

A Quasi-Polynomial Time Partition Oracle for Graphs with an Excluded Minor

Reut Levi *

Dana Ron †

July 24, 2018

Abstract

Motivated by the problem of testing planarity and related properties, we study the problem of designing efficient *partition oracles*. A *partition oracle* is a procedure that, given access to the incidence lists representation of a bounded-degree graph $G = (V, E)$ and a parameter ϵ , when queried on a vertex $v \in V$, returns the part (subset of vertices) which v belongs to in a partition of all graph vertices. The partition should be such that all parts are small, each part is connected, and if the graph has certain properties, the total number of edges between parts is at most $\epsilon|V|$. In this work we give a partition oracle for graphs with excluded minors whose query complexity is quasi-polynomial in $1/\epsilon$, thus improving on the result of Hassidim et al. (*Proceedings of FOCS 2009*) who gave a partition oracle with query complexity exponential in $1/\epsilon$. This improvement implies corresponding improvements in the complexity of testing planarity and other properties that are characterized by excluded minors as well as sublinear-time approximation algorithms that work under the promise that the graph has an excluded minor.

1 Introduction

An important and well studied family of graphs is the family of *Planar Graphs*. A natural problem is that of deciding whether a given graph $G = (V, E)$ is planar. Indeed, there is variety of linear-time algorithms for deciding planarity (e.g. [HT74, ET76, SH99, BM04, dOR06]). However, what if one is willing to relax the decision task while requiring that the algorithm be much more efficient, and run in *sub-linear* time? Namely, here we refer to the notion of *Property Testing* where the goal is to decide (with high success probability) whether a graph has the property (planarity) or is far from having the property (in the sense that relatively many edges-modifications are required in order to obtain the property). Such a task should be performed by accessing only small portions of the input graph.

Another type of problem related to planar graphs is that of solving a certain decision, search, or optimization problem, under the *promise* that the input graph is planar, where the problem may be hard in general. In some cases the problem remains hard even under the promise (e.g., Minimum Vertex-Cover [GJS76]), while in other cases the promise can be exploited to give more efficient

*School of Computer Science, Tel Aviv University. Tel Aviv 69978, Israel. E-mail: reuti.levi@gmail.com.

†School of Electrical Engineering, Tel Aviv University. Tel Aviv 69978, Israel. E-mail: danar@eng.tau.ac.il.

algorithms than are known for general graphs (e.g., graph Isomorphism [HW74]). Here too we may seek even more efficient, sublinear-time, algorithms, which are allowed to output approximate solutions.

The problem of testing planarity, and, more generally, testing any minor-closed property of graphs¹ was first studied by Benjamini, Schramm and Shapira [BSS08]. They gave a testing algorithm whose query complexity and running time are *independent* of $|V|$.² This result was later improved (in terms of the dependence on the distance parameter, ϵ) by Hassidim et al. [HKNO09], who also considered sublinear-time approximation algorithms that work under the promise that the graph has an excluded (constant size) minor (or more generally, for hyperfinite graphs as we explain subsequently). They show how to approximate the size of the minimum vertex cover, the minimum dominating set and the maximum independent set of such graphs, to within an additive term of $\epsilon|V|$ in time that depends only on ϵ and the degree bound, d , but not on $|V|$.

The main tool introduced by Hassidim et al. [HKNO09] for performing these tasks is *Partition Oracles*. Given query access to the incidence-lists representation of a graph, a partition oracle provides access to a partition of the vertices into small connected components. A partition oracle is defined with respect to a class of graphs, \mathcal{C} , and may be randomized. If the input graph belongs to \mathcal{C} , then with high probability the partition determined by the oracle is such that the number of edges between vertices in different parts of the partition is relatively small (i.e., at most $\epsilon|V|$). Such a bound on the number of edges between parts together with the bound on the size of each part lends itself to designing efficient testing algorithms and other sublinear approximation algorithms.

Hassidim et al. [HKNO09] provide a partition oracle for hyperfinite [Ele06] classes of graphs that makes $2^{d^{\text{poly}(1/\epsilon)}}$ queries to the graph, where d is an upper bound on the degree. A graph $G = (V, E)$ is (ϵ, k) -*hyperfinite* if it is possible to remove at most $\epsilon|V|$ edges of the graph so that the remaining graph has connected components of size at most k . A graph G is ρ -*hyperfinite* for $\rho : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ if for every $\epsilon \in (0, 1]$, G is $(\epsilon, \rho(\epsilon))$ -hyperfinite. For graphs with an excluded minor, (a special case of hyperfinite graphs), they provide a partition oracle with query complexity $d^{\text{poly}(1/\epsilon)}$ (as detailed in [Ona10, Sec. 2]). Hassidim et al. [HKNO09] leave as an open problem whether it is possible to design a partition oracle for graphs with an excluded minor that has query complexity polynomial in $1/\epsilon$. In particular, this would imply an algorithm for testing planarity whose complexity is polynomial in $1/\epsilon$.

1.1 Our Contribution

In this work we present a partition oracle for graphs with an excluded minor whose query complexity and running time are $(d/\epsilon)^{O(\log(1/\epsilon))} = d^{O(\log^2(1/\epsilon))}$, that is, quasi-polynomial in $1/\epsilon$.

IMPLICATIONS. Hassidim et al. [HKNO09] show how it is possible to reduce the problem of testing H -minor freeness (for a fixed graph H) to the problem of designing a partition oracle for H -minor free graphs. Using this reduction they obtain a testing algorithm for H -minor freeness (and more

¹For a fixed graph H , H is a *minor* of G if H is isomorphic to a graph that can be obtained by zero or more edge contractions on a subgraph of G . We say that a graph G is *H -minor free* (or *excludes H as a minor*) if H is not a minor of G . A property \mathcal{P} (class of graphs) is *minor-closed* if every minor of a graph in \mathcal{P} is also in \mathcal{P} . Any minor-closed property can be characterized by a finite family of excluded minors [RS04].

²Their algorithm has two-sided error. If one-sided error is desired, then for any fixed H that contains a simple cycle, the query complexity of one-sided error testing of H -minor freeness is $\Omega(\sqrt{|V|})$ [CGR⁺12]. On the positive side, if H is cycle-free, then there is a one-sided error algorithm whose complexity does not depend on $|V|$ [CGR⁺12].

generally, for any minor-closed property, since such properties can be characterized by a finite family of excluded minors [RS04]) whose query complexity and running time are $2^{\text{poly}(1/\epsilon)}$. As noted previously, this improves on the testing algorithm of Benjamini et al. [BSS08] for minor-closed properties, whose complexity is $2^{2^{\text{poly}(1/\epsilon)}}$. Using our partition oracle (and the reduction in [HKNO09]) we get a testing algorithm whose complexity is $2^{O(\log^2(1/\epsilon))}$.

Other applications of a partition oracle for a class of graphs \mathcal{C} are constant time algorithms that work under the promise that the input graph belongs to \mathcal{C} , where in our case \mathcal{C} is any class of graphs with an excluded minor. Under this promise, Hassidim et al. [HKNO09] provide constant time $\epsilon|V|$ -additive-approximation algorithms for the size of a minimum vertex cover, minimum dominating set and maximum independent set. They also obtain an ϵ -additive-approximation algorithm for the distance from not having any graph in \mathcal{H} as an induced subgraph where \mathcal{H} is a fixed subset of graphs. Combined with our partition oracle, the query complexity of these algorithms drops from $d^{\text{poly}(1/\epsilon)}$ to $(d/\epsilon)^{O(\log(1/\epsilon))} = d^{O(\log^2(1/\epsilon))}$.

TECHNIQUES. As in [HKNO09], our partition oracle runs a local emulation of a global partitioning algorithm. Hence, we first give a high-level idea of the global partitioning algorithm, and then discuss the local emulation. Our global partitioning algorithm is based on the global partitioning algorithm of [HKNO09] for graphs with an excluded minor, as described in [Ona10, Sec. 2], which in turn builds on a clustering method of Czygrinow, Hańćkowiak, and Wawrzyniak [CHW08]. The algorithm is also similar to the “Binary Borůvka” algorithm [PR08] for finding a minimum-weight spanning tree. The global algorithm works iteratively, coarsening the partition in each iteration. Initially each vertex is in its own part of the partition, and in each iteration some subsets of parts are merged into larger (connected) parts. The decisions regarding these merges are based on the numbers of edges between parts, as well as on certain random choices. Applying the analysis in [Ona10] it is possible to show that with high constant probability, after $O(\log(1/\epsilon))$ iterations, the number of edges between parts is at most $\epsilon|V|$, as required. Since the sizes of the parts obtained after the last merging step may be much larger than desired, as a final step it is possible to refine the partition without increasing the number of edges crossing between parts by too much by applying an algorithm of Alon Seymour and Thomas [AST90].

When turning to the local emulation of the global algorithm by the partition oracle, the query complexity and running time of the partition oracle depend on the sizes of the parts in the *intermediate* stages of the algorithm. These sizes are bounded as a function of $1/\epsilon$ and d , but can still be quite large. As an end-result, the partition oracle described in [Ona10] has complexity that grows exponentially with $\text{poly}(1/\epsilon)$. To reduce this complexity, we modify the global partition algorithm as follows: If (following a merging stage) the size of a part goes above a certain threshold, we ‘break’ it into smaller parts. To this end we apply to each large part (in each iteration) the abovementioned algorithm of Alon Seymour and Thomas [AST90]. This algorithm finds (in graphs with an excluded minor) a relatively small vertex separator whose removal creates small connected components. Each such refinement of the partition increases the number of edges crossing between parts. However, we set the parameters for the algorithm in [AST90] so that the decrease in the number of edges between parts due to the merging steps dominates the increase due to the ‘breaking’ steps. One could have hoped that since the sizes of the parts are now always bounded by some polynomial in $1/\epsilon$ (and d), the complexity of the partition oracle will be $\text{poly}(d/\epsilon)$ as well. However, this is not the case, since in order to determine the part that a vertex, v , belongs to after a certain iteration, it is necessary to determine the parts that other vertices in the local neighborhood of v belong to

in previous iterations. This leads to a recursion formula whose solution is quasi-polynomial in $1/\epsilon$.

1.2 Other Related Work

Yoshida and Ito [YI10] were the first to provide a testing algorithm for a minor-closed property whose complexity is polynomial in $1/\epsilon$ and d . Specifically, they give a testing algorithm for the property of being outerplanar. Their result was generalized by Edelman et al. [EHNO11] who design a partition oracle with complexity $\text{poly}(d/\epsilon)$ for the class of bounded treewidth graphs. Known families of graphs with bounded treewidth include cactus graphs, outerplanar graphs and series-parallel graphs. However, many graphs with an excluded minor do not have bounded treewidth. For example, planar graphs are known to have treewidth of $\Omega(\sqrt{n})$.

Building on the partition oracle of [HKNO09], Newman and Sohler [NS11] design an algorithm for testing any property of graphs under the promise that the input graph is taken from \mathcal{C} , for any \mathcal{C} that is a ρ -hyperfinite family of graphs. The number of queries their algorithm makes to the graph is independent of $|V|$ but is at least exponential in $1/\epsilon$. In a recent work, Onak [Ona12] proves that there exists a property such that testing this property requires performing $2^{\Omega(1/\epsilon)}$ queries even under the promise that the input graph is taken from a hyperfinite family of graphs. This family of graphs \mathcal{T} consists of graphs that are unions of bounded degree trees. Onak [Ona12] defines a subclass of \mathcal{T} and shows that every algorithm for testing the property of membership in this subclass must perform $2^{\Omega(1/\epsilon)}$ queries to the graph.

Czumaj, Shapira, and Sohler [CSS09] investigated another promise problem. They proved that any hereditary property, namely a property that is closed under vertex removal, can be tested in time independent of the input size if the input graph belongs to a hereditary and non-expanding family of graphs.

As for approximation under a promise, Elek [Ele10] proved that under the promise that the input graph, $G = (V, E)$, has sub-exponential growth and bounded degree, the size of the minimum vertex cover, the minimum dominating set and the maximum independent set, can be approximated up to an $\epsilon|V|$ -additive error with time complexity that is independent of the graph size. Newman and Sohler [NS11] showed how to obtain an $\epsilon|V|$ -additive approximation for a large class of graph parameters.

2 Preliminaries

In this section we introduce several definitions and some known results that will be used in the following sections. Unless stated explicitly otherwise, we consider simple graphs, that is, with no self-loops and no parallel edges. The graphs we consider have a known degree bound d , and we assume we have query access to their incidence-lists representation. Namely, for any vertex v and index $1 \leq i \leq d$ it is possible to obtain the i^{th} neighbor of v (where if v has less than i neighbors, then a special symbol is returned). If the graph is edge-weighted, then the weight of the edge is returned as well.

For a graph $G = (V, E)$ and two sets of vertices $V_1, V_2 \subseteq V$, we let $E(V_1, V_2)$ denote the set of edges in G with one endpoint in V_1 and one endpoint in V_2 . That is $E(V_1, V_2) \stackrel{\text{def}}{=} \{(v_1, v_2) \in E : v_1 \in V_1, v_2 \in V_2\}$.

Definition 1 Let $G = (V, E, w)$ be an edge-weighted graph and let $\mathcal{P} = (V_1, \dots, V_t)$ be a partition of the vertices of G such that for every $1 \leq i \leq t$, the subgraph induced by V_i is connected. Define the contraction G/\mathcal{P} of G with respect to the partition \mathcal{P} to be the edge-weighted graph $G' = (V', E', w')$ where:

1. $V' = \{V_1, \dots, V_t\}$ (that is, there is a vertex in V' for each subset of the partition \mathcal{P});
2. $(V_i, V_j) \in E'$ if and only if $i \neq j$ and $E(V_i, V_j) \neq \emptyset$;
3. $w'((V_i, V_j)) = \sum_{(u,v) \in E(V_i, V_j)} w((u, v))$.

As a special case of Definition 1 we get the standard notion of a single-edge contraction.

Definition 2 Let $G = (V, E, w)$ be an edge-weighted graph on n vertices v_1, \dots, v_n , and let (v_i, v_j) be an edge of G . The graph obtained from G by contracting the edge (v_i, v_j) is G/\mathcal{P} where \mathcal{P} is the partition of V into $\{v_i, v_j\}$ and singletons $\{v_k\}$ for every $k \neq i, j$.

Definition 3 For $\epsilon \in (0, 1]$, $k \geq 1$ and a graph $G = (V, E)$, we say that a partition $\mathcal{P} = (V_1, \dots, V_t)$ of V is an (ϵ, k) -partition (with respect to G), if the following conditions hold:

1. For every $1 \leq i \leq t$ it holds that $|V_i| \leq k$;
2. For every $1 \leq i \leq t$ the subgraph induced by V_i in G is connected;
3. The total number of edges whose endpoints are in different parts of the partition is at most $\epsilon|V|$ (that is, $|\{(v_i, v_j) \in E : v_i \in V_j, v_j \in V_j, i \neq j\}| \leq \epsilon|V|$).

Let $G = (V, E)$ be a graph and let \mathcal{P} be a partition of V . We denote by $g_{\mathcal{P}}$ the function from $v \in V$ to 2^V (the set of all subsets of V), that on input $v \in V$, returns the subset $V_{\ell} \in \mathcal{P}$ such that $v \in V_{\ell}$.

Definition 4 ([HKNO09]) An oracle \mathcal{O} is a partition oracle if, given query access to the incidence-lists representation of a graph $G = (V, E)$, the oracle \mathcal{O} provides query access to a partition $\mathcal{P} = (V_1, \dots, V_t)$ of V , where \mathcal{P} is determined by G and the internal randomness of the oracle. Namely, on input $v \in V$, the oracle returns $g_{\mathcal{P}}(v)$ and for any sequence of queries, \mathcal{O} answers consistently with the same³ \mathcal{P} . An oracle \mathcal{O} is an (ϵ, k) -partition oracle with respect to a class of graphs \mathcal{C} if the partition \mathcal{P} it answers according to has the following properties.

1. For every $V_{\ell} \in \mathcal{P}$, $|V_{\ell}| \leq k$ and the subgraph induced by V_{ℓ} in G is connected.
2. If G belongs to \mathcal{C} , then $|\{(u, v) \in E : g_{\mathcal{P}}(v) \neq g_{\mathcal{P}}(u)\}| \leq \epsilon|V|$ with high constant probability, where the probability is taken over the internal coin flips of \mathcal{O} .

³While the partition \mathcal{P} does not depend on the sequence of queries to the oracle, the oracle may keep in its memory the identity of the vertices it was previously queried on as well as any additional information it has acquired in previous queries and the outcome of coins it has flipped.

By the above definition, if $G \in \mathcal{C}$, then with high constant probability the partition \mathcal{P} is an (ϵ, k) -partition, while if $G \notin \mathcal{C}$ then it is only required that each part of the partition is connected and has size at most k . We are interested in partition oracles that have small query complexity, namely, that perform few queries to the graph (for each vertex they are queried on).

Recall that a graph H is called a *minor* of a graph G if H is isomorphic to a graph that can be obtained by zero or more edge contractions on a subgraph of G . A graph G is *H -minor free* if H is not a minor of G . We next quote two results that will play a central role in this work.

Fact 1 ([Ona10]) *Let H be a fixed graph with $c_1(H)$ edges. For every H -minor free graph $G = (V, E)$ it holds that:*

- $|E| \leq c_1(H) \cdot |V|$;
- E can be partitioned into at most $c_1(H)$ forests.

Proposition 1 ([AST90]) *Let $G = (V, E)$ be a K_h -minor free graph where the vertices of G are associated with non-negative weights that sum to 1. There exists a constant c such that for any $\beta \in (0, 1]$, there is a set of $ch^{3/2}|V|^{1/2}/\beta^{1/2}$ vertices of G whose removal leaves G with no connected component having weight greater than β . Such a set can be found in time $O(h^{1/2}|V|^{1/2}|E|)$.*

As a corollary we get:

Corollary 2 *Let H be a fixed graph. There is a constant $c_2(H) > 1$ such that for every $\gamma \in (0, 1]$, every H -minor free graph $G = (V, E)$ with degree bounded by d is $(\gamma, c_2(H)d^2/\gamma^2)$ -hyperfinite. Furthermore, a $(\gamma, c_2(H)d^2/\gamma^2)$ partition of V can be found in time $O(|V|^{3/2})$.*

Proof: We apply Proposition 1 with equal weights $1/|V|$ to all vertices and with $\beta = (c^2h^3d^2)/(\gamma^2|V|)$. By Proposition 1, we obtain a set S of vertices such that $|S| = (\gamma/d)|V|$ and such that the removal of S leaves G with connected components of size at most $c^2h^3d^2/\gamma^2 = c_2(H)d^2/\gamma^2$ each. Consider the partition \mathcal{P} of V that consists of a singleton subset for each vertex in S and a subset for each component (containing the vertices in the component). Since the number of edges incident to S is at most $\gamma|V|$, the partition \mathcal{P} is a $(\gamma, c_2(H)d^2/\gamma^2)$ partition, and since $|E| = O(|V|)$ (as the graph is H -minor free for a fixed H), the running time for finding \mathcal{P} is as stated. ■

3 A Global Partitioning Algorithm

Our partition oracle is local, in the sense that its output is determined by the local neighborhood of the vertex it is queried on. However, as in previous work, the oracle is based on a global partitioning algorithm, which accesses the whole graph, and the oracle emulates this algorithm locally. In this section we describe this global partition algorithm. As noted in the introduction, our algorithm and its analysis are based on [Ona10] (which in turn builds on a clustering method of Czygrinow, Hańćkowiak, and Wawrzyniak [CHW08], and is also similar to the “Binary Borůvka” algorithm [PR08] for finding a minimum-weight spanning tree).

The algorithm proceeds in iterations, where in iteration i it considers a graph G^{i-1} , where G^{i-1} is edge-weighted. The vertices of G^{i-1} correspond to (disjoint) subsets of vertices that induce connected subgraphs in G , and the weight of an edge between two vertices in G^{i-1} is the number of edges in G between the two corresponding subsets of vertices. Initially, the underlying graph

G^0 is G and all edges have weight 1. In each iteration the algorithm contracts a subset of the edges so that each vertex in G^i corresponds to a subset that is the union of subsets of vertices that correspond to vertices in G^{i-1} . When the algorithm terminates it outputs the partition into subsets that correspond to the vertices of the final graph.

Each iteration of the algorithm consists of two phases. In the first phase of iteration i , a subset of the edges of G^{i-1} are contracted, resulting in a graph \tilde{G}^i . In the second phase, some of the subsets that correspond to vertices in \tilde{G}^i remain as is, and some are ‘broken’ into smaller subsets. The vertices of G^i correspond to these subsets (both ‘broken’ and ‘unbroken’). Observe that if the graph G is H -minor free for some fixed graph H , then every G^i and \tilde{G}^i is H -minor free as well.

In the first phase of iteration i , the contracted edges are selected randomly as follows. Each vertex in G^{i-1} selects an incident edge with maximum weight and tosses a fair coin to be ‘Heads’ or ‘Tails’. Each selected edge is contracted if and only if it is selected by a ‘Heads’ vertex and its other endpoint is a ‘Tails’ vertex. This way, in each iteration, the contracted edges form stars (depth-1 trees). Therefore, a vertex in the graph \tilde{G}^i that results from the contraction of edges in G^{i-1} as described above, corresponds to a subset of vertices in V that induces a connected subgraph in G , and $\tilde{G}^i = G/\tilde{\mathcal{P}}^i$ (recall Definition 1) where $\tilde{\mathcal{P}}^i$ is this partition into subsets. Since each vertex in each \tilde{G}^i corresponds to a connected subgraph in G , we shall refer to the vertices of \tilde{G}^i as *connected components* (to be precise, they are connected components in the graph resulting from removing all edges in G the correspond to (weighted) edges in \tilde{G}^i). In each iteration, following the contraction of edges, if the size of a component goes above a certain threshold,⁴ $k = \text{poly}(d/\epsilon)$, then the component is ‘broken’ into smaller connected components, each of size at most k . This is done using the algorithm referred to in Corollary 2, and G^i is the (edge-weighted) graph whose vertices correspond to the new components.

Theorem 3 *Let H be a fixed graph. If the input graph G is H -minor free and has degree bounded by d , then for any given $\epsilon \in (0, 1]$, Algorithm 1 outputs an $(\epsilon, O(d^2/\epsilon^2))$ -partition of G with high constant probability.*

Proof: We first claim that in each iteration, after Step 2b, the total weight of the edges in the graph is decreased by a factor of $\left(1 - \frac{1}{8c_1(H)}\right)$ with probability at least $\frac{1}{8c_1(H)-1}$, where the probability is taken over the coin tosses of the algorithm. Fixing an iteration i , by Fact 1 we know that the edges of G^{i-1} can be partitioned into at most $c_1(H)$ forests. It follows that one of these forests contains edges with total weight at least $\frac{w(G^{i-1})}{c_1(H)}$ where $w(G^{i-1})$ denotes the total weight of the edges in G^{i-1} . Suppose we orient the edges of the forest from roots to leaves, so that each vertex in the forest has in-degree at most 1. Recall that for each vertex v , the edge selected by v in Step 2b is the heaviest among its incident edges. It follows that the expected total weight of edges contracted in Step 2b is at least $\frac{w(G^{i-1})}{c_1(H)}$ (recall that an edge (v, u) selected by v is contracted if the coin flip of v is ‘Heads’ and that of u is ‘Tails’, an event that occurs with probability $1/4$). Thus, the expected total weight of edges that are *not* contracted is at most $w(G^{i-1}) - \frac{w(G^{i-1})}{4c_1(H)} = \left(1 - \frac{1}{4c_1(H)}\right) w(G^{i-1})$. By Markov’s inequality the probability that the total weight of edges that are not contracted is at

⁴One may consider setting different thresholds for different iterations. However, in our analysis this does not seem to give a better bound on the complexity as compared to setting a common threshold.

Algorithm 1: A global $(\epsilon, c_2(H)d^2/\epsilon^2)$ -partition algorithm for an H -minor free graph $G = (V, E)$

1. Set $G^0 := G$
 2. For $i = 1$ to $\ell = \Theta(\log 1/\epsilon)$:
 - (a) Toss a fair coin for every vertex in G^{i-1} .
 - (b) For each vertex u let (u, v) be an edge with maximum weight that is incident to u (where ties are broken arbitrarily). If u 's coin toss is 'Heads' and v 's coin toss is 'Tails', then contract (u, v) .
 - (c) Let $\tilde{G}^i = (\tilde{V}^i, \tilde{E}^i, \tilde{w}^i)$ denote the graph resulting from the contraction of the edges as determined in the previous step. Hence, each vertex $\tilde{v}_j^i \in \tilde{V}^i$ corresponds to a subset of vertices in G , which we denote by \tilde{C}_j^i .
 - (d) Let $\gamma = \epsilon/(3\ell)$. For each \tilde{C}_j^i such that $|\tilde{C}_j^i| > c_2(H)/\gamma^2$, partition the vertices in \tilde{C}_j^i into connected subsets of size at most $k = c_2(H)d^2/\gamma^2$ each by running the algorithm referred to in Corollary 2 on the subgraph induced by \tilde{C}_j^i in G .
 - (e) Set $G^i := G/\mathcal{P}^i$, where \mathcal{P}^i is the partition resulting from the previous step.
 3. For each subset C_j^ℓ in \mathcal{P}^ℓ such that $|C_j^\ell| > c_2(H)d^2/\epsilon^2$, partition the vertices in C_j^ℓ into connected subsets each of size at most $3c_2(H)d^2/\epsilon^2$ by running the algorithm referred to in Corollary 2, and output the resulting partition.
-

least $\left(1 - \frac{1}{8c_1(H)}\right) w(G^{i-1})$ is at most

$$\frac{\left(1 - \frac{1}{4c_1(H)}\right) w(G^{i-1})}{\left(1 - \frac{1}{8c_1(H)}\right) w(G^{i-1})} = 1 - \frac{1}{8c_1(H) - 1}. \quad (1)$$

We say that an iteration i is *successful* if $w(\tilde{G}^i) \leq \left(1 - \frac{1}{8c_1(H)}\right) w(G^{i-1})$. By Equation 1, the probability of success is at least $\eta = \frac{1}{8c_1(H)-1}$, for every iteration. Let Y_i be the random variable that takes the value $-1/\eta + 1$ if the i -th iteration is successful and takes the value 1 otherwise. Let $Z_i = \sum_{j=1}^i Y_j$, then $E(Z_{k+1}|Z_1, \dots, Z_k) \leq Z_k + 1 \cdot (1 - \eta) + (1 - 1/\eta) \cdot \eta = Z_k$ and $|Z_k - Z_{k-1}| \leq \frac{1}{\eta}$ for every k . By Azuma's inequality [Azu67] we obtain that $\Pr(Z_\ell \geq t) \leq e^{-t^2/(2\ell(1/\eta)^2)}$ for every t . Setting $t = 3\sqrt{\ell}/\eta$, we obtain that with probability greater than 9/10, $Z_\ell < 3\sqrt{\ell}/\eta$. Let s denote the number of successful iterations, then $Z_\ell = (\ell - s) + s(1 - 1/\eta) = \ell - s/\eta$. Thus, we obtain that $s > \eta\ell - 3\sqrt{\ell}$. We conclude that with probability at least 9/10, the number of successful iterations is at least $\frac{\ell}{16c_1(H)-2}$.

Our second claim is that for every i , after Step 2d, it holds that $w(G^i) \leq w(\tilde{G}^i) + \frac{\epsilon n}{3\ell}$ (where $n = |V|$). This follows from Corollary 2: For each component \tilde{C}_j^i that we break, we increase the total weight of edges between components by an additive term of at most $\gamma|\tilde{C}_j^i| = \frac{\epsilon|\tilde{C}_j^i|}{3\ell}$. Thus, after

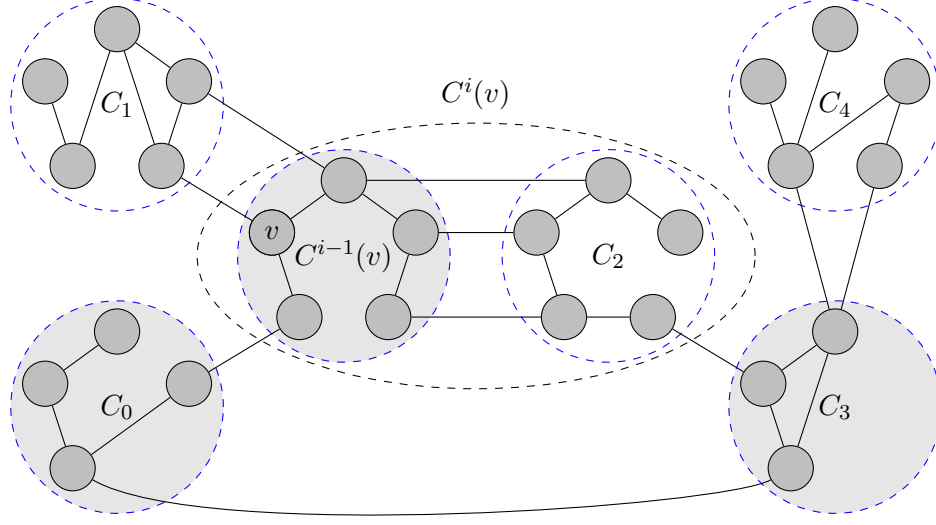


Figure 1: The dark components are ‘Heads’ vertices and the bright ones are ‘Tails’ vertices. Before the i^{th} iteration there are 3 edges in the cut between $C^{i-1}(v)$ and C_2 and 2 edges in the cut between $C^{i-1}(v)$ and C_1 . Since $C^{i-1}(v)$ is a ‘Heads’ vertex it merges with C_2 which is a ‘Tails’ vertex. C_3 does not merge with C_2 since the cut (C_3, C_4) is larger than the cut (C_3, C_2) . C_0 does not merge with any vertex since it is a ‘Tails’ vertex that is not connected to a ‘Heads’ vertex. This illustration is for a case that $C^i(v) = \tilde{C}^i(v)$.

Step 2d of the ℓ^{th} iteration, with probability at least 9/10, the weight of the edges in G^ℓ is at most

$$c_1(H) \cdot n \cdot \left(1 - \frac{1}{8c_1(H)}\right)^{\frac{\ell}{16c_1(H)-2}} + \ell \cdot \frac{\epsilon n}{3\ell} \leq 2\epsilon n/3. \quad (2)$$

Since we add at most $\epsilon n/3$ weight in Step 3 (when breaking the subsets corresponding to vertices in G^ℓ into subsets of size at most $3c_2(H)d^2/\epsilon^2$), we obtain the desired result. ■

4 The Partition Oracle

In this section we describe how, given query access to the incidence-lists representation of a graph $G = (V, E)$ and a vertex $v \in V$, it is possible to emulate Algorithm 1 locally and determine the part that v belongs to in the partition \mathcal{P} that the algorithm outputs. Namely, we prove the following theorem.

Theorem 4 *For any fixed graph H there exists an $(\epsilon, O(d^2/\epsilon^2))$ -partition-oracle for H -minor free graphs that makes $(d/\epsilon)^{O(\log(1/\epsilon))}$ queries to the graph for each query to the oracle. The total time complexity of a sequence of q queries to the oracle is $q \log q \cdot (d/\epsilon)^{O(\log(1/\epsilon))}$.*

Proof: Recall that the partition \mathcal{P} is determined randomly based on the ‘Heads’/‘Tails’ coin-flips of the vertices in each iteration. Also recall that $g_{\mathcal{P}}(v)$ denotes the subset of vertices that v belongs

to in \mathcal{P} . Since we want the oracle to be efficient, the oracle will flip coins “on the fly” as required for determining $g_{\mathcal{P}}(v)$. Since the oracle has to be consistent with the same \mathcal{P} for any sequence of queries it gets, it will keep in its memory all the outcomes of the coin flips it has made. For the sake of simplicity, whenever an outcome of a coin is required, we shall say that a coin is flipped, without explicitly stating that first the oracle checks whether the outcome of this coin flip has already been determined. We shall also explain subsequently how to break ties (deterministically) in the choice of a heaviest incident edge in each iteration of the algorithm.

Recall that the algorithm constructs a sequence of graphs $G^0 = G, \tilde{G}^1, G^1, \dots, \tilde{G}^\ell, G^\ell$, and that for each $0 \leq i \leq \ell$, the vertices in G^i correspond to connected subgraphs of G (which we refer to as components). For a vertex $v \in V$ let $C^i(v)$ denote the vertex/component that v belongs to in G^i , and define $\tilde{C}^i(v)$ analogously with respect to \tilde{G}^i . Indeed, we shall refer to vertices in G^i (\tilde{G}^i) and to the components that correspond to them, interchangeably. When the algorithm flips a coin for a vertex C in G^i , we may think of the coin flip as being associated with the vertex having the largest id (according to some arbitrary ordering) in the corresponding component in G . When the algorithm selects a heaviest edge incident to C and there are several edges $(C, C_1), \dots, (C, C_r)$ with the same maximum weight, it breaks ties by selecting the edge (C, C_j) for which C_j contains the vertex with the largest id (according to the same abovementioned arbitrary ordering). We can then refer to (C, C_j) as *the* heaviest edge incident to C . In particular, since in G^0 all edges have the same weight, the heaviest edge incident to a vertex u in G^0 is the edge (u, y) for which y is maximized.

Let $Q^i(v)$ denote the number of queries to G that are performed in order to determine $C^i(v)$, and let \bar{Q}^i denote an upper bound on $Q^i(v)$ that holds for any vertex v . We first observe that $\bar{Q}^1 \leq d^2$. In order to determine $C^1(v)$, the oracle first flips a coin for v . If the outcome is ‘Tails’ then the oracle queries the neighbors of v . For each neighbor u of v it determines whether (u, v) is the heaviest edge incident to u (by querying all of u ’s neighbors). If so, it flips a coin for u , and if the outcome is ‘Heads’, then the edge is contracted (implying that $u \in \tilde{C}^1(v)$). If v is a ‘Heads’ vertex, then it finds its heaviest incident edge, (v, u) by querying all of v ’s neighbors. If u is a ‘Tails’ vertex (so that (v, u) is contracted), then the oracle queries all of u neighbors, and for each neighbor it queries all of its neighbors. By doing so (and flipping all necessary coins) it can determine which additional edges (u, y) incident to u are contracted (implying for each that $y \in \tilde{C}^1(v) = \tilde{C}^1(u)$). In both cases (of the outcome of v ’s coin flip), the number of queries performed to G is at most d^2 . Recall that a component as constructed above is ‘broken’ if it contains more than $k = \tilde{O}(d^2/\epsilon^2)$ vertices. Since $|\tilde{C}^1(v)| \leq d + 1$ for every v , we have that $C^1(v) = \tilde{C}^1(v)$.

For general $i > 1$, to determine the connected component that a vertex v belongs to after iteration i , we do the following. First we determine the component it belongs to after iteration $i-1$, namely $C^{i-1}(v)$, at a cost of at most \bar{Q}^{i-1} queries. Note that by the definition of the algorithm, $|C^{i-1}(v)| \leq k$. We now have two cases (for an illustration see Figure 1):

Case 1: $C^{i-1}(v)$ is a ‘Tails’ vertex for iteration i . In this case we query all edges incident to vertices in $C^{i-1}(v)$, which amounts to at most $d \cdot k$ edges. For each endpoint u of such an edge we find $C^{i-1}(u)$. For each $C^{i-1}(u)$ that is ‘Heads’ we determine whether its heaviest incident edge connects to $C^{i-1}(v)$ and if so the edge is contracted (so that $C^{i-1}(u) \subset \tilde{C}^i(v)$). To do so, we need, again, to query all the edges incident to vertices in $C^{i-1}(u)$, and for each endpoint y of such an edge we need to find $C^{i-1}(y)$. The weight of each edge $(C^{i-1}(u), C^{i-1}(y))$ is $|E(C^{i-1}(u), C^{i-1}(y))|$ (and since all edges incident to vertices in $C^{i-1}(y)$ have been queried, this weight is determined). The total number of vertices x for which we need to find $C^{i-1}(x)$ is upper bounded by $d^2 k^2$, and this

is also an upper bound on the number of queries performed in order to determine the identity of these vertices.

Case 2: $C^{i-1}(v)$ is a ‘Heads’ vertex in iteration i . In this case we find its heaviest incident edge in G^{i-1} , as previously described for $C^{i-1}(u)$. Let C' denote the other endpoint in G^{i-1} . If C' is a ‘Tails’ vertex then we apply the same procedure to C' as described in Case 1 for $C^{i-1}(v)$ (that is, in the case that $C^{i-1}(v)$ is a ‘Tails’ vertex in G^{i-1}). The bound on the number of queries performed is also as in Case 1. In either of the two cases we might need to ‘break’ $\tilde{C}^i(v)$ (in case $|\tilde{C}^i(v)| > k$) so as to obtain $C^i(v)$. However, this does not require performing any additional queries to G since all edges between vertices in $\tilde{C}^i(v)$ are known, and this step only contributes to the running time of the partition oracle. We thus get the following recurrence relation for Q^i : $Q^i = d^2 \cdot k^2 + d^2 \cdot k^2 \cdot Q^{i-1}$. Since $k = \text{poly}(d/\epsilon)$ we get that

$$Q^\ell \leq (d \cdot \text{poly}(d/\epsilon))^{2\ell} = (d/\epsilon)^{O(\log(1/\epsilon))}, \quad (3)$$

as claimed.

Finally, we turn to the running time. Let $T^i(v)$ denote the running time for determining $C^i(v)$. By the same reasoning as above we have that $T^i \leq O(d^2 \cdot k^2) \cdot T^{i-1} + B$ where B is an upper bound on the running time of breaking a connected component at each iteration. From Corollary 2 we obtain that $B \leq (d \cdot k^2)^{3/2}$. Thus, the running time of the oracle is $(d/\epsilon)^{O(\log(1/\epsilon))}$ for a single query. As explained above, for the sake of consistency, the oracle stores its previous coin-flips. By using a balanced search tree to store the coin flips we obtain that the total running time of the oracle for a sequence of q queries is $q \log q \cdot (d/\epsilon)^{O(\log(1/\epsilon))}$, as claimed. ■

5 Applications

In this section we state the improved complexity for the applications, of the partition oracle, which are presented in [HKNO09]. We obtain an improvement either in the query complexity or in the time complexity for all their applications excluding the application of approximating the distance to hereditary properties in which case the improvement we obtain is not asymptotic.

- Hassidim et al. [HKNO09] show that for any fixed graph H there is a testing algorithm for the property of being H -minor free in the bounded-degree model that performs $O(1/\epsilon^2)$ queries to \mathcal{O} , where \mathcal{O} is an $(\epsilon d/4, k)$ -partitioning oracle for the class of H -minor free graphs with degree bounded by d , and has $O(dk/\epsilon + k^3/\epsilon^6)$ time complexity. By using the partition oracle from Theorem 4 we obtain that the query and time complexity of testing H -minor freeness (in the bounded-degree model) is improved from $d^{\text{poly}(1/\epsilon)}$ to $(d/\epsilon)^{O(\log 1/\epsilon)}$.
- Let \mathcal{P} be a minor-closed property. According to [RS95], \mathcal{P} can be characterized as a finite set of excluded minors. Let S denote this set. By taking the proximity parameter to be $\epsilon/|S|$ and applying the testing algorithm for minor-freeness on every minor in S we obtain that the query and time complexity of testing a minor-closed property in the bounded degree model is improved from $2^{\text{poly}(|S|/\epsilon)}$ to $(|S|/\epsilon)^{O(\log |S|/\epsilon)}$. In particular this implies a testing algorithm for planarity with complexity $(1/\epsilon)^{O(\log(1/\epsilon))}$.

The next approximation algorithms work under the promise that the input graph is a graph with an excluded minor (of constant size). Under this promise we obtain the following improvements in the query complexity while the time complexity remains unchanged (the former time complexity dominates the improvement in the time complexity of the partition oracle):

- Hassidim et al. [HKNO09] provide a constant time $\epsilon|V|$ -additive-approximation algorithm for minimum vertex cover size, maximum independent set size, and the minimum dominating set size for any family of graphs with an efficient partition oracle. The algorithm makes $O(1/\epsilon^2)$ queries to the partition oracle. By using the partition oracle from Theorem 4, the query complexity of the approximation algorithms is improved from $d^{\text{poly}(1/\epsilon)}$ to $(d/\epsilon)^{O(\log 1/\epsilon)}$.
- By Lemma 11 in [HKNO09], for any finite set of connected graphs \mathcal{H} , there is an ϵ -additive-approximation algorithm for the distance to the property of not having any graph in \mathcal{H} as an induced subgraph, which makes $O(1/\epsilon^2)$ queries to the partition oracle. Hence, the query complexity of the algorithm is improved from $d^{\text{poly}(1/\epsilon)}$ to $(d/\epsilon)^{O(\log 1/\epsilon)}$.

References

- [AST90] N. Alon, P. D. Seymour, and R. Thomas. A separator theorem for graphs with an excluded minor and its applications. In *Proceedings of STOC*, pages 293–299, 1990.
- [Azu67] K. Azuma. Weighted sums of certain dependent variables. *Tôhoku Math.*, 19(3):357–367, 1967.
- [BM04] J. M. Boyer and W. H. Myrvold. On the cutting edge: simplified $O(n)$ planarity by edge addition. *JGAA*, 8(3):241–273, 2004.
- [BSS08] I. Benjamini, O. Schramm, and A. Shapira. Every minor-closed property of sparse graphs is testable. In *Proceedings of STOC*, pages 393–402, 2008.
- [CGR⁺12] A. Czumaj, O. Goldreich, D. Ron, C. Seshahadri, A. Shapira, and C. Sholer. Finding cycles and trees in sublinear time. To appear in *RSA*, also available as <http://arxiv.org/abs/1007.4230>, 2012.
- [CHW08] A. Czygrinow, M. Hańćkowiak, and W. Wawrzyniak. Fast distributed approximations in planar graphs. In *Proceedings of DISC*, pages 78–92, 2008.
- [CSS09] A. Czumaj, A. Shapira, and C. Sohler. Testing hereditary properties of nonexpanding bounded-degree graphs. *SICOMP*, 38(6):2499–2510, 2009.
- [dOR06] H. de Fraysseix, P. Ossona de Mendez, and P. Rosenstiehl. Trémaux trees and planarity. *IJFCS*, 17(5):1017–1030, 2006.
- [EHNO11] A. Edelman, A. Hassidim, H. N. Nguyen, and K. Onak. An efficient partitioning oracle for bounded-treewidth graphs. In *Proceedings of RANDOM*, pages 530–541, 2011.
- [Ele06] G. Elek. The combinatorial cost. Technical Report math/0608474, ArXiv, 2006.
- [Ele10] G. Elek. Parameter testing in bounded degree graphs of subexponential growth. *RSA*, 37(2):248–270, 2010.
- [ET76] S. Even and R. E. Tarjan. Computing an st -numbering. *TCS*, 2(3):339–344, 1976.
- [GJS76] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified NP-complete graphs problems. *TCS*, 1(3):237–267, 1976.

- [HKNO09] A. Hassidim, J. A. Kelner, H. N. Nguyen, and K. Onak. Local graph partitions for approximation and testing. In *Proceedings of FOCS*, pages 22–31, 2009.
- [HT74] J. Hopcroft and R. E. Tarjan. Efficient planarity testing. *JACM*, 21(4):549–568, 1974.
- [HW74] J. E. Hopcroft and J. K. Wong. Linear time algorithm for isomorphism of planar graphs (preliminary report). In *Proceedings of STOC*, pages 172–184, 1974.
- [NS11] I. Newman and C. Sohler. Every property of hyperfinite graphs is testable. In *Proceedings of STOC*, pages 675–684, 2011.
- [Ona10] K. Onak. *New Sublinear Methods in the Struggle Against Classical Problems*. PhD thesis, MIT, 2010.
- [Ona12] K. Onak. On the complexity of learning and testing hyperfinite graphs. Available from the author’s website, 2012.
- [PR08] S. Pettie and V. Ramachandran. Randomized minimum spanning tree algorithms using exponentially fewer random bits. *TALG*, 4(1), 2008.
- [RS95] N. Robertson and P. D. Seymour. Graph minors. XIII. the disjoint paths problem. *J. Comb. Theory Ser. B*, 63(1):65–110, 1995.
- [RS04] N. Robertson and P. D. Seymour. Graph minors. XX. Wagner’s conjecture. *J. Comb. Theory Ser. B*, 92(1):325–357, 2004.
- [SH99] W. K. Shih and W. L. Hsu. A new planarity test. *TCS*, 223(1-2):179–191, 1999.
- [YI10] Y. Yoshida and H. Ito. Testing outerplanarity of bounded degree graphs. In *Proceedings of RANDOM*, pages 642–655, 2010.