S| \leq \alpha \cdot \text{OPT}_\Pi(G, N^k[S])$. We say that the d -local problem Π is ε -weakly uniformizable on \mathcal{C} if for every integer $d > 0$ and r and any real α , there exists some integers k, r' and real β such that for every r -round α -approximation A of Π on \mathcal{C} in constant time, there exists some r' -round LOCAL algorithm B that is a ε -weakly k -uniform β -approximation of Π on \mathcal{C} in constant time. The function $(r, \alpha) \mapsto (k, \beta)$ is called the (weak) *binding function*.

It is a corollary of the proof of Proposition C.7 that any cuttable problem is weakly uniformizable, in the LOCAL model this time (see Corollary C.10). Therefore, we can use Lemma C.12 to only work on color classes of linear size (which preserves the polynomial-ness of identifiers), and prove the following.

THEOREM A.1 (META-THEOREM II). *Let d, k be integers $\alpha > 0$ a real, A be a $\frac{1}{d+1}$ -weakly k -uniform α -approximation LOCAL algorithm for a ρ -local problem Π in a class of graphs \mathcal{C} with round complexity r , where \mathcal{C} is hereditary and stable by disjoint union. Let \mathcal{D} be a graph class of asymptotic dimension d with d -dimensional control function f , that is T -locally- \mathcal{C} for $T = f(6k+6\rho) + 4k + 4\rho + \max\{k + \rho, r\}$. Then A is a k -uniform $\alpha(d+1)$ -approximation algorithm on \mathcal{D} .*

The proof of this can be found in Section C.11.

When the error set is not empty, we require further assumptions. For a discussion on why those assumptions are required, we refer to Section B.4 A problem Π is called ω -paddable if for any integer N , there exists some $G \in \mathcal{C}$ on at least N vertices such that $\text{OPT}_\Pi(G) \leq \omega$. We also add a small property that the problem needs to satisfy: a d -local problem Π is *dense* in \mathcal{C} if there exists some unbounded function g that for any $G \in \mathcal{C}$ of diameter D , we have $\text{OPT}_\Pi(G) \geq \beta \cdot D$. We require that²⁴ $\liminf f = +\infty$. This essentially mean that there are no large (in diameter) portion of the graph with no vertices of the solution; every vertex is close to a vertex of the solution. Under the both additional assumptions of paddability and density, we can prove an analogue of Meta-Theorem II. See Meta-Theorem III for the full statement.

Summary Table of the Results. Here is a simplified table of the assumptions necessary to apply our meta-theorem that transform a constant-round constant-factor approximation algorithm A on the class \mathcal{C} to another constant-round constant-factor approximation algorithm on the class \mathcal{D} .

²⁴It is possible to get rid of this property if one adds that the graphs in the definition of ω -paddable also have bounded diameter. Then, the rough idea is that, in the proof of Meta-Theorem III, the algorithm will be able to compute the optimal solution on the graph \widehat{G} of the proof. Then one gets Eq. (6) without the ω term.

We assume that \mathcal{C} is hereditary and stable by disjoint union, and that \mathcal{D} is of bounded asymptotic dimension.

	T -error set is empty (\mathcal{D} is locally- \mathcal{C})	T -error set is not empty (\mathcal{D} is locally nice)
<i>A does not assume polynomial IDs</i>	Π is cuttable (Meta-Theorem I and Proposition C.7)	Π is cuttable and additive (Meta-Theorem I and Proposition C.7)
<i>A assumes polynomial IDs</i>	Π is cuttable (Meta-Theorem II and Corollary C.10)	Π is cuttable, additive, paddable and dense ²⁵ (Meta-Theorem III and Corollary C.10)

Table 1. Simplified table of the assumptions necessary to apply Meta-Theorem I, Meta-Theorem II and Meta-Theorem III, along with their references.

B Discussions

B.1 Definitions for Locally Verifiable Problems

Let us define ρ -local verifiable problems in more detail. Let G be a graph with a set Σ_{in} of input labels on vertices or edges, which includes a unique identifier for each vertex. If $v \in V(G)$, A is a LOCAL algorithm, and X is a set of vertices of G , we use $A(G, X, v)$ to represent the output of vertex v when A is run on G with input labels $\Sigma_{\text{in}} \sqcup X$ (the local inputs form an encoding of X , i.e. every vertex knows if it is contained or not in X). A simple graph problem Π is ρ -local if there exists some LOCAL algorithm R which runs in ρ rounds such that for all subsets of vertices X , the following holds:

- if X is a feasible solution of the problem Π on (G, Σ_{in}) , then for all vertices v of G , we have $A(G, X, v) = 1$, or
- if X is not a feasible solution of the problem Π on (G, Σ_{in}) , then there exists $v \in V(G)$ such that $A(G, X, v) = 0$.

B.2 Why cutability seems unavoidable

Cutability is not a trivial property. Consider the following very artificial problems: we want to output a set (as small as possible) which contains every vertex of degree 0 or 1 in the graph. Of course, this problem admits a trivial 1-round algorithm: take all vertices of degree 0 or 1. However, our goal is to be able to use algorithms in a black-box manner. Instead, we consider the 1-round algorithm A that outputs every vertex of degree not equal to 2. Take \mathcal{C} the class of forests, on which A is a 2-approximation. This is because there are at least as many leaves as vertices of degree 3 or more.

Now, let \mathcal{D} be the closure (under subgraphs) of q -subdivisions of ladders²⁶, where q is as large as we like and a ladder is the Cartesian product of an edge and an arbitrarily long path. Clearly,

²⁵It is possible to remove the dense condition, at the expense of making the paddable condition stricter. See Section A.

²⁶The q -subdivision of a graph G is a copy of G where every edge is replaced by a path with q inner vertices.

every graph in \mathcal{D} is q -locally in \mathcal{E} . Additionally \mathcal{D} has asymptotic dimension 1, (one can check that they are K_4 -minor-free, and so have bounded treewidth hence asymptotic dimension 1 [BBE⁺24]). However, A fails catastrophically on \mathcal{D} . Indeed, there are graphs in \mathcal{D} that do not have any leaves or isolated vertices (one can just take the q -subdivision of a ladder), so that the best solution on this graph is empty. Of course, in this very specific scenario there does exist an alternative algorithm which gives a 1-approx in forests and actually works in \mathcal{D} . However, the point is to avoid modifying the black-box A . More generally, such degree-based problems seem to have uncontrollable behavior across local-global boundaries: in \mathcal{E} the algorithm can afford to take more vertices than necessary because of the existence of some other vertices (such as the large amount of leaves in the case of forests), whereas in \mathcal{D} these vertices do not exist.

B.3 Cutability of DISTANCE-EXACTLY- d DOMINATING SET

In some graph classes such as planar graphs and minor-free classes, Proposition C.6 can be extended to MINIMUM DISTANCE-EXACTLY- d DOMINATING SET. Given a graph G and d a positive integer, a subset of vertices D is called a *distance-exactly- d dominating set* of G if every vertex v of $V(G) \setminus D$ is at distance exactly d from some vertex in D . In the MINIMUM DISTANCE-EXACTLY- d DOMINATING SET problem, one asks for *distance-exactly- d dominating set* of minimum size. In graph classes that are closed under the addition of leaves to the graphs, DISTANCE- $(\leq d)$ DOMINATING SET is cuttable with parameter d . To give a rough intuition, in these classes, one can attach a path of length $d - 1$ to a vertex v , put v and all vertices in the path to the dominating set, and thus simulate a distance- $(\leq d)$ dominating set that way.

B.4 On the additional assumptions of the Meta-Theorem for paddable problems

In the statement of Meta-Theorem III, we require problems that are paddable, i.e., we want to exclude graph classes that admit arbitrarily large graphs whose difficulty is concentrated in a bounded-diameter region with bounded optimum. Otherwise, a padding trick can inflate instance size (hence identifier ranges) without increasing hardness. One cannot hope for a version of Meta-Theorem III with non-empty error sets and for problems without this property. Here is the scenario we want to avoid. One could take any approximation algorithm on a class of graph \mathcal{E} , and construct a class \mathcal{D} where graphs of size n are formed by taking a graph of \mathcal{E} of size at most \sqrt{n} (or n^c for some $0 < c < 1$), and attaching somewhere an error set X of size at least $n - \sqrt{n}$, which has small radius and small (bounded) optimum solution. Then, this version of Meta-Theorem III would give a good approximation on this class \mathcal{D} . However, solving the problem Π efficiently on \mathcal{D} would require solving it on \mathcal{E} , as the error set is of small radius and has a small optimum solution. Put in a nutshell, there is roughly a reduction from the class \mathcal{D} to the class \mathcal{E} . However, as the error set is large in size, one could attribute very large identifiers to the portion of the graph in \mathcal{E} : setting $|X| \geq n - \sqrt{n}$ can artificially inflate the identifiers by a square (and one can get a polynomial with power C by choosing $|X| \geq n^{1/C}$). If the algorithm does not need the polynomial identifiers, everything still works. If it is not the case however, we have transformed an easy instance to another easy instance where identifiers are larger. This is not possible if the algorithm truly needs the polynomial identifiers, i.e., if the approximation guarantee depends on the vertex identifiers.²⁷

²⁷For these reasons, this might hint towards considering the size of the largest identifier as complexity parameter of its own, just like the size of the graph [BBK⁺26].

Thus, we require \mathcal{C} to be paddable or \mathcal{D} to be not paddable. We prove the result for the case where \mathcal{C} is paddable only, as the proof is similar for the other case.²⁸

C Proofs

C.1 Approximation for planar graphs

Proposition 3.1. *There is a 205-approximation LOCAL algorithm A for MINIMUM DOMINATING SET in planar graphs with round complexity 6.*




Fig. 2. Depicted in green is the dominating set D_u for $u = 2$. Note that $d(u)$ is 1 or 5.

PROOF. Let G be a planar graph, let Y be a nice dominating set of G and let D be the dominating set of G returned by the algorithm. Let $u \in V(G)$ and let D_u be the dominating set computed by u . That means

- D_u is nice
- D_u is lexicographically smallest nice set dominating $N^4[u]$

The first key observation is that given disjoint subsets A, B, C of $V(G)$ such that $G = G[A \cup C] \cup G[B \cup C]$, that is, C separates A from B (see Figure 3). If $G[B \cup C]$ is in the fourth neighbourhood of $u \in B \cup C$, the value of $D_u \cap B$ is completely determined by $D_u \cap C$. In other words, there is a unique best way to extend $D_u \cap C$ to B . Therefore, for such vertices u , there are at most $2^{|C|} \cdot \text{MDS}(G, B \cup C)$ different candidates for $d(u)$.

We now discuss how to use that observation to reach the desired conclusion. We associate to each vertex $u \in V \setminus Y$ an arbitrary vertex $y_u \in Y \cap N(u)$. We consider the plane multigraph H obtained from an arbitrary planar embedding of G by contracting every edge uy_u for $u \in V \setminus Y$, but keeping any resulting loop or multiple edge. Note that $V(H) = Y$ by construction. Let H' be the plane multigraph obtained from H by iteratively deleting any edge incident to a face bounded by one edge (i.e. a loop) or to two faces bounded by two edges (i.e. there is a multiedge of multiplicity 3), see Figure 4. In particular, H' may contain loops and multiple edges, but in the embedding they separate different parts of $V(H')$ or, in the case of multiple edges, are incident on at least one side to a face of size at least 3. We observe that by Euler's formula, there are fewer than $2 \cdot 3|V(H')| = 6|Y|$ edges in H' . To every edge e of H or H' we associate an edge e_G of G such that e_G turns into e in the contraction process. We let Z be the subset of vertices of $V(G) \setminus Y$ which are an endpoint of e_G

²⁸If we require the algorithm to have knowledge the size of the graph, we can have Meta-Theorem III with this additional property. However, to achieve this, we need to change the definition of paddable: we now request, for any n , a graph of size n (instead of size $\geq n$) with the desired properties.




Fig. 3. The depiction of A, B, C in the proof of the key observation.




Fig. 4. The plane graphs H and H' obtained from G .

for some edge $e \in E(H')$. Since every edge e_G contains at most 2 vertices of Z , we have $|Z| \leq 12|Y|$. Therefore, there are at most $13|Y|$ vertices of $Y \cup Z$ in D .

It remains to bound $|D \setminus (Y \cup Z)|$. Note that there are two types of vertices $v \in V(G) \setminus (Y \cup Z)$.

- $N[v] \subseteq N[y_v]$
- $N[v] \not\subseteq N[y_v]$ and v is between two edges e_G, e'_G in G such that e and e' bound a face of size 2 in H'

There cannot be any other case, as if the first case does not apply then v has degree at least 2, so is adjacent to some vertex w such that w is not a neighbour of y_u . But then w gets contracted to a vertex $s \neq y_u$ when creating H . Note that vw does not get contracted and let e be the edge in H such that $e_G = vw$. But since v is not in Z , the edge e gets deleted when creating H' , i.e. e is in between two parallel edges to e and we are in the second case.

We now assume w is a vertex such that $d(w) = v$. If the first case applies then any neighbour of v is also a neighbour of y_v , in particular, the distance of a vertex $w \neq v$ to v is the same as to y_v . However, either y_v has strictly more neighbours than v , or y_v has a smaller label than v , as $y_v \in Y$, hence w would prefer y_v over v .

We analyse now the second case, which is slightly more interesting. Let e_G be incident to x_1, x_2 and e'_G to w_1, w_2 , where for each i , either $d(x_i) = y_i$ or $x_i = y_i$ and similarly for w_i and without loss of generality we assume $y_u = y_1$. For a depiction, see Figure 5. Note that $C = \{x_1, x_2, y_1, y_2, w_1, w_2\}$ separates u from the rest of $V(H')$, and that vertices involved in some edge inside f in H are all




Fig. 5. A depiction of x_1, x_2, w_1, w_2 for the vertex u_1 .

adjacent to y_1 or to y_2 hence are at distance at most 3 to all vertices in f (with respect to u , since u is a neighbour of y_2 or has a neighbour which is a neighbour of y_2 as we are in the second case). But this means that w is at distance at most 4 to all vertices in f . Therefore, there are at most 2^6 choices for D_u and $2^6 \cdot 2$ choices for $d(u)$ if $d(u) \notin Y \cup Z$. This is because not more than 2 vertices from the set contained in the cycle $y_2 x_1 w_1 y_1 w_2 x_2 y_2$ can be chosen as otherwise y_1, y_2 would be a better choice. We can note that in the case where both y_1 and y_2 get selected in D_u there is no extra vertex to add as all of the inside is dominated and when one of y_1 or y_2 is chosen, then there is at most one vertex picked on the inside, otherwise y_1, y_2 would be a better choice. Therefore, we reduce the number of choices to $1 \cdot 2^4 \cdot 2 + 2 \cdot 2^4 \cdot 1 = 64$.

Since there are at most $3|Y|$ options for f , this adds up to $192|Y|$.

All together, we get $13|Y|$ for the set Y or the set Z being chosen, and $192|Y|$ for the choices of the second case of vertices not in Z or Y being chosen. This sums up to $205|Y|$, as claimed. \square

C.2 Meta-theorem for the LOCAL model with no polynomial identifiers requirement

THEOREM 4.2 (META-THEOREM I). *Let A be a k -uniform α -approximation LOCAL algorithm that does not require polynomial identifiers, for an additive ρ -local problem Π in a class of graphs \mathcal{C} with round complexity r , where \mathcal{C} is hereditary and stable by disjoint union. Let \mathcal{D} be a graph class of asymptotic dimension d with d -dimensional control function f that is T -locally δ -nice w.r.t. \mathcal{C} and Π , where $T = f(2k + 2\rho) + \max\{k + \rho, r\}$. Then, there exists a $\max\{k, \rho\}$ -uniform $(\alpha(d + 1) + 1)$ -approximation algorithm B on \mathcal{D} with round complexity $T + \delta + 2$. Furthermore, if \mathcal{D} is T -locally \mathcal{C} , the approximation ratio is $\alpha(d + 1)$ and the additivity condition can be dropped.*

This proof is essentially a simplified proof of Meta-Theorem III. However, we include it for completeness.

PROOF. Let $G \in \mathcal{D}$. As \mathcal{D} has d -dimensional control function f , then $G^{2k+2\rho}$ admits a $(d + 1)$ -coloring $f(2k + 2\rho)$ -bounded in G . Fix such a coloring, and, for each $i \in \{0, 1, \dots, d\}$, let C_i be the set of color- i vertices in $G^{2k+2\rho}$ (and also in G). We denote by \mathfrak{C}_i the set of $(2k + 2\rho)$ -components of C_i (i.e. connected components of $G^{2k+2\rho}[C_i]$). So, by definition, all components of \mathfrak{C}_i have weak diameter in G at most $f(2k + 2\rho)$, and are pairwise at distance at least $2k + 2\rho + 1$ in G (because distinct components of \mathfrak{C}_i cannot be adjacent in $G^{2k+2\rho}$).

We describe the Algorithm B we want to use in Algorithm 3, which is very similar to Algorithm 2.

We run Algorithm B on G . It is easy to check that $B(G)$ is a valid solution for G , because Π is additive. By definition, $|B(G)| \leq |S| + |S'|$, where $S = A(G) \setminus X$ (Step 3) and S' is a brute-forced set of minimum size satisfying the constraints of vertices of G not satisfied (Step 4). Note that every vertex not in $N^\rho[X]$ is satisfied by S , so $S' \subseteq N^{2\rho}[X]$. Clearly, $|S'| \leq \text{OPT}_\Pi(G)$, and $|S'| = 0$ if there

Algorithm 3 Generic algorithm B from Meta-Theorem I.

Require: An ε -weakly k -uniform α -approximate LOCAL algorithm A for Π in \mathcal{C} with round complexity r , a graph classes \mathcal{C} and \mathcal{D} with properties of the statement of Meta-Theorem III, the d -dimensional control function f of \mathcal{D} and $G \in \mathcal{D}$ with T -errors X where $T = f(2k + 2\rho) + \max\{k + \rho, r\}$.

Ensure: A solution S for the graph G to the problem Π with the guarantee that $|S| \leq (\alpha(d + 1) + 1) \cdot \text{OPT}_\Pi(G)$.

- 1: $S \leftarrow \emptyset$
- 2: Each vertex u computes $G[N^T[u]]$ and checks whether it belongs to \mathcal{C} ; as a consequence, it decides whether it belongs to X .
- 3: Each vertex u runs A for r rounds, and gets added to S only if $u \in A(G) \setminus X$.
- 4: $S \leftarrow S \cup S'$, where S' is a brute-forced minimum set of G that satisfies the constraints of vertices that are non-satisfied.

is no T -errors ($X = \emptyset$). Moreover, by definition of the coloring, $|A(G) \setminus X| = \sum_{i=0}^d |(A(G) \cap C_i) \setminus X|$. In other words,

$$|B(G)| \leq \sum_{i=0}^d |(A(G) \cap C_i) \setminus X| + \begin{cases} 0 & \text{if } X = \emptyset \\ \text{OPT}_\Pi(G) & \text{otherwise} \end{cases} \quad (1)$$

To upper bound the term $|(A(G) \cap C_i) \setminus X|$, consider some color $i \in \{0, 1, \dots, d\}$ and some component $C \in \mathcal{C}_i$ such that $C \not\subseteq X$. Consider some $v_C \in C \setminus X$, and denote $G_C = G[N^T[v_C]]$. Note that, by definition of X and as $v \notin X$, we have $G_C \in \mathcal{C}$. Moreover, let $\mathcal{C}'_i = \{C \in \mathcal{C}_i \mid C \not\subseteq X\}$ and let $G' = G[\bigcup_{C \in \mathcal{C}'_i} N^T[v_C]]$ be the union of the G_C 's for $C \in \mathcal{C}_i$ such that $C \not\subseteq X$.

Along the proof, we will use several times the following fact:

Fact C.1. *For any graph H and $W \subseteq V(H)$, to satisfy the constraints of vertices of W in H , it is enough to select vertices taken from $N^\rho[W]$.*

In the following, to differentiate the neighborhoods in G and in G_C or G' , we write them using either N_G , N_{G_C} , or $N_{G'}$ notations. Notice that, for each non-negative integer $t \leq \max\{k + \rho, r\}$, we have $G[N_G^t[C]] \subseteq G_C$. This is because C has weak diameter in G at most $f(2k + 2\rho)$ and thus $G[N_G^t[C]]$ has weak diameter in G at most $f(2k + 2\rho) + t \leq T$, by the choice of T . Moreover, as \mathcal{C} is hereditary, we have that $G[N_G^t[C]] \in \mathcal{C}$.

As $G[N_G^t[C]] \subseteq G_C$, vertices in C have the same distance t -neighborhood in G and G_C . In particular, $G[N_G^t[C]] = G_C[N_{G_C}^t[C]]$ and $G[N_G^t[\bigcup \mathcal{C}'_i]] = G'[N_{G'}^t[\bigcup \mathcal{C}'_i]]$. From Fact C.1, any minimum set of vertices of G that satisfies the constraints of $N_G^k[C]$ is contained in $G[N_G^{k+\rho}[C]]$, and similarly if we replace G by G_C . It follows that $\text{OPT}_\Pi(G, N_G^k[C]) = \text{OPT}_\Pi(G_C, N_{G_C}^k[C])$. Recall that any two connected components $C, C' \in \mathcal{C}_i$ are at distance in G at least $2k + 2\rho + 1$. So $G[N_G^{k+\rho}[C]]$ and $G[N_G^{k+\rho}[C']]$ are disjoint subgraphs, and, as to satisfy all constraints of vertices in $V(G)$, one needs to satisfy constraints of vertices in $N_G^k[C]$ and $N_G^k[C']$, by disjunction of these sets and by Fact C.1, we have by summing over all $C \in \mathcal{C}_i$ that

$$\text{OPT}_\Pi \left(G, N_G^k \left[\bigcup \mathcal{C}'_i \right] \right) = \text{OPT}_\Pi \left(G', N_{G'}^k \left[\bigcup \mathcal{C}'_i \right] \right).$$

Moreover, because \mathcal{C} is stable by disjoint union and that $G[N_G^r[C]] \in \mathcal{C}$ for every $C \in \mathcal{C}'_i$, we get $G[N_G^r[\bigcup \mathcal{C}'_i]] \in \mathcal{C}$. Therefore, as A is a k -uniform α -approximation on \mathcal{C} , we get

$$\left| A(G') \cap \bigcup \mathcal{C}'_i \right| \leq \alpha \cdot \text{OPT}_\Pi \left(G', N_{G'}^k \left[\bigcup \mathcal{C}'_i \right] \right) = \text{OPT}_\Pi \left(G, N_G^k \left[\bigcup \mathcal{C}'_i \right] \right).$$

Since $G[N_G^r[\bigcup \mathcal{C}'_i]] = G'[N_{G'}^r[\bigcup \mathcal{C}'_i]]$, Algorithm A , which runs in r rounds, returns the same vertex set in G and G' for vertices in $\bigcup \mathcal{C}'_i$. Therefore, $|A(G) \cap \bigcup \mathcal{C}'_i| = |A(G') \cap \bigcup \mathcal{C}'_i|$. Combining with the previous inequality, we get

$$\left| A(G) \cap \bigcup \mathcal{C}'_i \right| \leq \alpha \cdot \text{OPT}_\Pi \left(G, N_G^k \left[\bigcup \mathcal{C}'_i \right] \right) \leq \alpha \cdot \text{OPT}_\Pi(G, V(G)) = \alpha \cdot \text{OPT}_\Pi(G).$$

We remark that $|(A(G) \cap C) \setminus X| = 0$ if $C \subseteq X$. Because all vertices in C_i but not in X are in some $C \in \mathcal{C}'_i$, we have

$$|(A(G) \cap C_i) \setminus X| \leq \left| A(G) \cap \bigcup \mathcal{C}'_i \right| \leq \alpha \cdot \text{OPT}_\Pi(G). \quad (2)$$

Putting Eq. (2) in Eq. (1), we get that Algorithm B is an $(\alpha(d+1) + 1)$ -approximation, and even an $\alpha(d+1)$ -approximation if $X = \emptyset$.

To see that B is also a k -uniform approximation, it is sufficient to see that, for any $W \subseteq V(G)$, the bound obtained is $|B(G) \cap W| \leq \alpha(d+1) \cdot \text{OPT}_\Pi(G, N_G^k[W]) + \text{OPT}_\Pi(G, N_G^\rho[W])$ where the first two terms are due to running A and the last term is due to running the brute-force. Indeed, if the brute-force computed a set whose intersection with W was smaller than $\text{OPT}_\Pi(G, N_G^\rho[W])$, we could replace it by the minimum size set satisfying the constraints of the vertices in $N_G^\rho[W]$ and obtain a smaller set, a contradiction. This proves that B is a $\max\{k, \rho\}$ -uniform with an approximate ratio $\alpha(d+1) + 1$ (or $\alpha(d+1)$ if $X = \emptyset$, as the brute-force is not required in that case). Now, let us prove that Algorithm B has the desired round complexity.

- Computing X in Step 2 takes $T + 1$ rounds.
- Running A in Step 3 takes r rounds.
- For Step 4, consider the set S computed by B at Step 3, before the brute-force. Observe that the vertices in the set $W = V(G) \setminus N_G^\rho[X]$ are satisfied by S . This is because vertices of W are at distance at least $\rho + 1$ from X and thus vertices of $N_G^\rho[W]$ cannot be in X . Thus, A applies to all vertices of $N_G^\rho[W]$, and so W is indeed satisfied by S (Fact C.1). Therefore, to satisfy $N_G^\rho[X]$ it is enough to select a set S' from $N_G^{2\rho}[X]$ (Fact C.1) as done in Step 4. By assumption, the connected components of $N_G^{2\rho}[X]$ have weak diameter in G at most δ (obviously, if $X = \emptyset$ then $\delta = 0$). Therefore, the brute-force (Step 5) will take at most $\delta + 1$ rounds.

Naively, Steps 2 and 3 together take $(T + 1) + r$ rounds. As Algorithm A is not guaranteed to work when executed on vertices of X , Step 3 must be run only after Step 2. However, both steps can be run in parallel as follows. We run A for r rounds exactly and stop it just after (since its running time on a vertex of X could result into more than r rounds). Then, the decision to add u in S (if selected by A) is delayed for the next $T + 1 - r$ rounds (this is non-negative from the choice of T). In parallel, $G[N^T[u]]$ is computed and is checked to be in \mathcal{C} or not after $T + 1$ rounds. So, after $\max\{r, T + 1\} = T + 1$ rounds, set S in Step 3 has been completed. Step 4 takes $\delta + 1$ steps, so that the total round complexity of B is $T + \delta + 2$, completing the proof. \square

C.3 MINIMUM DOMINATING SET is uniformizable

Proposition C.2. *Let \mathcal{C} be a hereditary class of unbounded degree²⁹ and that is closed by disjoint union. Then, the MINIMUM DOMINATING SET problem is uniformizable in the LOCAL model with binding function $(r, \alpha) \mapsto (r + 1, \alpha + \varepsilon)$ for any real $\varepsilon > 0$.*

PROOF. We first prove this statement in the LOCAL model with $\varepsilon = 0$. Assume for sake of contradiction that there exists integers d, r and a real α such that for all integers k and real β , there exists an r -round α -approximation A of Π in constant time on the hereditary graph class \mathcal{C} such that A is not k -uniform with β -approximation of Π in constant time. Here, we will take $k = r + 1$ and $\beta = \alpha$. By the definition of uniformity, this means that there exists some subset S of vertices for which:

$$|A(G) \cap S| > \alpha \cdot \text{MDS}(G, N^k[S]).$$

Let us cut G around S : set $G' := G[N^k[S] \cup X]$ where $X \subseteq N^{k+1}[S]$ is a minimum dominating set of $N^k[S]$ in G . We keep the same identifiers in G' as in G . We first bound the size of the minimum dominating set of G' .

Claim C.3. $\text{MDS}(G') \leq \text{MDS}(G, N^k[S])$.

PROOF OF CLAIM. By definition of X , it follows that $|X| = \text{MDS}(G, N^k[S])$. X is also a dominating set of G' , because each of its vertices is either in X or in $N^k[S]$, and that in the second case, it is dominated by some vertex of X by definition of X . Therefore, $\text{MDS}(G') \leq |X| = \text{MDS}(G, N^k[S])$. \diamond

As our goal is to get a contradiction on $G' \in \mathcal{C}$ ³⁰ instead of G , we prove the following.

Claim C.4. *On S , algorithm A returns the same set in G and in G' .*

This will prove that $|A(G) \cap S| = |A(G') \cap S|$.

PROOF OF CLAIM. A is an r -round algorithm. Therefore, vertices in S only see vertices and edges in $G[N^r[S]]$, and edges between vertices of $N^r(S)$ and vertices of $N^{r+1}(S)$ ³¹. All of those are the same in G and G' , therefore $A(G) \cap S = A(G') \cap S$, as wanted. \diamond

By Claim C.4,

$$|A(G')| \geq |A(G') \cap S| = |A(G) \cap S| > \alpha \text{MDS}(G, N^k[S]).$$

By Claim C.3, we have $\alpha \text{MDS}(G, N^k[S]) \geq \alpha \cdot \text{MDS}(G')$. It follows that $|A(G')| \geq \alpha \cdot \text{MDS}(G')$, a contradiction as $G' \in \mathcal{C}$ and A is an α -approximation on \mathcal{C} . Let us now adapt this proof for the LOCAL model, for $\varepsilon > 0$. For now, by taking $\gamma = \alpha + \varepsilon$, we have proved $|A(G')| \geq (\alpha + \varepsilon) \cdot \text{MDS}(G')$. The only thing we need to take care of is that nodes in G have identifiers that are bounded by a polynomial of $|V(G')|$. Instead of changing the identifiers in G' without impacting the output of A , which is tricky, we consider a bigger graph. Let N be the biggest identifier in G' . \mathcal{C} is a hereditary graph class of unbounded degree, so there exists a graph $H \in \mathcal{C}$ and a vertex v of H with degree at least N .³² Take $H' = H[N_H[v]]$ and give the vertices identifiers in $[N + 1, 2N]$. Let $G'' = G' \sqcup H'$.

²⁹For a graph class in which every graph has maximum degree Δ , the LOCAL algorithm that includes all nodes is a $\Delta + 1$ -approximation. This implies that the bounded-degree case is less interesting than the unbounded-degree one, which is therefore the focus of our study.

³⁰because \mathcal{C} is hereditary

³¹The definition of an r -round LOCAL algorithm running on vertex v includes edges at distance $r + 1$ of v but not vertices at distance $r + 1$ of v in G .

³²Here, if the algorithm requires access to n (the size of the graph), one needs to choose $N \geq n$ and delete neighbors of v so that, in the following, we get $|V(G'')| = n$.

Clearly, H' can be dominated by choosing the single vertex v , i.e. $\text{MDS}(G'') = \text{MDS}(G') + 1$. Furthermore, $A(G'') = A(G') \sqcup A(H')$ and it follows

$$|A(G'')| \geq (\alpha + \varepsilon) \cdot \text{MDS}(G') + 1 = (\alpha + \varepsilon) \cdot \text{MDS}(G'') + 1 - (\alpha + \varepsilon).$$

Now, we can apply the same trick as in Proposition D.2. Observe that the diameter of each connected component of the graph G' is at most $3 \text{MDS}(G')$. We construct an algorithm B where each vertex checks whether the diameter D of its component is less than $3B$, for $B = (\alpha - 1)/\varepsilon$. This can be done by collecting its radius- $3B$ neighborhood in $3B$ rounds. If true, the vertex locally brute forces an optimal solution for its component, which it already knows about. If false, the vertex applies the approximate algorithm A . If $D \leq B$, then the algorithm is obviously a good approximation (with approximation ratio 1). Otherwise, $D > 3B$ and it follows that $\text{MDS}(G') > B$. One can verify that $|A(G'')| > \alpha \cdot \text{MDS}(G'')$, a contradiction, as $G'' \in \mathcal{C}$. \square

C.4 MINIMUM k -TUPLE DOMINATING SET is cuttable and additive

Proposition C.5. *MINIMUM k -TUPLE DOMINATING SET is additive and cuttable with parameter 1.*

PROOF. By definition, the problem is additive. Let us now prove cutability. Let G be a graph and $X \subseteq V(G)$, and let S be a k -Tuple Dominating Set of the set X . Set $Y = X \cup S$. S is a k -Tuple Dominating Set of $G[Y]$. Indeed, every vertex of $G[Y]$ that is in $Y \setminus S$ is in $X \setminus S$, and therefore is k -Tuple Dominated by S . Thus for $\Pi = \text{MINIMUM } k\text{-TUPLE DOMINATING SET}$, we have $\text{OPT}_{\Pi}(G, X) \geq \text{OPT}_{\Pi}(G[Y])$. \square

C.5 DISTANCE- $\leq d$ DOMINATING SET is cuttable

Proposition C.6. *DISTANCE- $(\leq d)$ DOMINATING SET is cuttable with parameter 1.*

PROOF. Let Π be the DISTANCE- $(\leq d)$ DOMINATING SET problem, let G be a graph and X a subset of vertices. Let D be a minimum-size set that is distance- d dominating for the vertices in X , so that $|D| = \text{OPT}_{\Pi}(G, X)$. Note that $D \subseteq N^d[X]$. We are going to construct a set $X \subseteq Y \subseteq N^d[X]$ such $\text{OPT}_{\Pi}(G[Y]) \leq \text{OPT}_{\Pi}(G, X)$. We do this recursively. Start with $Y := X$. We maintain the invariant that every vertex of Y that is not dominated by $D \cap Y$ in $G[Y]$ is in X . This invariant holds trivially for $Y = X$. While some vertex v of X is not dominated in $G[Y]$ by $D \cap Y$, do the following. Let $y \in D$ be a vertex that dominates v in G , i.e., with the distance to y at most d . Note that here, we do not have necessarily $y \notin Y$: it could be the case that v is dominated by some $y \in X$ in G because of a path with a portion outside X . As $y \in D \subseteq N^d[X]$, there is a path P in $G[N^d[X]]$ from v to y with length at most d . Actually, one can just take the shortest vy -path. It has length at most d and is in $N^d[v] \subseteq N^d[X]$. Set $Y := Y \cup V(P)$. All vertices that are in P (including v) are dominated by y , and therefore are now dominated in $G[Y]$. The invariant is thus satisfied. Each time we do this operation, at least one more vertex of X is dominated in $G[Y]$. As X is finite, this process terminates and we have constructed a set Y that is dominated by $D \cap Y$ in $G[Y]$. It follows that $\text{OPT}_{\Pi}(G[Y]) \leq \text{OPT}_{\Pi}(G, X)$, and this completes the proof. \square

C.6 Cutting implies uniform algorithms in the LOCAL model with no polynomial identifiers requirement

A problem Π is *uniformizable* on \mathcal{C} if for all integers r and all reals α , there exists some integers k and real β such that every r -round α -approximation A of Π on \mathcal{C} is a k -uniform β -approximation of Π on \mathcal{C} . The function $(r, \alpha) \mapsto (k, \beta)$ is called the *binding* function.

Proposition C.7. *Let \mathcal{C} be a hereditary class that is closed by disjoint union, and let Π be a ρ -local problem which is cuttable with parameter β . Then, Π is uniformizable in the LOCAL model with no polynomial identifiers requirement with binding function $(r, \alpha) \mapsto (r + 1, \alpha \cdot \beta)$.*

PROOF. Assume for sake of contradiction that there exist an integer r and a real α such that for all integers k, r' and real γ , there exists an r -round α -approximation A of Π on the hereditary graph class \mathcal{C} but no r' -round k -uniform γ -approximation algorithm B for Π on \mathcal{C} . Here, we take $k = r + 1$, $\gamma = \alpha \cdot \beta$, $r' = r$ and $B = A$, i.e. suppose A is not a k -uniform γ -approximation of Π . By the definition of uniformity, this means that there exists some subset S of vertices for which:

$$|A(G) \cap S| > \alpha\beta \cdot \text{OPT}_{\Pi}(G, N^k[S]).$$

As Π is cuttable with the particular case $X = N^k[S]$, there exists a set Y such that $\text{OPT}_{\Pi}(G[Y]) \leq \beta \cdot \text{OPT}_{\Pi}(G, N^k[S])$. We now study the graph $G' := G[Y]$, with the same identifiers as in G . We prove the following.

Claim C.8. *We have $A(G) \cap S = A(G') \cap S$.*

PROOF OF CLAIM. This is the same proof as in Claim C.4, except with a set Y instead. As A is an r -round algorithm, vertices in S only see vertices and edges in $G[N^r[S]]$, and edges between vertices of $N^r(S)$ and vertices of $N^{r+1}(S)$. Those are the same in G and G' as $N^{r+1}[S] = N^k[S] \subseteq Y$, and so we have $A(G) \cap S = A(G') \cap S$, as wanted. \diamond

By Claim C.8,

$$|A(G')| \geq |A(G') \cap S| = |A(G) \cap S| > \alpha\beta \cdot \text{OPT}_{\Pi}(G, N^k[S]).$$

By the cutability, we have $\text{OPT}_{\Pi}(G') \leq \beta \cdot \text{OPT}_{\Pi}(G, N^k[S])$. It follows that $|A(G')| > \alpha \cdot \text{OPT}_{\Pi}(G')$, a contradiction as $G' \in \mathcal{C}$ and A is an α -approximation on \mathcal{C} . \square

C.7 The algorithm of Heydt et al. is uniform

Observation C.9. *The algorithm described in [HKOdM⁺25, Theorem 2.3] is a $k(\varepsilon)$ -uniform $(11 + \varepsilon)$ -approximation LOCAL algorithm for MINIMUM DOMINATING SET in planar graphs with round complexity $C(\varepsilon)$, for every $\varepsilon > 0$, and for some functions C and k .*

PROOF. Let us consider the algorithm described in [HKOdM⁺25, Theorem 2.3]. It has a number of rounds bounded by a function of ε , and does not need polynomial identifiers so Observation C.9 follows by Proposition C.2. \square

For the interested reader, here is a more detailed analysis of the running time of the algorithm. It returns a solution $D_1 \cup D_2 \cup D_3^2 \cup D_3^1$. D_1 can be computed in two rounds, D_2 in two additional rounds, and D_3^1 in one more additional round. The set D_3^2 can be computed in time $O(\log \Delta / \log(1 + \varepsilon))$ where $\Delta = 4\nabla_1 \cdot (4^{\nabla_1} + 2\nabla_1)(\Delta_R + 1)/\varepsilon$ with (with the particular case of planar graphs) $\nabla_1 < 3$ and $\Delta_R = \kappa^{s-1}(t + s - 1 + (s - 1)\kappa^3)$ with $\kappa = \max\{2\nabla_0, 2\nabla\}$. Again, for planar graphs, $\nabla_0 < 3$ and $\nabla \leq 2$ so $\kappa \leq 6$. It follows that $\Delta_R \leq 15732$ because $s = t = 3$ on planar graphs. So $\Delta \leq 13593312/\varepsilon = O(1/\varepsilon)$ and the running time is $O(\log(1/\varepsilon)/\log(1 + \varepsilon/2)) = O(\log(1/\varepsilon)/\varepsilon)$.

C.8 Cutability implies weakly uniform algorithms

Corollary C.10. *Let \mathcal{C} be a hereditary class closed by disjoint union, and let Π be a ρ -local problem which is cutable with parameter β . Then, for any real $\varepsilon > 0$, the problem Π is ε -weakly uniformizable in the LOCAL model with (weak) binding function $(r, \alpha) \mapsto (r + 1, \alpha \cdot \beta)$.*

PROOF. Let us go through the proof of Proposition C.7 one more time. For sake of brevity, we do not copy the whole identical proof, but instead give a summary of the proof. Assume for sake of contradiction that there exists integers d, r and a real α such that for all integers k and real γ , there exists an r -round α -approximation A of Π in constant time on the hereditary graph class \mathcal{C} such that A is not ε -weakly k -uniform with γ -approximation of Π in constant time. Here, we will take $k = r + 1$ and $\gamma = \alpha \cdot \beta$. By the definition of weak uniformity, this means that there exists some subset $S \subseteq V(G)$ of size at least $\varepsilon|V(G)|$ for which:

$$|A(G) \cap S| > \alpha\beta \cdot \text{OPT}_{\Pi}(G, N^k[S]).$$

We have showed there exists a set Y such that $\beta \cdot \text{OPT}_{\Pi}(G, N^k[S]) \geq \text{OPT}_{\Pi}(G[Y])$. Let $G' := G[Y]$. We keep the same identifiers in G' as in G . We have showed that $|A(G')| \geq \alpha \cdot \text{OPT}_{\Pi}(G')$, which would be a contradiction with the fact that α -approximation on \mathcal{C} , if A was an algorithm in the LOCAL model with no polynomial identifiers requirement. For this to be a contradiction in the LOCAL model, we have to show that the IDs in G' are also polynomial in $|V(G')|$. Now, we just use the additional property that S has large size. All the identifiers in G' are polynomial in $|V(G)|$, say smaller than $D \cdot |V(G)|^c$ for some real $c > 0$. As $|V(G')| \geq |S| \geq \varepsilon \cdot |V(G)|$, the identifier in G' are smaller than $D \cdot |V(G')|^c / \varepsilon^c$, which is still polynomial. This finishes the proof. \square

C.9 Bounded Euler genus graphs are locally nice

Proposition C.11. *Let Π be a ρ -local problem. Then, for any $g \geq 1$, the class of Euler genus- g graphs is T -locally δ -nice w.r.t. planar graphs and Π , where $\delta < g \cdot (2T + 4\rho + 1)$.*

PROOF. To derive a contradiction, consider a connected component H of $G[N^{2\rho}[X]]$, and assume that H has weak diameter $\delta \geq g \cdot (2T + 4\rho + 1)$ in G . As $g \geq 1$, X and H are not empty, and H is not planar. So, H must contain some path P such that the distance in G between its endpoints, say s to t , is δ . In particular, for each $d \in \{0, 1, \dots, \delta\}$, there must exist a vertex of P that is at distance exactly d in G from s . This is because the distance in G from s , when moving along an edge of P , is a function that can vary by at most one unit, and this distance goes from 0 (at s) to δ (at t). For every $i \in \{0, 1, \dots, g\}$, let u_i be any vertex of P at distance exactly $d_i = i \cdot (2T + 4\rho + 1)$ in G from s . So $u_0 = s$, $u_g = t$, and all the u_i 's exist in P (so in H) since $d_i \in \{0, 1, \dots, \delta\}$. By definition of H , each u_i intersects $N^{2\rho}[X]$. So, for each u_i one can select a vertex $x_i \in X$ at distance in G at most 2ρ from u_i . By the triangle inequality, for all $0 \leq i < j \leq g$, x_i and x_j are at distance in G at least $(d_j - 2\rho) - (d_i + 2\rho) = (j - i) \cdot (2T + 4\rho + 1) - 4\rho \geq 2T + 1$. It follows that $N^T[x_i]$ and $N^T[x_j]$ are disjoint. By definition of X and H , the subgraphs $G[N^T[x_i]]$'s are not planar, thus of Euler genus ≥ 1 . By additivity of the Euler genera³³ of these $g + 1$ pairwise disjoint subgraphs of G (cf. [MT01, Theorem 4.4.3]), we get a contradiction that G has Euler genus g . Thus $\delta < g \cdot (2T + 4\rho + 1)$ as claimed. \square

³³The plural of genus.

C.10 Balanced asymptotic dimension

Lemma C.12. *Let \mathcal{G} be a graph class with asymptotic dimension at most d and control function f . For every graph $G \in \mathcal{G}$ with n vertices and every integer $r \geq 1$, there exists a cover C_0, C_1, \dots, C_d of $V(G)$ such that*

- every r -component of each C_i has weak diameter in G at most $f(3r) + 2r$, and
- for every $i \in \{0, \dots, d\}$ we have $|C_i| \geq \lfloor \frac{n}{d+1} \rfloor$.

PROOF. Let $q := \lfloor \frac{n}{d+1} \rfloor$ and r be any integer. By Proposition 2.1 applied to scale $3r$, there exists a $(d+1)$ -coloring of G^{3r} whose monochromatic components have weak diameter at most $f(3r)$ (as they have weak diameter $f(3r)$ in G^{3r}) in G . Let C_0, C_1, \dots, C_d be the corresponding color classes, with $|C_0| \leq |C_1| \leq \dots \leq |C_d|$ without loss of generality. They form a cover of $V(G)$, and every $3r$ -component of each C_i has weak diameter at most $f(3r)$ in G .

We construct sets C'_0, \dots, C'_d by proving the following by induction of i . For every i , there exists sets C'_0, \dots, C'_{i-1} such that

(I1) for every $j \leq i$, we have $|C'_j| \geq q$, and if $|C'_j| > q$ then $C'_j = C_j$,

(I2) every r -component of each C'_j (for $j \leq i$) has weak diameter at most $f(3r) + 2r$ in G , and

(I3) $\bigcup_{j \leq i} C'_j \supseteq \bigcup_{j \leq i} C_j$.

Note that (I3) says that the sets C'_0, \dots, C'_i cover at least every vertex of C_0, \dots, C_i . For $i = 0$ these conditions are trivially satisfied. Assume the conditions hold for some $i \in \{0, \dots, d\}$. Let $A := N^r[C_i] \setminus C_i$ and $B := (V(G) \setminus N^r[C_i]) \setminus C_i$, such that $A \sqcup B = V(G) \setminus C_i$. Intuitively, we will try to add vertices to U from A whenever possible (these are within distance r of the original C_i), and otherwise take vertices from B (which are farther from C_i).

If $|C_i| \geq q$, then take $C'_j = C_j$ for all $j \geq i$. (I1) is satisfied because $q \leq |C_i| \leq \dots \leq |C_d|$. (I2) is satisfied because the C_i 's have all weak diameter at most $f(3r)$, and (I3) is trivially satisfied. So assume $|C_i| < q$ and let $s := q - |C_i| > 0$ be the number of new vertices we must add to C_i to reach size q . There are two possibilities.

- **Case 1:** $|A| \geq s$. Choose an arbitrary set $X \subseteq A$ of size s and let $C'_i := C_i \cup X$; this satisfies (I1) and (I3). Let us prove (I2). Let $x, x' \in A$ be two vertices in the same r -component of C'_i . As $C'_i \subseteq N^r[C_i]$, there exist $y, y' \in C_i$ such that $d(x, y), d(x', y') \leq r$. We first prove that y and y' are in the same $3r$ -component of C_i . Because x and x' are in the same r -component of C'_i , exists a sequence $x = x_1, x_2, \dots, x_t = x'$ in C'_i with $d(x_j, x_{j+1}) \leq r$ for all j . As $C'_i \subseteq N^r[C_i]$, we can choose, for each j , some $y_j \in C_i$ such that $d(x_j, y_j) \leq r$. Without loss of generality, choose $y_1 := y$ and $y_t := y'$. Then for any j , we have that

$$d(y_j, y_{j+1}) \leq d(y_j, x_j) + d(x_j, x_{j+1}) + d(x_{j+1}, y_{j+1}) \leq 3r,$$

and so y and y' are in the same $3r$ -component of C_i , which means that $d(y, y') \leq f(3r)$. Hence

$$\begin{aligned} d(x, x') &\leq d(x, y) + d(y, y') + d(y', x') \\ &\leq r + f(3r) + r = f(3r) + 2r, \end{aligned}$$

as required. This proves (I2).

- **Case 2:** $|A| < s$. Then by definition of B ,

$$|B| \geq |V(G)| - |A| - |C_i| \geq n - s + 1 - |C_i| = n - q + 1.$$

As the sets C_j for $j \neq i$ cover B , there exists some $j \neq i$ such that $|C_j \cap B| \geq (n - q + 1)/d \geq q$. Choose an arbitrary set $X \subseteq C_j \cap B$ of size s and let $C'_i := C_i \cup X$; this satisfies (I1) and (I3). (I2) is true all r -components of C'_i are either r -components of C_i (so they have diameter at most $f(3r)$) or are subsets of r -components of C_j (and the same diameter bound applies).

In both cases we have constructed sets C'_i that satisfies (I1), (I2) and (I3). By (I3), $\bigcup_{i=0}^d C'_i = \bigcup_{i=0}^d C_i = V(G)$, so the C'_i form a cover. (I1) and (I2) prove the two wanted conditions on the cover. This finishes the proof. \square

C.11 Meta-theorem for non-paddable problems with empty error set

THEOREM A.1 (META-THEOREM II). *Let d, k be integers $\alpha > 0$ a real, A be a $\frac{1}{d+1}$ -weakly k -uniform α -approximation LOCAL algorithm for a ρ -local problem Π in a class of graphs \mathcal{C} with round complexity r , where \mathcal{C} is hereditary and stable by disjoint union. Let \mathcal{D} be a graph class of asymptotic dimension d with d -dimensional control function f , that is T -locally- \mathcal{C} for $T = f(6k + 6\rho) + 4k + 4\rho + \max\{k + \rho, r\}$. Then A is a k -uniform $\alpha(d + 1)$ -approximation algorithm on \mathcal{D} .*

PROOF. Let $G \in \mathcal{D}$ and set $\ell := f(6k + 6\rho) + 4k + 4\rho$. As \mathcal{D} has d -dimensional control function f , then by Lemma C.12, one can assume $G^{2k+2\rho}$ admits a $(d + 1)$ -coloring ℓ -bounded in G where every color class has size at least $\lfloor \frac{n}{d+1} \rfloor$. Fix such a coloring, and, for each $i \in \{0, 1, \dots, d\}$, let C_i be the set of color- i vertices in $G^{2k+2\rho}$ (and also in G), so that $|C_i| \geq \lfloor \frac{n}{d+1} \rfloor$. We denote by \mathfrak{C}_i the set of $(2k + 2\rho)$ -components of C_i (i.e. connected components of $G^{2k+2\rho}[C_i]$). So, by definition, all components of \mathfrak{C}_i have weak diameter in G at most ℓ , and are pairwise at distance at least $2k + 2\rho + 1$ in G (because distinct components of \mathfrak{C}_i cannot be adjacent in $G^{2k+2\rho}$).

We run Algorithm A on G . It is easy to check that $A(G)$ is a valid solution for G , because \mathcal{D} is T -locally- \mathcal{C} . By definition of the coloring,

$$|A(G)| = \sum_{i=0}^d |A(G) \cap C_i|. \quad (3)$$

To upper bound the term $|A(G) \cap C_i|$, consider some color $i \in \{0, 1, \dots, d\}$ and some component $C \in \mathfrak{C}_i$. Consider some $v_C \in C$, and denote $G_C = G[N^T[v_C]]$. Note that as \mathcal{D} is T -locally- \mathcal{C} , we have $G_C \in \mathcal{C}$. Let $G' = G[\bigcup_{C \in \mathfrak{C}_i} N^T[v_C]]$ be the union of the G_C 's for $C \in \mathfrak{C}_i$.

Along the proof, we will use several times the following fact:

Fact C.1. *For any graph H and $W \subseteq V(H)$, to satisfy the constraints of vertices of W in H , it is enough to select vertices taken from $N^\rho[W]$.*

In the following, to differentiate the neighborhoods in G and in G_C or G' , we write them using either N_G, N_{G_C} , or $N_{G'}$ notations. Notice that, for each non-negative integer $t \leq \max\{k + \rho, r\}$, we have $G[N_G^t[C]] \subseteq G_C$. This is because C has weak diameter in G at most $\ell = f(6k + 6\rho) + 4k + 4\rho$ and thus $G[N_G^t[C]]$ has weak diameter in G at most $f(6k + 6\rho) + 4k + 4\rho + t \leq T$, by the choice of T . Moreover, as \mathcal{C} is hereditary, we have that $G[N_G^t[C]] \in \mathcal{C}$.

As $G[N_G^t[C]] \subseteq G_C$, vertices in C have the same distance t -neighborhood in G and G_C . In particular, $G[N_G^t[C]] = G_C[N_{G_C}^t[C]]$ and $G[N_G^t[C_i]] = G'[N_{G'}^t[C_i]]$. From Fact C.1, any minimum set of vertices of G that satisfies the constraints of $N_G^k[C]$ is contained in $G[N_G^{k+\rho}[C]]$, and similarly

if we replace G by G_C . It follows that $\text{OPT}_\Pi(G, N_G^k[C]) = \text{OPT}_\Pi(G_C, N_{G_C}^k[C])$. Recall that any two connected components $C, C' \in \mathfrak{C}_i$ are at distance in G at least $2k + 2\rho + 1$. So $G[N_G^{k+\rho}[C]]$ and $G[N_G^{k+\rho}[C']]$ are disjoint subgraphs, and, as to satisfy all constraints of vertices in $V(G)$, one needs to satisfy constraints of vertices in $N_G^k[C]$ and $N_G^k[C']$, by disjunction of these sets and by Fact C.1, we have by summing over all $C \in \mathfrak{C}_i$ that

$$\text{OPT}_\Pi(G, N_G^k[C_i]) = \text{OPT}_\Pi(G', N_{G'}^k[C_i]).$$

Moreover, because \mathcal{C} is stable by disjoint union and that $G[N_G^t[C]] \in \mathcal{C}$ for every $C \in \mathfrak{C}'_i$, we get $G[N_G^t[C_i]] \in \mathcal{C}$. As Algorithm A is a $\frac{1}{d+1}$ -weakly k -uniform α -approximation on \mathcal{C} , that $G' \in \mathcal{C}$ and that $|C_i| \geq \lfloor \frac{|V(G)|}{d+1} \rfloor \geq \lfloor \frac{|V(G')|}{d+1} \rfloor$ (and therefore G' indeed has polynomial identifiers), we get by definition of weak uniformity that

$$|A(G') \cap C_i| \leq \alpha \cdot \text{OPT}_\Pi(G', N_{G'}^k[C_i]) = \alpha \cdot \text{OPT}_\Pi(G, N_G^k[C_i]).$$

Since $G[N_G^r[C_i]] = G'[N_{G'}^r[C_i]]$, Algorithm A, which runs in r rounds, returns the same vertex set in G and G' for vertices in C_i . Therefore, $|A(G) \cap C_i| = |A(G') \cap C_i|$. Combining with the previous inequality, we get

$$\begin{aligned} |A(G) \cap C_i| &\leq \alpha \cdot \text{OPT}_\Pi(G, N_G^k[C_i]) \leq \alpha \cdot \text{OPT}_\Pi(G, V(G)) \\ &= \alpha \cdot \text{OPT}_\Pi(G). \end{aligned} \quad (4)$$

Putting Eq. (4) in Eq. (3), we get that Algorithm A is an $\alpha(d+1)$ -approximation on \mathcal{D} . \square

C.12 Meta-theorem for paddable problems with non-empty error set

THEOREM C.13 (META-THEOREM III). *Let d, k be integers $\alpha > 0$ a real, A be a ε -weakly³⁴ k -uniform α -approximation LOCAL algorithm for an additive ρ -local problem Π in a class of graphs \mathcal{C} with round complexity r , where \mathcal{C} is hereditary, stable by disjoint union, and $O(1)$ -paddable for Π . Let \mathcal{D} be a graph class of asymptotic dimension d with d -dimensional control function f such that Π is dense in \mathcal{D} , and such that \mathcal{D} is T -locally δ -nice w.r.t. \mathcal{C} and Π , where $T = f(2k + 2\rho) + \max\{k + \rho, r\}$. Then there exists a function g such that for any $\eta > 0$, there exists a $\max\{k, \rho\}$ -uniform $(\alpha(d+1) + 1 + \eta)$ -approximation algorithm B on \mathcal{D} with round complexity $\max\{T + \delta + 2, g(\eta)\}$.*

PROOF. Let $G \in \mathcal{D}$ and X be the set of T -errors of G with respect to \mathcal{C} . As \mathcal{D} has d -dimensional control function f , then $G^{2k+2\rho}$ admits a $(d+1)$ -coloring $f(2k+2\rho)$ -bounded in G . Fix such a coloring, and, for each $i \in \{0, 1, \dots, d\}$, let C_i be the set of color- i vertices in $G^{2k+2\rho}$ (and also in G). We denote by \mathfrak{C}_i the set of $(2k+2\rho)$ -components of C_i (i.e. connected components of $G^{2k+2\rho}[C_i]$). So, by definition, all components of \mathfrak{C}_i have weak diameter in G at most $f(2k+2\rho)$, and are pairwise at distance at least $2k + 2\rho + 1$ in G (because distinct components of \mathfrak{C}_i cannot be adjacent in $G^{2k+2\rho}$).

We describe the Algorithm B we want to use in Algorithm 4, which is very similar to Algorithm 2.

We run Algorithm B on G . It is easy to check that $B(G)$ is a valid solution for G , because Π is additive. By definition, $|B(G)| \leq |S| + |S'|$, where $S = A(G) \setminus X$ (Step 4) and S' is a brute-forced set of minimum size satisfying the constraints of vertices of G not satisfied (Step 5). Note that every vertex not in $N^\rho[X]$ is satisfied by S , so $S' \subseteq N^{2\rho}[X]$. Clearly, $|S'| \leq \text{OPT}_\Pi(G)$, and $|S'| = 0$ if there

³⁴Note that, for the proof, we could also prove as a preliminary that weakly uniform and paddable implies uniform. However, for sake of brevity, we prove everything directly in the proof.

Algorithm 4 Generic algorithm B from Meta-Theorem III.

Require: An ε -weakly k -uniform α -approximate LOCAL algorithm A for Π in \mathcal{C} with round complexity r , a graph classes \mathcal{C} and \mathcal{D} with properties of the statement of Meta-Theorem III, the d -dimensional control function f of \mathcal{D} and $G \in \mathcal{D}$ with T -errors X where $T = f(2k + 2\rho) + \max\{k + \rho, r\}$.

Ensure: A solution S for the graph G to the problem Π with the guarantee that $|S| \leq (\alpha(d + 1) + 1) \cdot \text{OPT}_\Pi(G)$.

- 1: Let B be such that any graph G of \mathcal{D} with $\text{diam}(G) \geq B$ has $\text{OPT}_\Pi(G) \geq \alpha(d + 1)\omega/\eta$. If $\text{diam}(G) < B$, brute-force an optimal solution. $\triangleright B$ exists as Π is dense in \mathcal{D} .
- 2: $S \leftarrow \emptyset$
- 3: Each vertex u computes $G[N^T[u]]$ and checks whether it belongs to \mathcal{C} ; as a consequence, it decides whether it belongs to X .
- 4: Each vertex u runs A for r rounds, and gets added to S only if $u \in A(G) \setminus X$.
- 5: $S \leftarrow S \cup S'$, where S' is a brute-forced minimum set of G that satisfies the constraints of vertices that are non-satisfied.

is no T -errors ($X = \emptyset$). Moreover, by definition of the coloring, $|A(G) \setminus X| = \sum_{i=0}^d |(A(G) \cap C_i) \setminus X|$. In other words,

$$|B(G)| \leq \sum_{i=0}^d |(A(G) \cap C_i) \setminus X| + \begin{cases} 0 & \text{if } X = \emptyset \\ \text{OPT}_\Pi(G) & \text{otherwise} \end{cases} \quad (5)$$

To upper bound the term $|(A(G) \cap C_i) \setminus X|$, consider some color $i \in \{0, 1, \dots, d\}$ and some component $C \in \mathfrak{C}_i$ such that $C \not\subseteq X$. Consider some $v_C \in C \setminus X$, and denote $G_C = G[N^T[v_C]]$. Note that, by definition of X and as $v \notin X$, we have $G_C \in \mathcal{C}$. Moreover, let $\mathfrak{C}'_i = \{C \in \mathfrak{C}_i \mid C \not\subseteq X\}$ and let $G' = G[\bigcup_{C \in \mathfrak{C}'_i} N^T[v_C]]$ be the union of the G_C 's for $C \in \mathfrak{C}'_i$ such that $C \not\subseteq X$.

Along the proof, we will use several times the following fact:

Fact C.1. *For any graph H and $W \subseteq V(H)$, to satisfy the constraints of vertices of W in H , it is enough to select vertices taken from $N^p[W]$.*

In the following, to differentiate the neighborhoods in G and in G_C or G' , we write them using either N_G , N_{G_C} , or $N_{G'}$ notations. Notice that, for each non-negative integer $t \leq \max\{k + \rho, r\}$, we have $G[N_G^t[C]] \subseteq G_C$. This is because C has weak diameter in G at most $f(2k + 2\rho)$ and thus $G[N_G^t[C]]$ has weak diameter in G at most $f(2k + 2\rho) + t \leq T$, by the choice of T . Moreover, as \mathcal{C} is hereditary, we have that $G[N_G^t[C]] \in \mathcal{C}$.

As $G[N_G^t[C]] \subseteq G_C$, vertices in C have the same distance t -neighborhood in G and G_C . In particular, $G[N_G^t[C]] = G_C[N_{G_C}^t[C]]$ and $G[N_G^t[\bigcup \mathfrak{C}'_i]] = G'[N_{G'}^t[\bigcup \mathfrak{C}'_i]]$. From Fact C.1, any minimum set of vertices of G that satisfies the constraints of $N_G^k[C]$ is contained in $G[N_G^{k+\rho}[C]]$, and similarly if we replace G by G_C . It follows that $\text{OPT}_\Pi(G, N_G^k[C]) = \text{OPT}_\Pi(G_C, N_{G_C}^k[C])$. Recall that any two connected components $C, C' \in \mathfrak{C}_i$ are at distance in G at least $2k + 2\rho + 1$. So $G[N_G^{k+\rho}[C]]$ and $G[N_G^{k+\rho}[C']]$ are disjoint subgraphs, and, as to satisfy all constraints of vertices in $V(G)$, one needs to satisfy constraints of vertices in $N_G^k[C]$ and $N_G^k[C']$, by disjunction of these sets and by Fact C.1, we have by summing over all $C \in \mathfrak{C}_i$ that

$$\text{OPT}_\Pi \left(G, N_G^k \left[\bigcup \mathfrak{C}'_i \right] \right) = \text{OPT}_\Pi \left(G', N_{G'}^k \left[\bigcup \mathfrak{C}'_i \right] \right).$$

Moreover, because \mathcal{E} is stable by disjoint union and that $G[N_G^r[C]] \in \mathcal{E}$ for every $C \in \mathcal{C}'_i$, we get $G[N_G^r[\bigcup \mathcal{C}'_i]] \in \mathcal{E}$.

In the following, we will choose a graph \widehat{G} with size linear in $|V(G)|$ so that adding \widehat{G} to any subgraph H of G' will form a graph with size linear in $|V(\widehat{G} \sqcup H)|$. This serves two purposes: we are now able to apply weak uniformity on $\widehat{G} \sqcup H$ and we can attribute unique identifiers to \widehat{G} so that the identifiers of $\widehat{G} \sqcup H$ are polynomial (we keep the same identifiers for H). Let $N := \lceil \varepsilon |V(G)| / (1 - \varepsilon) \rceil$. By paddability, there exists a graph $\widehat{G} \in \mathcal{E}$ on at least N vertices such that $\text{OPT}_\Pi(\widehat{G}) \leq \omega$. Let I be the largest identifier of G . We give identifiers to vertices of \widehat{G} from the set $\{I + 1, I + 2, \dots, I + |V(\widehat{G})|\}$. Let G'' be the disjoint union of G' and \widehat{G} and let $Y = \bigcup \mathcal{C}'_i \sqcup V(\widehat{G})$. The identifiers of vertices in G'' are indeed polynomial in the size of G'' , as \widehat{G} has size at least linear in $|V(G)|$. Our goal now is to apply weak uniformity in G'' to the set Y . Note that we only care about $|A(G') \cap \bigcup \mathcal{C}'_i|$ but need Y to apply weak uniformity. We first prove that $|Y| \geq \varepsilon |V(G'')|$.

$$\begin{aligned} |Y| &\geq |V(\widehat{G})| \geq \varepsilon |V(\widehat{G})| + (1 - \varepsilon) |V(\widehat{G})| \\ &\geq \varepsilon |V(\widehat{G})| + (1 - \varepsilon) N \geq \varepsilon |V(\widehat{G})| + \varepsilon |V(G)| \geq \varepsilon |V(G'')|. \end{aligned}$$

Therefore, by ε -weak k -uniformity α -approximation of A on \mathcal{E} , we get

$$|A(G'') \cap Y| \leq \alpha \cdot \text{OPT}_\Pi(G'', N^k[Y]).$$

We have $|A(G'') \cap Y| = |A(G') \cap \bigcup \mathcal{C}'_i| + |A(\widehat{G})|$ and $\text{OPT}_\Pi(G'', N^k[Y]) = \text{OPT}_\Pi(\widehat{G}) + \text{OPT}_\Pi(G', N_{G'}^k[\bigcup \mathcal{C}'_i]) \leq \omega + \text{OPT}_\Pi(G', N_{G'}^k[\bigcup \mathcal{C}'_i])$ because \widehat{G} and G' are in disjoint components of G'' . Forgetting about the term $|A(\widehat{G})|$, it follows that

$$\left| |A(G') \cap \bigcup \mathcal{C}'_i| \right| \leq \alpha \cdot \omega + \alpha \cdot \text{OPT}_\Pi(G', N_{G'}^k[\bigcup \mathcal{C}'_i]). \quad (6)$$

Since $G[N_G^r[\bigcup \mathcal{C}'_i]] = G'[N_{G'}^r[\bigcup \mathcal{C}'_i]]$, Algorithm A , which runs in r rounds, returns the same vertex set in G and G' for vertices in $\bigcup \mathcal{C}'_i$. Therefore, $|A(G) \cap \bigcup \mathcal{C}'_i| = |A(G') \cap \bigcup \mathcal{C}'_i|$. Combining with the previous inequality, we get

$$\begin{aligned} \left| |A(G) \cap \bigcup \mathcal{C}'_i| \right| &\leq \alpha \cdot \omega + \alpha \cdot \text{OPT}_\Pi(G, N_G^k[\bigcup \mathcal{C}'_i]) \leq \alpha \cdot \omega + \alpha \cdot \text{OPT}_\Pi(G, V(G)) \\ &= \alpha \cdot \omega + \alpha \cdot \text{OPT}_\Pi(G). \end{aligned}$$

We remark that $|(A(G) \cap C) \setminus X| = 0$ if $C \subseteq X$. Because all vertices in C_i but not in X are in some $C \in \mathcal{C}'_i$, we have

$$|(A(G) \cap C_i) \setminus X| \leq \left| |A(G) \cap \bigcup \mathcal{C}'_i| \right| \leq \alpha \cdot \omega + \alpha \cdot \text{OPT}_\Pi(G). \quad (7)$$

Now, we can apply the same trick as in Proposition D.2. As Π is dense in \mathcal{D} , there exists some B such that any graph of \mathcal{D} of diameter at least B has optimum value at least $\alpha(d+1)\omega/\eta$. Algorithm B checks (Step 1) whether the diameter D of its component is less than B . This can be done by collecting the radius- B neighborhood of every vertex in B rounds. If true, the vertex locally brute-forces an optimal solution for its component, which it already knows about. If false, the vertex continues the algorithm with Step 2. If $D < B$, then the algorithm is obviously a good approximation (with approximation ratio 1). Otherwise, $D \geq B$ and it follows that $\text{OPT}_\Pi(G) \geq \alpha(d+1)\omega/\eta$. Putting in Eq. (7) in Eq. (5), we get that Algorithm B is an $(\alpha(d+1) + 1 + \eta)$ -approximation, and even an $(\alpha(d+1) + \eta)$ -approximation if $X = \emptyset$.

To see that B is also a k -uniform approximation, it is sufficient to see that, for any $W \subseteq V(G)$, the bound obtained is $|B(G) \cap W| \leq \alpha(d+1)\omega + \alpha(d+1) \cdot \text{OPT}_{\Pi}(G, N_G^k[W]) + \text{OPT}_{\Pi}(G, N_G[W])$ where the first two terms are due to running A and the last term is due to running the brute-force. Indeed, if the brute-force computed a set whose intersection with W was smaller than $\text{OPT}_{\Pi}(G, N_G^{\rho}[W])$, we could replace it by the minimum size set satisfying the constraints of the vertices in $N_G^{\rho}[W]$ and obtain a smaller set, a contradiction. This proves that B is $\max\{k, \rho\}$ -uniform with approximate ratio $\alpha(d+1) + 1$ (or $\alpha(d+1)$ if $X = \emptyset$, as the brute-force is not required in that case). Now, let us prove that Algorithm B has the desired round complexity.

- Step 1 takes $B = g(\eta)$ rounds for computing if the diameter is small, and no additional round for the brute-force.
- Computing X in Step 3 takes $T + 1$ rounds.
- Running A in Step 4 takes r rounds.
- For Step 5, consider the set S computed by B at Step 4, before the brute-force. Observe that the vertices in the set $W = V(G) \setminus N_G^{\rho}[X]$ are satisfied by S . This is because vertices of W are at distance at least $\rho + 1$ from X and thus vertices of $N_G^{\rho}[W]$ cannot be in X . Thus, A applies to all vertices of $N_G^{\rho}[W]$, and so W is indeed satisfied by S (Fact C.1). Therefore, to satisfy $N_G^{\rho}[X]$ it is enough to select a set S' from $N_G^{2\rho}[X]$ (Fact C.1) as done in Step 5. By assumption, the connected components of $N_G^{2\rho}[X]$ have weak diameter in G at most δ (obviously, if $X = \emptyset$ then $\delta = 0$). Therefore, the brute-force (Step 5) will take at most $\delta + 1$ rounds.

Naively, Steps 3 and 4 together take $(T + 1) + r$ rounds. As Algorithm A is not guaranteed to work when executed on vertices of X , Step 4 must be run only after Step 3. However, both steps can be run in parallel as follows. We run A for r rounds exactly and stop it just after (since its running time on a vertex of X could result into more than r rounds). Then, the decision to add u in S (if selected by A) is delayed for the next $T + 1 - r$ rounds (this is non-negative from the choice of T). In parallel, $G[N^T[u]]$ is computed and is checked to be in \mathcal{C} or not after $T + 1$ rounds. So, after $\max\{r, T + 1\} = T + 1$ rounds, set S in Step 4 has been completed. The same trick can be applied for Step 1. Step 5 takes $\delta + 1$ steps, so that the total round complexity of B is $\max\{T + \delta + 2, g(\eta)\}$, completing the proof. \square

D Miscellaneous propositions

Proposition D.1. *Every graph class \mathcal{C} including trees has no 0-uniform α -approximations for MINIMUM DOMINATING SET, for every ratio $\alpha \geq 1$.*

PROOF. Consider a depth-2 tree T_{α} , with $\alpha + 1$ vertices at depth-1, each having $\alpha^2 + 1$ neighbors at depth-2. See Fig. 6 for an example. Clearly, $\text{MDS}(T_{\alpha}) = \alpha + 1$. Moreover, any α -approximation algorithm A in T_{α} has to select all depth-1 vertices. If not, each of the t non-selected vertex at depth 1 would force the selection of $\alpha^2 + 1$ extra vertices at depth 2, creating by this way a solution with at least $(\alpha+1-t) + t \cdot (\alpha^2+1) = \alpha \cdot (t\alpha+1) + 1$ vertices. For $t \geq 1$, this is at least $\alpha \cdot (\alpha+1) + 1 > \alpha \cdot \text{MDS}(T_{\alpha})$. However, for a subset S composed of all depth-1 vertices of T_{α} (so with $|S| = \alpha + 1$), we have on one side $|A(T_{\alpha}) \cap S| = |S| = \alpha + 1$, whereas $\text{MDS}(T_{\alpha}, N^0[S]) = \text{MDS}(T_{\alpha}, S) = 1$ (considering the root). So $|A(T_{\alpha}) \cap S| \not\leq \alpha \cdot \text{MDS}(T_{\alpha}, N^0[S])$. Therefore, A cannot be 0-uniform in T_{α} . \square




Fig. 6. The tree T_α has no 0-uniform α -approximations (here with $\alpha = 2$). Indeed, any α -approximation A of its minimum dominating set must contain the $\alpha + 1$ vertices of S (green), and thus $|A(T_\alpha) \cap S| = |S| \not\leq \alpha \cdot \text{MDS}(T_\alpha, N^0[S])$ since S can be dominated by a single vertex (the root).

Proposition D.2. *For every $\varepsilon > 0$, any r -round algorithm that returns a dominating set of size at most $\alpha \cdot \text{MDS}(G) + \beta$ on G , can be converted into an $(\alpha + \varepsilon)$ -approximation LOCAL algorithm with round complexity $r + O(\beta/\varepsilon)$.*

PROOF. Observe that the diameter of each connected component of the graph is at most $3 \cdot \text{MDS}(G)$. So, each vertex can check whether the diameter D of its component is less than B , for $B = 3\beta/\varepsilon$. This can be done by collecting its radius- B neighborhood in $B = O(\beta/\varepsilon)$ rounds. If true, the vertex locally brute forces an optimal solution for its component that, which it already knows about. If false, the vertex applies the approximate algorithm. In that case, $B \leq D$ and $D \leq 3 \text{MDS}(G)$, which together imply $\beta \leq \varepsilon \cdot \text{MDS}(G)$. Thus, the returned set has size at most $\alpha \cdot \text{MDS}(G) + \beta \leq (\alpha + \varepsilon) \cdot \text{MDS}(G)$. \square

Proposition D.3. *Every n -vertex graph G of Euler genus g satisfies $\nabla_1(G) < \sqrt{3g/2} + 3$. This bound is best possible up to an additive constant, for each $g \geq 0$.*

PROOF. By definition, we have $\nabla_1(G) = \max_H \{|E(H)|/|V(H)|\}$, where H is a depth-1 minor. For every H , we have $|E(H)|/|V(H)| < 3g/|V(H)| + 3$ from Euler's formula. As $|E(H)| \leq \binom{|V(H)|}{2}$, $\nabla_1(G) \leq \max_{m \in [1, n]} \min \{(m-1)/2, 3g/m + 3\} < \sqrt{3g/2} + 3$.

The latter inequality holds, because either $3g/m + 3 < \sqrt{3g/2} + 3$, and we are done, or $3g/m + 3 \geq \sqrt{3g/2} + 3$, which implies $m \leq \sqrt{6g}$. In that case $(m-1)/2 \leq (\sqrt{6g})/2 - 1/2 < \sqrt{3g/2} + 3$ as wanted.

Consider a graph G composed of a clique K_m with $m \geq 4$ and $m \neq 7$ where each edge is replaced by a path of length three. The graph G has same Euler genus as K_m , which is $g = \lceil (m-3)(m-4)/6 \rceil$ (cf. [MT01, Theorem 4.4.5]). Moreover, it has K_m as depth-1 minor. Thus, $\nabla_1(G) \geq |E(K_m)|/m = (m-1)/2 \geq \sqrt{3g/2} + 3/4$.

The latter inequality holds because, one can check that $(m-1)^2 \geq (m-3)(m-4) + 9$ for all $m \geq 4$. Since $(m-3)(m-4) + 9 \geq 6 \lceil (m-3)(m-4)/6 \rceil + 3 = 6g + 3$, we have $(m-1)^2/4 \geq (6g+3)/4$, and thus $(m-1)/2 \geq \sqrt{3g/2} + 3/4$ as claimed. \square