

Every Minor-Closed Property of Sparse Graphs is Testable*

Itai Benjamini [†] Oded Schramm [‡] Asaf Shapira [§]

Abstract

Suppose G is a graph of bounded degree d , and one needs to remove ϵn of its edges in order to make it planar. We show that in this case the statistics of local neighborhoods around vertices of G is far from the statistics of local neighborhoods around vertices of any planar graph G' . In fact, a similar result is proved for any minor-closed property of bounded degree graphs.

The main motivation of the above result comes from theoretical computer-science. Using our main result we infer that for any minor-closed property \mathcal{P} , there is a constant time algorithm for detecting if a graph is “far” from satisfying \mathcal{P} . This, in particular, answers an open problem of Goldreich and Ron [STOC 1997], who asked if such an algorithm exists when \mathcal{P} is the graph property of being planar. The proof combines results from the theory of graph minors with results on convergent sequences of sparse graphs, which rely on martingale arguments.

1 Introduction

Suppose we are given an n -vertex graph of bounded degree and are asked to decide if it is planar. This problem is well known to be solvable in time $\Theta(n)$ [26]. But suppose we are only asked to distinguish with probability $2/3$ between the case that the input is planar from the case that an ϵ -fraction of its edges should be removed in order to make it planar. Can we design a faster algorithm for this relaxed version of the problem, with running time $o(n)$? A special case of our main result is that in this case we can even design an algorithm whose running time is a *constant* that depends only on ϵ (and the bound on the degrees) and is independent of the size of the input.

1.1 Background on property testing

Before stating our main result, which is a significant generalization of the above mentioned result on planarity testing, let us briefly introduce the basic notions in the field of property testing. The

*A Preliminary version of this paper appeared in the Proc. of the 40th ACM Symposium on Theory of Computing (STOC) 2008.

[†]The Weizmann Institute and Microsoft Research

[‡]Microsoft Research

[§]School of Mathematics and School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332. Supported in part by NSF Grant DMS-0901355.

meta problem is the following: given a combinatorial structure S , distinguish if S satisfies some property \mathcal{P} or if S is ϵ -far from satisfying \mathcal{P} , where S is said to be ϵ -far from satisfying \mathcal{P} if an ϵ -fraction of its representation should be modified in order to make S satisfy \mathcal{P} . The main goal is to design randomized algorithms, which look at a very small portion of the input, and using this information distinguish with high probability between the above two cases. Such algorithms are called *property testers* or simply *testers* for the property \mathcal{P} . Preferably, a tester should look at a portion of the input whose size is a function of ϵ only. Blum, Luby and Rubinfeld [6] were the first to formulate a question of this type, and the general notion of property testing was first formulated by Rubinfeld and Sudan [40]. The study of testing properties of combinatorial structures, and in particular properties of graphs, was initiated by Goldreich Goldwasser and Ron [20].

The main focus of this paper is on testing properties of graphs in the *bounded degree* model, which was first introduced and studied by Goldreich and Ron [21]. In this model, we fix a degree bound d and represent graphs using adjacency lists. More precisely, we assume that a graph G is represented as a function $f_G : [n] \times [d] \mapsto [n] \cup \{*\}$, where given a vertex $v \in V(G)$ and $1 \leq i \leq d$ the function $f(v, i)$ returns the i^{th} neighbor of v , in case v has at least i vertices. If v has less than i vertices then $f(v, i) = *$. A graph of bounded degree d is said to be ϵ -far from satisfying \mathcal{P} if one needs to execute at least ϵdn edge operations of deleting or adding an edge to G in order to turn it into a graph satisfying \mathcal{P} . (Since d is treated as a constant in this paper, the ϵdn above is just proportional to ϵn .)

A testing algorithm (or tester) \mathcal{T} for graph property¹ \mathcal{P} and accuracy ϵ is a (possibly randomized) algorithm that distinguishes with probability at least $2/3$ between graphs satisfying \mathcal{P} from graph that are ϵ -far from satisfying it. More precisely, if the input graph satisfies \mathcal{P} , then \mathcal{T} accepts it with probability at least $2/3$, where the probability is taken over the coin tosses of \mathcal{T} . Similarly, if the graph is ϵ -far from satisfying \mathcal{P} , then \mathcal{T} should reject it with probability at least $2/3$. The tester is given n as input² and is provided with access to the function f_G as a black box. We define the query complexity $q_{\mathcal{T}}(n, \epsilon)$ of the tester \mathcal{T} as the maximal number of f_G -calls the tester executes on any graph G with n vertices. The most surprising aspect of property testing is that for many properties one can design a corresponding testing algorithm whose running time is independent of the size of the input, and only depends on the error parameter ϵ . Let us define this notion of efficient testing:

Definition (Testable). *A graph property \mathcal{P} is testable if there exists a function $q_{\mathcal{P}}(n, \epsilon)$ satisfying $\sup_n q_{\mathcal{P}}(n, \epsilon) < c(\epsilon)$ for every $0 < \epsilon < 1$ such that the following holds: there is a tester $\mathcal{T}_{\mathcal{P}}$ that can test \mathcal{P} on n -vertex graphs with accuracy ϵ using at most $q_{\mathcal{P}}(n, \epsilon)$ queries. In other words, if for any $\epsilon > 0$, there is a constant time randomized algorithm that can distinguish with high probability between graphs satisfying \mathcal{P} from those that are ϵ -far from satisfying it.*

¹As usual, a graph property is simply a family of graphs closed under graph isomorphism.

²Our tester needs n as the input only to be able to pick a vertex at random from the tested graph.

As we have mentioned above, our main result deals with testing properties of bounded degree graphs. Let us briefly mention some results on testing properties of *dense* graphs, a related model of property testing model that was first introduced and studied by Goldreich, Goldwasser and Ron [20]. In this model, a graph G is said to be ϵ -far from satisfying a property \mathcal{P} , if one needs to add/delete at least ϵn^2 edges to/from G in order to turn it into a graph satisfying \mathcal{P} . The tester can ask an oracle whether a pair of vertices, say i and j , are adjacent in the input graph G . It was shown in [20] that a very general family of graph “partition problems” are all testable in dense graphs. This family includes properties like being k -colorable, having a large cut, and having a large clique. Alon and Shapira [3] have shown that every hereditary graph property³ is testable in dense graphs. This also gave an (essential) characterization of the graph properties that are testable with one-sided error⁴. A characterization of the properties that are testable in dense graphs was obtained by Alon et al. [1]. Note that in this model (as its name suggests) we implicitly assume that the input graph is dense, because the definition of ϵ -far is relative to n^2 . Therefore, some properties are trivially testable in this model. In particular, minor-closed properties are trivially testable in this model with $O(1/\epsilon)$ queries and even with one-sided error. This follows from the results of Kostochka and Thomason [28, 29, 43, 44] that every finite graph with average degree $\Omega(r\sqrt{\log r})$ contains every graph on r vertices as a minor. Therefore, every large enough finite graph with $\Omega(n^2)$ edges does not satisfy a minor-closed property.

As the above mentioned results indicate, testing properties in dense graphs is relatively well understood. In sharp contrast, our current understanding of testing properties in the bounded-degree model is much more limited. For example, while every hereditary property is testable in dense graphs, in bounded degree graphs some properties are testable (e.g. being triangle-free [21]), some require $\tilde{\Theta}(\sqrt{n})$ queries (e.g. being bipartite [21, 22]) and some require $\Theta(n)$ queries (e.g. 3-colorability [8]). Besides the above mentioned results, it was also shown in [21] that k -connectivity (for any fixed k), being Eulerian and being Cycle-free are testable in bounded degree graphs. Czumaj, Sohler and Shapira [10] have recently shown that every hereditary property is testable if the input graph is guaranteed to be “non-expanding”. Some of the arguments in the present paper are motivated by some of the ideas from [10]. Another property of bounded degree graphs, which has received a lot of attention is that of being an expander, see [10, 23, 27, 34].

One reason for the fact that we understand testing of dense graphs more than we understand testing of bounded-degree graphs is that there are structural results “describing” dense graphs, primarily Szemerédi’s regularity lemma [42], while there are no similar results for arbitrary sparse graphs. Our main result here, Theorem 1.1 below, is the first to show that a general (and natural) family of properties are all testable in bounded-degree graphs.

³A graph property is hereditary if it is closed under removal of vertices. Therefore, any minor closed property is hereditary

⁴A tester has one-sided error if it always accepts graphs satisfying the property.

For more details on property testing, see the surveys [4, 11, 19, 38, 39] and Section 4 where we discuss another model of testing graph properties.

1.2 Minor-closed properties and the main result

Our main result deals with the testing of minor closed graph properties. Let us briefly introduce the basic notions in this area, which is too rich to thoroughly survey here. For more details, see Chapter 12 of [14] and the recent article of Lovász [33] on the subject. A graph H is said to be a *minor* of a graph G , if H can be obtained from G using a sequence of vertex removals, edge removals and edge contractions⁵. Equivalently, a graph G contains an h -vertex graph H as a minor, if G contains h pairwise disjoint vertex sets V_1, \dots, V_h such that the graph induced by G on each of these sets is connected, and if $(i, j) \in E(H)$ then G contains at least one edge connecting a vertex of V_i to a vertex of V_j . If H is not a minor of G , then G is said to be *H -minor free*. A graph property \mathcal{P} is *minor-closed* if every minor of a graph in \mathcal{P} is also in \mathcal{P} , or equivalently if \mathcal{P} is closed under removal of edges, removal of vertices and contraction of edges. The topic of graph minors is among the most (or perhaps the single most) studied concepts in graph theory. Our main result in this paper is accordingly the following:

Theorem 1.1 (Main Result). *For every (finite) graph H , the property of being H -minor free is testable. More generally, every minor-closed graph property is testable.*

Perhaps the most well-known result in the area of graph minors is the Kuratowski-Wagner Theorem [30, 45], which states that a graph is planar if and only if it is K_5 -minor free and $K_{3,3}$ -minor free. This fundamental result raised the natural question if a similar characterization, using a finite family of forbidden minors, also holds for embedding graphs in other fixed surfaces? Observe that a graph remains planar if we remove one of its edges or vertices and if we contract one of its edge. In fact, this closure property under these three basic operations holds for the property of a graph being embeddable in any specific surface. Thus the property of being embeddable in a specific surface is minor-closed. In one of the deepest results in graph theory, Robertson and Seymour proved the so called Graph-Minor Theorem [37], which states that for every minor-closed graph property \mathcal{P} , there is a *finite* family of graphs $\mathcal{H}_{\mathcal{P}}$ such that a graph satisfies \mathcal{P} if and only if it is H -minor free for all $H \in \mathcal{H}_{\mathcal{P}}$. Note that this in particular answers the above mentioned problem regarding the characterization of graphs embeddable in a fixed surface using a finite number of forbidden minors.

We note that besides the above (mainly) graph theoretic motivation, the graph minors and minor closed properties have also received a considerable amount of attention by the Computer-Science community, because many natural graph properties (e.g., planarity) are described by excluded

⁵Contracting an edge connecting vertices u, v is the result of replacing u and v by a new vertex w , and connecting w to all the vertices that were connected to either u or v . Loops and multiple edges are removed.

minors, and graph problems that are NP-hard in general can often be solved on excluded minor families in polynomial time, or at least be approximated better than on arbitrary graphs. See [13] and its references.

Let us mention some well studied minor-closed graph properties. Of course, the most well known such property is **Planarity**. A well-studied variant of planarity is **Outer-planarity**, the property of being embeddable in the plane in such a way that all vertices lie on the outer face. Another generalization of planarity is being embeddable in a surface of genus at most k . Graphs satisfying this property are said to have **genus** k . The **Tree-width** of a graph is one of the most important invariants of graphs. This notion, which measures how close a graph is to being a tree, was introduced by Robertson and Seymour as part of their proof of the Graph-Minor Theorem; see [7]. It also has numerous applications in the area of fixed-parameter algorithms; see [16] for more details. As it turns out, having bounded tree-width is also a minor-closed graph property. Another well-known minor-closed property is being **Series-parallel**. Series parallel graphs are graphs that can be obtained from a single edge by sequence of parallel extensions (adding an edge parallel to an edge that already exists) and series extensions (subdividing an edge by a new node). We conclude with the property of being **Knotlessly-embeddable**, which is the property of being embeddable in \mathbb{R}^3 in a way that no two cycles are linked and no cycle is knotted.

By Theorem 1.1 we infer that the above mentioned properties are all testable with a constant number of queries. It is interesting to note that prior to this work none of the above properties was even known to be testable with $q_T(n) = o(n)$ queries. In fact, in the paper which introduced property testing of bounded degree graphs [21], Goldreich and Ron asked if planarity can be tested in constant time, a result which follows from Theorem 1.1.

One important aspect of Theorem 1.1, which we have neglected to address thus far, is the actual dependence of the query-complexity of the testers on ϵ . As it does not seem like one can achieve a query complexity that is sub-exponential in $1/\epsilon$ using our approach, we opted in several places not to give explicit bounds. We further discuss this issue in Section 4, where we show how can one derive explicit (but rather large) upper bounds on the query complexity of the testers as a function of ϵ .

1.3 Hyper-finiteness and testing monotone hyper-finite properties

The following notion of hyper-finiteness was defined by Elek [17] (though it is implicit in [32], for example).

Definition (Hyper-Finite). *A graph $G = (V, E)$ is (δ, k) -hyper-finite if one can remove $\delta|V|$ edges from G and obtain a graph with connected components of size at most k . A collection of graphs \mathcal{G} is hyper-finite if for every $\delta > 0$ there is some finite $k = k(\delta)$ such that every graph in \mathcal{G} is (δ, k) -hyper-finite.*

Our main theorem, Theorem 1.1, will be a corollary of the following more general result, where as usual, a graph property \mathcal{P} is *monotone* if every subgraph of a graph in \mathcal{P} is also in \mathcal{P} .

Theorem 1.2. *Every monotone hyper-finite graph property is testable.*

Consider the property that the number of vertices at distance r around every vertex is at most some fixed function $g(r)$ satisfying $g(r) = 2^{o(r)}$. Such graphs are said to have *sub-exponential growth*. Then, Theorem 1.2 can be used to show that this property is testable. This does not seem to follow directly from Theorem 1.1.

1.4 Recent results

Elek [18] has studied the problem of testing properties of graphs of sub-exponential growth (defined at the end of the previous subsection). Besides being able to handle some hereditary properties like the ones studied in [10], his approach can also handle properties like the size of the largest independent-set or the size of the smallest dominating-set.

As we mention in Section 4, the query complexity of our algorithm for testing if a graph is H -minor free is triply-exponential in $\text{poly}(1/\epsilon)$. Very recently, Hassidim et al. [25] have simplified our proof and also obtained an improved upper bound for the query complexity, which is only singly exponential in $\text{poly}(1/\epsilon)$. They have also showed that one can approximate (rather than only test) the distance of a minor-free graph from satisfying a hereditary property, thus strengthening the result of [10].

1.5 Techniques and overview of the paper

In the next section, we introduce a metric condition for testability. Basically, we define a sequence of pseudometrics ρ_r indexed by an integer $r \in \mathbb{N}_+$, where $\rho_r(G, G')$ measures the difference between G and G' in the frequency of isomorphism types of r -neighborhoods of vertices. We show that if for some fixed $r > 0$ the ρ_r distance between two graph families is positive, then these graph families can be distinguished by a tester.

Next, we state a theorem (Theorem 2.2) that roughly says that (ϵ, k) -hyper-finite graphs are far away in some pseudometric ρ_R from graphs that are not (ϵ', k) -hyper-finite, where ϵ' depends on ϵ and tends to zero as $\epsilon \searrow 0$. This result can be deduced from a recent result of Schramm [41], concerning properties of convergent sequences of bounded degree graphs. However, we provide an alternative self-contained proof in Section 3. In contrast with [41], our proof is finitary and gives an explicit upper bound on R , while the proof in [41] did not supply such a bound. We note that convergent sequences of graphs have been previously used in the study of property testing of *dense* graphs [9].

How does a tester for a monotone hyperfinite graph property \mathcal{P} work? A first guess might be the following: as the property is monotone, we may expect to sample enough vertices such

that with high probability the neighborhood of at least one of them will not satisfy the property, thus establishing that the graph itself does not satisfy the property. This works for properties like triangle-freeness, however, it is not difficult to see (see item 2, in the concluding remarks) that this approach of aiming for a one-sided error tester is bound to fail in our case. The reason is that a graph can be far from being (say) planar, yet locally be planar. But still, the local neighborhoods of vertices do tell us something about its global properties. For example, in [21], it was shown that the local density around vertices can be used to test the property of being cycle-free. By the above discussion we must look for another property that the graph must satisfy if it satisfies a minor closed property. As it turns out, the right property to look for is hyper-finiteness.

We now proceed with an informal and not very precise description of how and why our tester works. Given an input graph G , the tester first tests for hyper-finiteness. The test for hyper-finiteness proceeds by picking a bounded number of vertices at random and exploring the R -neighborhood of each one (where the number of vertices chosen and R depend on ϵ). It then checks if the observed frequency of the isomorphism types of these neighborhoods looks approximately like some (η, k) -hyper-finite graph, with appropriate values for η and k . We do not have an explicit characterization of the approximate frequencies occurring in hyper-finite graphs, but since only an approximation is needed (with known accuracy), there is a finite table listing the occurring frequencies, which can be part of the algorithm. The point here is that the table does not depend on the size of the tested graph. If the graph fails the hyper-finiteness test, then it is rejected. If it passes, then we know that by removing a small proportion of the edges it is broken down to pieces of size k . As these small pieces are still far from satisfying the property, we can use a bounded number of random samples to actually find a subgraph of the input that does not satisfy the property.

The rest of the paper is organized as follows. In Section 2 we prove all of the results stated above, with the exception of Theorem 2.2, which is proved in Section 3. In particular, by combining a theorem of Lipton and Tarjan with a result of Alon, Seymour and Thomas [2], regarding separators in minor-free graphs, we get that H -minor free graphs are hyper-finite. This facilitates a proof of Theorem 1.1 from Theorem 1.2. In Section 4 we discuss several open problems and conjectures for future research.

2 A Metric Criterion for Testability

Let us slightly generalize the notion of property testing, and say that two graph properties A and B are *distinguishable* if there is a randomized algorithm that makes a bounded number of queries to an input graph and accepts every graph in A with probability at least $2/3$ and rejects every graph in B with probability at least $2/3$. If the input does not belong to either A or B the algorithm is allowed to return an arbitrary answer. Then ϵ -testing A is equivalent to taking B to be the set of graphs that are ϵ -far from A .

Shortly, we will define a useful pseudo-metric⁶ on the set of graphs of bounded degree d . But first, we need to set some terminology. A *rooted graph* is a pair (G, v) , where G is a graph and $v \in V(G)$. When the root v of a pair (G, v) will not be important, we will sometimes refer to a rooted graph (G, v) simply as G .

An *isomorphism between rooted graphs* (H, u) and (G, v) is an isomorphism between the underlying graphs H and G that maps u to v . For a vertex $v \in V(G)$, we denote by $B_G(v, r)$ the induced subgraph of G whose vertices consist of the vertices of G at distance at most r from v . Consider a rooted graph H and a finite graph G . Let $m_r^G(H)$ denote the number of vertices $v \in V(G)$ such that there is a rooted-graph isomorphism from $(B_G(v, r), v)$ onto H . (Of course, this is often zero.) Set

$$\mu_r^G(H) := \frac{m_r^G(H)}{|V(G)|},$$

and define a pseudometric ρ_r by

$$\rho_r(G, G') := \sum_H |\mu_r^G(H) - \mu_r^{G'}(H)|,$$

where the sum extends over all isomorphism types of rooted graphs. Clearly, the number of terms that are nonzero is bounded by a constant that depends only on r and a bound on the degrees in G and G' . Observe that μ_r^G defines a probability measure on the set of rooted graphs and that ρ_r is monotone non-decreasing in r . If A and B are graph families, we define

$$\rho_r(A, B) := \inf\{\rho_r(G, G') : G \in A, G' \in B\}.$$

The following proposition gives a metric condition for distinguishability.

Proposition 2.1. *Let A and B be two graph properties having only graphs with degrees at most d . If there is some integer $R > 0$ such that $\rho_R(A, B) > 0$, then A and B are distinguishable.*

This gives a sufficient condition for A to be testable: if for every $\epsilon > 0$ the set $B(\epsilon)$ of graphs that are ϵ -far from A satisfies $\sup_R \rho_R(A, B(\epsilon)) > 0$, then A is testable.

The converse to Proposition 2.1 does not hold. For example, if A is the collection of graphs with an even number of vertices and B is the collection with an odd number of vertices, then A and B are distinguishable in the current model in which the number of vertices of the graph is given as the input. However, it is not hard to come up with a natural model for property testing in which the criterion given by the proposition is necessary and sufficient.

The primary reason that makes hyper-finiteness so useful for property testing is the following theorem.

⁶A pseudometric on a set X differs from a metric on X in that a pseudometric is allowed to be zero on pairs (x, y) with $x \neq y$.

Theorem 2.2. Fix $d, k \in \mathbb{N}_+$ and $\epsilon > 0$. Let A be the set of (ϵ, k) -hyper-finite graphs with vertex degrees bounded by d , and let B be the set of finite graphs with vertex degrees bounded by d that are not $(4\epsilon \log(4d/\epsilon), k)$ -hyper-finite. Then there is some $R = R(d, k, \epsilon) \in \mathbb{N}_+$ such that $\rho_R(A, B) > 0$.

Theorem 2.2 can actually be deduced from a recent result of Schramm [41], which deals with infinite unimodular graphs. The proof of Theorem 2.2 is postponed to Section 3, where we present a proof that is adapted to the finite setting, gives quantitative bounds on R and $\rho_R(A, B)$, and would hopefully be more accessible. We now turn to prove Proposition 2.1

Proof of Proposition 2.1. Suppose that for some integer $R > 0$ and some positive $\delta > 0$ we have $\rho_R(A, B) > \delta$. In that case, we can distinguish between A and B using the following algorithm. Let \mathcal{H} denote the set of (isomorphism types of) rooted graphs (H, v) of radius R around v and maximum degree at most d , and set $h := |\mathcal{H}|$. Then, given an input graph G of size n , we estimate $\mu_R^G(H)$ for every $H \in \mathcal{H}$, up to an additive error of $\delta/(2h)$, with success probability $1 - \frac{1}{4h}$. Here, we can apply an additive Chernoff bound to deduce that to this end it is enough to sample $O(\frac{h^2}{\delta^2} \log h)$ vertices, explore their R -neighborhood, and compute the fraction of these vertices whose neighborhood is isomorphic to (H, v) . Let $\hat{\mu}_R^G(H)$ be the estimated values of $\mu_R^G(H)$, and recall that when computing $\rho_R(G, G')$ we only need to consider rooted graphs (H, v) of radius R . The algorithm now checks if there exists a graph $G_A \in A$ for which $\sum_H |\hat{\mu}_R^{G_A}(H) - \mu_R^{G_A}(H)| \leq \frac{1}{2}\delta$. Observe that to this end, the algorithm does not have to actually search all possible graphs. It can just store a $\delta/4$ -net of the set of all possible h -tuples of values of $\mu^{G_A}(H)$, taken over all graphs $G_A \in A$ (this is a finite list). If that is the case, the algorithm declares that G belongs to A , and otherwise it declares that G belongs to B . Note that since $\rho_R(A, B) > \delta$ and $\sum_H |\hat{\mu}_R^G(H) - \mu_R^G(H)| \leq \frac{1}{2}\delta$ with high probability, the algorithm is unlikely to misclassify graphs in A or in B . \square

We now turn to prove Theorems 1.2 using Theorem 2.2 and Proposition 2.1.

Proof of Theorem 1.2. We first describe the algorithm for testing if a graph is ϵ -far from satisfying \mathcal{P} . Given some $0 < \epsilon < 1$, let $\epsilon_0 = \epsilon_0(\epsilon)$ be sufficiently small so that $4\epsilon_0 \log(4d/\epsilon_0) < \frac{1}{2}\epsilon$. Let k be such that each graph in \mathcal{P} is (ϵ_0, k) -hyper-finite. Given an input graph G the algorithm performs the following two steps. In the first step, it invokes Proposition 2.1 in order to distinguish (with high probability and using a constant number of queries) between the case that the input is (ϵ_0, k) -hyper-finite from the case that it is not $(\frac{1}{2}\epsilon, k)$ -hyper-finite. If this procedure declares that G is not (ϵ_0, k) -hyper-finite the algorithm declares that G does not satisfy \mathcal{P} . If this procedure declares that G is $(\frac{1}{2}\epsilon, k)$ -hyper-finite then the algorithm samples (with repetitions) a constant number $m = m(\epsilon)$ of vertices of G denoted v_1, \dots, v_m . For each $1 \leq i \leq m$ the algorithm explores the neighborhood of v_i of radius k , that is, it explores $B_G(v_i, k)$. If the union of these neighborhoods forms a graph that does not satisfy \mathcal{P} the algorithm declares that G does not satisfy \mathcal{P} , otherwise it declares that G satisfies \mathcal{P} .

We now prove that the above algorithm uses a constant number of queries and that with probability at least $2/3$ it distinguishes between graphs satisfying \mathcal{P} and those that are ϵ -far from satisfying it. We start with the first step. Since $4\epsilon_0 \log(4d/\epsilon_0) < \frac{1}{2}\epsilon$ we deduce from Theorem 2.2 that there is some fixed $R = R(\mathcal{P}, \epsilon)$ such that if G_1 is (ϵ_0, k) -hyper-finite and G_2 is not $(\frac{1}{2}\epsilon, k)$ -hyper-finite then $\rho_R(G_1, G_2) > \delta = \delta(\epsilon)$. Therefore, by Proposition 2.1 we can indeed perform the first step of the proof using a constant number of queries (the depends only on ϵ).

Suppose first that the input G satisfies \mathcal{P} . Then by our choice of k , the input is (ϵ_0, k) -hyper-finite. Therefore, with high probability it will pass the first step of the algorithm. Since \mathcal{P} is monotone every subgraph of G satisfies \mathcal{P} , therefore the input will pass the second step (with probability 1). We conclude that every graph satisfying \mathcal{P} will be accepted with high probability.

Assume now that G is ϵ -far from satisfying \mathcal{P} . To handle this case we will need to introduce some notation. Let $\mathcal{K}(k)$ be the family of all (non-isomorphic) connected graphs of size at most k . Since k depends only on ϵ we have that $\mathcal{K}(k)$ is a finite set whose size depends only on ϵ . For a collection of graphs $S \subseteq \mathcal{K}(k)$ define the graph $G(S)$ to be the disjoint union of the graphs of S , and let $g = g(S)$ be the smallest integer such that the graph obtained by taking g vertex disjoint copies of $G(S)$ does not satisfy \mathcal{P} . If no such integer exists then we set $g(S) = \infty$. Finally, let

$$g(k) = \max_{S \subseteq \mathcal{K}(k) : g(S) < \infty} g(S) .$$

Note that the above is well defined since $\mathcal{K}(k)$ is finite so the maximum if over a finite set. Since k is determined by ϵ we have that $g(k)$ is bounded from above by a function of ϵ . We now get back to analyzing the algorithm. If G is not $(\frac{1}{2}\epsilon, k)$ -hyper-finite then it will be rejected with high probability by the first step of the algorithm. So assume that G is ϵ -far from satisfying \mathcal{P} and is $(\frac{1}{2}\epsilon, k)$ -hyper-finite, and let G' be the graph obtained from G by removing (not more than) $\frac{1}{2}\epsilon dn$ edges in a way that all the connected components of G' are of size at most k . Observe that G' is still at least $\frac{1}{2}\epsilon$ -far from satisfying \mathcal{P} . Let $S' = S'(G')$ be the subset of $\mathcal{K}(k)$ that contains a graph $K \in \mathcal{K}(k)$ if and only if one of the connected components of G' is isomorphic to K . We now define a subset $S'' \subseteq S'$ and a subgraph of G' denoted G'' as follows; for every $K \in S'$ if G' has less than $\frac{1}{2}\epsilon n / |\mathcal{K}(k)| k$ connected components that are isomorphic to K then remove K from S' and remove all the edges from all the connected components of G' that are isomorphic to K . Note that since every graph $K \in \mathcal{K}(k)$ has at most k vertices and the maximum degree is d , we thus remove for every $K \in S'$ not more than $\frac{1}{4}\epsilon n / |\mathcal{K}(k)|$ edges, and since $|S'| \leq |\mathcal{K}(k)|$ the total number of edges removed is bounded by $\frac{1}{4}\epsilon dn$. Hence G'' does not satisfy \mathcal{P} (actually, it is $\frac{1}{4}\epsilon$ -far from satisfying \mathcal{P}). By definition, every connected components of G'' is isomorphic to one of the graphs in S'' . Since G'' does not satisfy \mathcal{P} , and \mathcal{P} is monotone, we get that $g(S'') < \infty$, implying that $g(S'') \leq g(k)$. Since every graph $K \in S''$ is isomorphic to at least $\frac{1}{2}\epsilon n / |\mathcal{K}(k)| k$ of the connected components of G'' , we get that a randomly chosen vertex from G'' belongs to a connected component isomorphic to K with probability at least $\frac{1}{2}\epsilon / |\mathcal{K}(k)| k$. Therefore, if we sample $4g(k)|\mathcal{K}(k)|^2 k / \epsilon$ vertices then

the expected number of vertices that belong to a connected component isomorphic to K is at least $2g(k)|\mathcal{K}(k)|$. Hence, by a Chernoff bound the probability of having less than $g(k)$ such vertices is at most $1/4|\mathcal{K}(k)|$ (actually, much smaller than that). If we repeat this process $|\mathcal{K}(k)|$ times, then with probability at least $3/4$ we obtain a collection of $|S''| \cdot g(k)$ connected components which form the disjoint union of $g(k) \geq g(S'')$ copies of $G(S'')$. Hence we can take $m(\epsilon)$ in the description of the algorithm to be $4g(k)|\mathcal{K}(k)|^3 k/\epsilon$. Furthermore, assuming G is large enough⁷ (as a function of ϵ) these $g(S'') \cdot |S''|$ connected components will be distinct with high probability. By definition, this graph does not satisfy \mathcal{P} . Finally, since G'' is a subgraph of G and \mathcal{P} is monotone, this means that with high probability the union of the neighborhoods of the vertices v_1, \dots, v_m will not satisfy \mathcal{P} , and so G will be rejected with high probability at the second stage of the algorithm. \square

Recall the Lipton-Tarjan planar separator theorem [31], which says that every planar graph G has a set of vertices V_0 of size $O(|V(G)|^{1/2})$, such that every connected component of $G \setminus V_0$ has at most $2|V(G)|/3$ vertices. Lipton and Tarjan used this theorem to show that the set of planar graphs with degree bounded by d is hyper-finite [32, Theorem 3]. Alon, Seymour and Thomas [2] proved that the planar separator theorem as stated above holds more generally for the class of H -minor free, where H can be any finite graphs⁸. The proof of [32, Theorem 3], along with the main result of [2], therefore gives the following result.

Proposition 2.3. *Fix $d \in \mathbb{N}_+$. Let H be a finite graph, and let A be the set of graphs that are H -minor free and have degrees bounded by d . Then A is hyper-finite.* \square

The proof of Theorem 1.1 is now an easy corollary of the above results.

Proof of Theorem 1.1. Since being H -minor free is a minor-closed property, it is enough to prove the second claim of the theorem: that minor closed graph properties are testable. Let \mathcal{P} be a minor closed graph property. If \mathcal{P} includes all graphs, then it is clearly testable. Otherwise, suppose that $H \notin \mathcal{P}$. Then \mathcal{P} is H -minor free. By Proposition 2.3, \mathcal{P} is hyperfinite. Since \mathcal{P} is monotone, the theorem now follows from Theorem 1.2. \square

We finally note that the proofs of Theorems 1.2 and 1.1 given above also imply the following corollary which was stated in the abstract.

Corollary 2.4. *For every $\epsilon > 0$ and $d \in \mathbb{N}_+$ there is an $R = R(\epsilon, d) \in \mathbb{N}_+$ such that if G and G' are finite graphs with vertex degrees bounded by d , G is planar and G' is ϵ -far from being planar, then $\rho_R(G, G') \geq 1/R$.*

⁷If G is not large enough than some function of ϵ , then we can simply ask about all the edges of G and thus return an exact answer.

⁸The implicit constant in the $O(\cdot)$ notation above is explicit in [31]. Also [2] gives an estimate for this constant, which naturally depends on H .

3 Proof of Theorem 2.2

Let us start with an overview of the proof of Theorem 2.2. Assuming that G is (ϵ, k) -hyper-finite and that G' is a graph satisfying $\rho_R(G, G') \leq \delta$ with some appropriate $\delta > 0$, we will show that G' must be $(4\epsilon \log(4d/\epsilon), k)$ -hyper-finite. First, since G is (ϵ, k) -hyper-finite, there is a subset S of the edges of G , of size at most $\epsilon|V(G)|$, that partitions G into connected components of size at most k . We then use the existence of S to construct a random $\tilde{S} \subset E(G)$, where each connected component of $G \setminus \tilde{S}$ has at most k vertices and the expected size of \tilde{S} is at most $4\epsilon \log(3d/\epsilon)|V(G)|$. The most important feature of the way we will pick \tilde{S} is that it is local, in the sense that there is a finite bound $R = R(d, \epsilon, k)$, such that the probability that an edge e is in \tilde{S} only depends on the isomorphism type of the pair⁹ $(B_G(e, R), e)$. Now, if G' is a graph satisfying $\rho_R(G, G') \leq \delta$, then locally it behaves almost exactly like G does. Therefore, choosing a set of edges \tilde{S}' from G' using the same process that was used for G , and removing it from G' should also partition G' into connected components of size at most k . Furthermore, the expected relative size of \tilde{S}' in G' is close to that of \tilde{S} in G , implying that G' is $(4\epsilon \log(4d/\epsilon), k)$ -hyper-finite.

We now turn to the formal proof of Theorem 2.2. Let $G \in A$. Since G is (ϵ, k) -hyper-finite, there is a set $S \subset E(G)$ such that $|S| \leq \epsilon|V(G)|$ and all the connected components of $G \setminus S$ are of size at most k . Fix such a set S . With no loss of generality, we assume that none of the edges of S connects two vertices from the same connected component of $G \setminus S$. The proof strategy is to replace the set S with a random set \tilde{S} that has no long-range dependencies and the probability of an edge being in \tilde{S} depends only on the local structure of G near the endpoints of the edge. It will then be easy to see that a similar \tilde{S} exists for G_0 if $\rho_R(G, G_0)$ is small and R is large.

Let us first introduce some notation. If K is a subgraph of a graph H and K' is a subgraph of a graph H' , then we say that the pair (K, H) is isomorphic to the pair (K', H') if there is an isomorphism of H onto H' that takes K onto K' . Isomorphisms of triples (H, K, J) , $K, J \subset H$, are similarly defined. For a graph G , we define $\mathcal{K}(G)$ as the set of vertex sets $K \subset V(G)$ of size at most k which span a connected graph in G . Given the set S that disconnects G into connected components of size at most k , we let $\mathcal{K}_S(G)$ consist of those elements of $\mathcal{K}(G)$ that span a connected component in $G \setminus S$. We stress that the elements in $\mathcal{K}(G)$ are only required to span a connected subgraph in G , while the elements of $\mathcal{K}_S(G)$ are required to span a *maximally* connected subgraph in $G \setminus S$. In what follows it will be convenient to sometimes identify an element $K \in \mathcal{K}(G)$ with the (connected) graph spanned by K .

Given a subgraph $H \subset G$ and $r \in \mathbb{N}$, let $N_r(H)$ denote the subgraph of G induced by the vertices at distance at most r from H . Let N be some connected graph and let K be a connected subgraph of N of size $|V(K)| \leq k$. For $r \in \mathbb{N}_+$, let $\Gamma_r(K, N)$ denote the set that consists of all

⁹If e is an edge, then $B_G(e, R)$ and (H, e) are the natural generalizations of $B_G(v, R)$ and (H, v) .

subgraphs $K' \in \mathcal{K}(G)$ such that $(K', N_r(K'))$ is isomorphic to (K, N) . Set

$$p_r(K, N) := \begin{cases} \frac{|\Gamma_r(K, N) \cap \mathcal{K}_S(G)|}{|\Gamma_r(K, N)|} & \text{if } \Gamma_r(K, N) \neq \emptyset, \\ 1 & \text{otherwise.} \end{cases}$$

If $K \in \mathcal{K}(G)$, we abbreviate,

$$p_r(K) := p_r(K, N_r(K)). \quad (3.1)$$

For $v \in V(G)$, let $\mathcal{K}(v)$ denote the elements of $\mathcal{K}(G)$ that contain v . Set

$$q_r(v) := \sum_{K \in \mathcal{K}(v)} p_r(K).$$

Intuitively, $p_r(K)$ is an approximation of the conditional probability that $K \in \mathcal{K}_S(G)$ given the isomorphism type of the pair $(K, N_r(K))$. In the following, we will need the fact that for some not too large R , very few vertices have $q_R(v)$ small. To this end, we have the following lemma.

Lemma 3.1. *There exists some finite $R_1 = R_1(d, \epsilon, k)$ and some $R \in \mathbb{N}_+ \cap [1, R_1]$ such that the number of vertices $v \in V(G)$ with $q_R(v) < 1/2$ is at most $\epsilon |V(G)|/(2d)$.*

We stress that R itself may depend on G and S , but it is bounded by R_1 , which is not allowed to depend on G and S . We postpone the proof of the lemma and continue with the proof of Theorem 2.2. Let \mathcal{K}' denote a random subset of $\mathcal{K}(G)$, where each $K \in \mathcal{K}(G)$ is in \mathcal{K}' with probability $\min(2 \log(2d/\epsilon) p_R(K), 1)$, independently, where R is as provided by Lemma 3.1. Let

$$S' := \bigcup_{K \in \mathcal{K}'} \partial K,$$

where ∂K denotes the set of edges of G that are not in K but neighbor with some vertex in K . Let W denote the set of vertices that are contained in some $K \in \mathcal{K}'$, and S'' denote the set of all edges of G having both vertices in $V(G) \setminus W$. Finally, define $\tilde{S} := S' \cup S''$. Clearly, the connected components of $G \setminus \tilde{S}$ are of size at most k . We now have to estimate $\mathbf{E}|\tilde{S}|$.

Let us start by estimating $\mathbf{E}|S''|$. Consider an arbitrary vertex $v \in V(G)$ such that $q_R(v) \geq 1/2$. Then

$$\begin{aligned} \mathbf{P}[v \notin W] &= \prod_{K \in \mathcal{K}(v)} (1 - \mathbf{P}[K \in \mathcal{K}']) \\ &\leq \exp\left(- \sum_{K \in \mathcal{K}(v)} 2 \log(2d/\epsilon) p_R(K)\right) \\ &= (\epsilon/(2d))^{2q_R(v)}. \end{aligned}$$

Therefore, if $q_R(v) \geq 1/2$, we have $\mathbf{P}[v \notin W] \leq \epsilon/(2d)$. On the other hand, the number of vertices satisfying $q_R(v) < 1/2$ is at most $\epsilon|V(G)|/(2d)$, by Lemma 3.1. Thus, we have

$$\begin{aligned} \mathbf{E}[|V(G) \setminus W|] &\leq \frac{\epsilon|V(G)|}{2d} + \sum \left\{ \mathbf{P}[v \notin W] : v \in V(G), q_R(v) \geq \frac{1}{2} \right\} \\ &\leq \frac{\epsilon|V(G)|}{d}. \end{aligned}$$

Since every edge in S'' is incident to a vertex in $V(G) \setminus W$, we get $\mathbf{E}|S''| \leq \epsilon|V(G)|$.

We now estimate $\mathbf{E}|S'|$. For $K \in \mathcal{K}(G)$ set $L(K) := \Gamma_R(K, N_R(K))$. Observe that $K' \in L(K)$ if and only if $L(K') = L(K)$. Let $\{K_1, K_2, \dots, K_m\}$ be a set of elements in $\mathcal{K}(G)$, one from each equivalence class of the relation $L(K') = L(K)$. Set $L_i := L(K_i)$. Note that $p_R(K) = p_R(K_i)$ for every $K \in L_i$, and since $R > 0$ also $|\partial K| = |\partial K_i|$ for every $K \in L_i$. Therefore, the definition of p_R gives

$$\sum_{K \in L_i} p_R(K) |\partial K| = \sum_{K \in L_i} \frac{|L_i \cap \mathcal{K}_S(G)|}{|L_i|} |\partial K_i| = \sum_{K \in L_i \cap \mathcal{K}_S(G)} |\partial K|.$$

By summing over i , we get,

$$\begin{aligned} \sum_{K \in \mathcal{K}(G)} p_R(K) |\partial K| &= \sum_i \sum_{K \in L_i} p_R(K) |\partial K| \\ &= \sum_i \sum_{K \in L_i \cap \mathcal{K}_S(G)} |\partial K| \\ &= \sum_{K \in \mathcal{K}_S(G)} |\partial K| \\ &= 2|S|. \end{aligned}$$

As each element of $\mathcal{K}(G)$ is chosen with probability $\min(2 \log(2d/\epsilon) p_R(K), 1)$ to be in \mathcal{K}' , we infer that

$$\mathbf{E}[|S'|] \leq 4 \log(2d/\epsilon) |S| \leq 4\epsilon \log(2d/\epsilon) |V(G)|.$$

Putting this together with our previous bound on $\mathbf{E}|S''|$, we obtain

$$\mathbf{E}[|\tilde{S}|] \leq 4\epsilon \log(3d/\epsilon) |V(G)|. \quad (3.2)$$

Now suppose that G_0 is any finite graph with degrees bounded by d . We can define a random set of edges $S_0 \subset E(G_0)$, as follows. Let \mathcal{K}'_0 be a random set of elements of $\mathcal{K}(G_0)$, where each $K \in \mathcal{K}(G_0)$ is placed in \mathcal{K}'_0 with probability $\min(2 \log(2d/\epsilon) p_R(K, N_R(K)), 1)$, independently. (Here, $N_R(K)$ refers to the neighborhood in G_0 , of course.) Define $S'_0 := \bigcup_{K \in \mathcal{K}'_0} \partial K$ and define S''_0 based on G_0 and S'_0 as S'' was defined based on G and S' .

In order to estimate $\mathbf{E}[|\tilde{S}_0|]$, we briefly consider the situation in G again. Note that for every vertex $v \in V(G)$ the expected number of edges that are incident with v and belong to \tilde{S} is completely determined by its neighborhood of radius $r = R+k+1$, or, more precisely, by the isomorphism type

of the rooted graph $H(v) := (B(v, r), v)$. Let t_H be this expectation when $H(v)$ is isomorphic to H , that is, the expected number of edges incident with v that will belong to \tilde{S} if the neighborhood of v is isomorphic to $H(v)$. Recall that $\mu_r^G(H) \cdot |V(G)|$ is the number of vertices $v \in V(G)$ such that $H(v)$ is isomorphic to H . Then by linearity of expectation we can write

$$\frac{2\mathbf{E}[|\tilde{S}|]}{|V(G)|} = \sum_H t_H \cdot \mu_r^G(H).$$

Similar considerations apply to G_0 and \tilde{S}_0 , and hence

$$\begin{aligned} \frac{2\mathbf{E}[|\tilde{S}_0|]}{|V(G_0)|} &= \sum_H t_H \cdot \mu_r^{G_0}(H) \\ &= \sum_H t_H \cdot \mu_r^G(H) + \sum_H t_H \cdot (\mu_r^{G_0}(H) - \mu_r^G(H)) \\ &= \frac{2\mathbf{E}[|\tilde{S}|]}{|V(G)|} + \sum_H t_H \cdot (\mu_r^{G_0}(H) - \mu_r^G(H)) \\ &\leq \frac{2\mathbf{E}[|\tilde{S}|]}{|V(G)|} + d \cdot \rho_r(G, G_0), \end{aligned}$$

where we have used the fact that for every H we have $t_H \leq d$ since G_0 is of bounded degree d . Thus, $G_0 \notin B$ if

$$\rho_r(G, G_0) < (8/d)\epsilon \log(4/3). \quad (3.3)$$

This completes the proof of Theorem 2.2. \square

Before turning to the proof of Lemma 3.1 we prove the following simple fact

Lemma 3.2. *Let T be a rooted tree with maximum degree d . Then the number of subtrees of T of size k containing the root is at most d^{2k} .*

Proof. Every subtree of T of size k can be realized as a closed walk of length $2k$; just double each edge and take an Euler tour. Therefore, the number of such sub-trees is at most the number of closed walks of length $2k$ which is d^{2k} . \square

Proof of Lemma 3.1. Let o be a uniformly random vertex chosen from $V(G)$. Let \mathcal{G} be a set of finite rooted graphs, which has exactly one representative for each isomorphism class of rooted graphs. For $r \in \mathbb{N}$, let H_r denote the random element from \mathcal{G} that is isomorphic to the rooted graph $(B(o, r), o)$. Let \mathcal{Z} denote the set of pairs (H, K) such that $H \in \mathcal{G}$ and K is a connected subset of $V(H)$ of cardinality at most k which contains the root of H . It is important to note that unlike \mathcal{G} , which does not contain two isomorphic copies of the same graph, the set \mathcal{Z} may contain two isomorphic pairs (H, K) and (H, K') .

For $Z = (H, K) \in \mathcal{Z}$, we would like to define $\hat{p}_r(Z)$ as the conditional probability given H_r that $H = H_k$ and the isomorphic image of K in $B(o, k)$ is in $\mathcal{K}_S(G)$. But here is a slightly annoying technical point. If H_k has nontrivial automorphisms, then there is no unique isomorphic image of K in $B(o, k)$. For this reason, given o and r , we take a random isomorphism γ_r , chosen uniformly among all rooted-graph isomorphism from H_r onto $(B(o, r), o)$, with independent choices for each r . Let β_r denote the sequence of maps $(\gamma_{j+1}^{-1} \circ \gamma_j : j = 0, 1, \dots, r-1)$. Intuitively, β_r tells us for each $r' < r$ how $H_{r'}$ sits inside H_r .

For $Z = (H, K) \in \mathcal{Z}$ let A_Z denote the event that $H_k = H$ and $\gamma_k(K) \in \mathcal{K}_S(G)$, and set for every $Z \in \mathcal{Z}$

$$\hat{p}_r(Z) := \mathbf{P}[A_Z \mid H_r, \beta_r]. \quad (3.4)$$

Clearly, exactly one of the events $A_Z, Z \in \mathcal{Z}$, holds. Therefore,

$$\sum_{Z \in \mathcal{Z}} \hat{p}_r(Z) = 1 \quad (3.5)$$

holds for every $r \in \mathbb{N}$. Since (H_r, β_r) can be determined from (H_{r+1}, β_{r+1}) , it is clear that

$$\hat{p}_0(Z), \hat{p}_1(Z), \hat{p}_2(Z), \dots$$

is a martingale for any fixed $Z \in \mathcal{Z}$. (This is always the case when conditioning on finer and finer σ -fields.) Since $\hat{p}_r(Z) \in [0, 1]$ and $\mathbf{E}[\hat{p}_r(Z)] = \mathbf{P}[A_Z] = \hat{p}_0(Z)$, we have $\mathbf{E}[\hat{p}_r(Z)^2] \leq \hat{p}_0(Z)$. Therefore, the variance of $\hat{p}_r(Z)$ is at most $\hat{p}_0(Z)$.

Now recall that for a martingale M_1, M_2, \dots we have the so called ‘‘orthogonality for martingale differences’’, which says that

$$\sum_{j=1}^{n-1} \mathbf{E}[(M_{j+1} - M_j)^2] = \mathbf{E}[M_n^2] - \mathbf{E}[M_1^2] \leq \text{Var}(M_n). \quad (3.6)$$

We apply this to the sequence $\hat{p}_{ki}(Z), i = 0, 1, 2, \dots$, (which is also a martingale), sum over $Z \in \mathcal{Z}$ and use our above bound on the variance of $\hat{p}_r(Z)$, to obtain for every $j \in \mathbb{N}_+$

$$\sum_{i=0}^{j-1} \sum_{Z \in \mathcal{Z}} \mathbf{E}[(\hat{p}_{(i+1)k}(Z) - \hat{p}_{ik}(Z))^2] \leq \sum_{Z \in \mathcal{Z}} \text{Var}(\hat{p}_{jk}(Z)) \leq \sum_{Z \in \mathcal{Z}} \hat{p}_0(Z) = 1. \quad (3.7)$$

It is now time to make the connection between \hat{p}_r and p_r . Fix some $i \in \mathbb{N}_+$ and $Z = (H, K) \in \mathcal{Z}$. Suppose that $H_k = H$. Then

$$\gamma_{(i+1)k}(H_{(i+1)k}) = B(o, (i+1)k) \supset N_{ik}(\gamma_k(K)) \supset \gamma_{ik}(H_{ik}) = B(o, ik),$$

where, as before $N_r(U)$ denotes the r -neighborhood of U . Let Y denote the isomorphism type of the triple $(N_{ki}(\gamma_k(K)), \gamma_k(K), o)$. Then clearly Y can be determined from $Z, H_{(i+1)k}$ and $\beta_{(i+1)k}$.

Moreover, Y determines the isomorphism type of $(H_{ik}, \gamma_{ik}^{-1} \circ \gamma_k(K), o_{ik})$ and therefore determines $\hat{p}_{ik}(Z)$. Consequently, the three term sequence

$$\hat{p}_{ik}(Z), 1_{\{H_k=H\}} \mathbf{P}[A_Z \mid Y], \hat{p}_{(i+1)k}(Z) \quad (3.8)$$

is a martingale. We abbreviate

$$\bar{p}_{ik}(Z) := 1_{\{H_k=H\}} \mathbf{P}[A_Z \mid Y].$$

It is easy to see that

$$\bar{p}_{ik}(Z) = \begin{cases} p_{ik}(\gamma_k(K)) & H = H_k, \\ 0 & H \neq H_k. \end{cases}$$

Therefore,

$$q_{ik}(o) = \sum_{Z \in \mathcal{Z}} \bar{p}_{ik}(Z). \quad (3.9)$$

Since the three random variables $\hat{p}_{ik}(Z)$, $\bar{p}_{ik}(Z)$, $\hat{p}_{(i+1)k}(Z)$ form a martingale, we get from (3.6) that

$$\begin{aligned} \mathbf{E}\left[\left(\hat{p}_{(i+1)k}(Z) - \hat{p}_{ik}(Z)\right)^2\right] &= \mathbf{E}\left[\left(\hat{p}_{(i+1)k}(Z) - \bar{p}_{ik}(Z)\right)^2\right] + \mathbf{E}\left[\left(\bar{p}_{ik}(Z) - \hat{p}_{ik}(Z)\right)^2\right] \\ &\geq \mathbf{E}\left[\left(\bar{p}_{ik}(Z) - \hat{p}_{ik}(Z)\right)^2\right]. \end{aligned} \quad (3.10)$$

Recall the definition of $\mathcal{K}(v)$ from just below (3.1). There is clearly a finite upper bound $t = t(k, d)$ on $|\mathcal{K}(v)|$. In fact, by Lemma 3.2, we have

$$t \leq d^{2k}. \quad (3.11)$$

Observe that

$$|\{(H, K) \in \mathcal{Z} : H_k = H\}| = |\mathcal{K}(o)| \leq t.$$

Thus, there are at most t different $Z \in \mathcal{Z}$ for which $\bar{p}_{ik}(Z) - \hat{p}_{ik}(Z) \neq 0$. Cauchy-Schwarz therefore gives

$$t \sum_{Z \in \mathcal{Z}} \left(\bar{p}_{ik}(Z) - \hat{p}_{ik}(Z)\right)^2 \geq \left(\sum_{Z \in \mathcal{Z}} \bar{p}_{ik}(Z) - \sum_{Z \in \mathcal{Z}} \hat{p}_{ik}(Z)\right)^2. \quad (3.12)$$

Taking expectation on both sides and using linearity of expectation we get

$$t \sum_{Z \in \mathcal{Z}} \mathbf{E}\left[\left(\bar{p}_{ik}(Z) - \hat{p}_{ik}(Z)\right)^2\right] \geq \mathbf{E}\left[\left(\sum_{Z \in \mathcal{Z}} \bar{p}_{ik}(Z) - \sum_{Z \in \mathcal{Z}} \hat{p}_{ik}(Z)\right)^2\right]. \quad (3.13)$$

We now sum (3.10) over all $Z \in \mathcal{Z}$ and apply (3.13) and then (3.5) and (3.9), to get

$$\begin{aligned} t \sum_{Z \in \mathcal{Z}} \mathbf{E} \left[\left(\hat{p}_{(i+1)k}(Z) - \hat{p}_{ik}(Z) \right)^2 \right] &\geq \mathbf{E} \left[\left(\sum_{Z \in \mathcal{Z}} \bar{p}_{ik}(Z) - \sum_{Z \in \mathcal{Z}} \hat{p}_{ik}(Z) \right)^2 \right] \\ &= \mathbf{E} \left[(q_{ik}(o) - 1)^2 \right]. \end{aligned}$$

Now sum over i and apply (3.7), to obtain

$$t \geq \sum_{i=1}^{j-1} \mathbf{E} \left[(q_{ik}(o) - 1)^2 \right] = |V(G)|^{-1} \sum_{v \in V(G)} \sum_{i=1}^{j-1} (q_{ik}(v) - 1)^2. \quad (3.14)$$

Set $j = 8dt/\epsilon$ and observe that (3.14) implies that it cannot be the case that for every $1 \leq i \leq j-1$ we have more than $\epsilon|V(G)|/2d$ vertices satisfying $q_{ik}(v) \leq 1/2$. Therefore, recalling 3.11, we can take R_1 in the statement of the lemma to be

$$R_1(d, \epsilon, k) = kj \leq 10k d^{2k+1}/\epsilon, \quad (3.15)$$

thus completing the proof. \square

4 Concluding Remarks and Open Problems

1. Let us give an explicit upper bound for the dependence on ϵ of the number of queries required to test the property of being H -minor free for any fixed connected graph H . We assume here that the maximum degree d is a constant and H is fixed. Given a graph G and $\epsilon > 0$, let us choose ϵ_0 and k as in the proof of Theorem 1.2. It is clear that $1/\epsilon_0 = \text{poly}(1/\epsilon)$ and the result of [2] together with [32, Theorem 3] give that $k = \text{poly}(1/\epsilon)$. By Theorem 2.2 and its proof we know that if G is not $(\epsilon/2, k)$ -hyper-finite then for every G' that is H -minor-free we have $\rho_R(G, G') = \Omega(\epsilon)$ for some fixed $R = 2^{O(k)}/\epsilon = 2^{\text{poly}(1/\epsilon)}$. Here we are using the bounds that appear in (3.3) and (3.15). The algorithm first tries to distinguish between the case that the input is (ϵ_0, k) -hyper-finite from the case that it is not $(\epsilon/2, k)$ -hyper-finite. By the proof of Proposition 2.1 this can be done with $\text{poly}(h/\epsilon)$ queries, where h is the number of graphs of bounded degree d and radius R . Clearly $h = 2^{2^{O(R)}} = 2^{2^{\text{poly}(1/\epsilon)}}$. If the algorithm finds that G is not (ϵ_0, k) -hyper-finite it rejects G . Otherwise, the algorithm samples $\text{poly}(1/\epsilon)$ vertices, and for each vertex it explores its neighborhood of radius k . This requires $2^{O(k)} = 2^{\text{poly}(1/\epsilon)}$ queries. If any of the neighborhoods explored is not H -minor free the algorithm rejects, otherwise it accepts. The algorithm clearly accepts G with high probability if it is H -minor free, as such a G is (ϵ_0, k) -hyper-finite. If G is ϵ -far from being H -minor free and not $(\epsilon/2, k)$ -hyper-finite then it is rejected with high probability in the first step. If not then it is easy to see that $\Omega(\epsilon n)$ of its vertices belong to connected subgraphs of G of radius at most k that are not H -minor free. Hence, G is rejected in the second step.

Having handled the case of connected H , observe that if H is not connected, then after testing hyper-finiteness, we can test for being H' -minor free for each connected component H' of H . It is then easy to determine if G is H -minor free or not. Similar reasoning applies in the case of a property determined by a finite list of forbidden minors. By the Graph-Minor Theorem [37], this includes all minor closed properties. Thus, every minor closed property can be tested using $2^{2^{\text{poly}(1/\epsilon)}}$ queries.

2. Another model of property testing is when the number of edges is arbitrary, and the error is relative to the number of edges [35]; that is, a graph with m edges is ϵ -far from \mathcal{P} if we have to modify ϵm of its edges to get a graph satisfying \mathcal{P} . This raises the following:

Problem 4.1. *What is the query complexity of testing minor-closed properties in graphs with bounded average degree?*

It is clear that $o(n^{1/2})$ queries are not enough, for they will not suffice to distinguish between an edgeless graph and a clique on $|V(G)|^{1/2}$ of the vertices of the graph.

3. As we have argued in Subsection 1.1, minor closed properties are trivially testable in the dense graph model using $O(1/\epsilon)$ queries even with one-sided error. The intuition is that if a hereditary graph property (every minor closed property is also hereditary) is testable, then the reason must be that one can find a “proof” that a graph does not satisfy the property, in the form of a subgraph that does not satisfy it (which implies that the graph does not satisfy it). This intuition turns out to be correct in dense graphs [3]. However, it is easy to see that in the bounded degree model, we cannot expect to have such small proofs. For example, if G is a bounded degree expander of girth $\Omega(\log n)$, then for any fixed H with at least one cycle, we have that G is $\Omega(1)$ -far from being H -minor free¹⁰ when it is sufficiently large, but on the other hand, every subgraph of G of size $o(\log n)$ is a tree, and in particular H -minor free. Thus if H is not a tree, then one cannot test H -minor freeness with $o(\log n)$ queries and one-sided error. In fact, a much stronger $\Omega(\sqrt{n})$ lower bound can be deduced by adapting an argument from [21]. We raise the following conjecture, stating that the $\Omega(\sqrt{n})$ lower bound is tight.

Conjecture 4.2. *For every H , being H -minor free can be tested in the bounded degree setting with one-sided error and query complexity $\tilde{O}(\sqrt{n})$.*

If the conjecture is true, then the Graph Minor Theorem [37] implies that the same is true for any minor-closed graph property.

¹⁰This follows, for example, from the separator theorem of [2] for H -minor free graphs, mentioned earlier.

4. As we have briefly mentioned earlier, it does not seem like our approach can lead to testing algorithms for the property of being H -minor free whose query complexity is polynomial in $1/\epsilon$. It seems interesting to further investigate the possibility of coming up with such a tester.
5. Our main result here is that we can distinguish in constant time between graphs satisfying a minor closed property, from those that are far from satisfying it. Can we also estimate in constant time the fraction of edges that need to be removed in order to make the graph satisfy the property?
6. It would be interesting to obtain an explicit description of the difference between the frequencies of local neighborhoods that one sees in (ϵ, k) -hyper-finite graphs versus graphs that are not (ϵ', k) -hyper-finite, where $\epsilon \ll \epsilon'$.

Acknowledgements: We would like to thank David Wilson and Gábor Elek for helpful discussions, and an anonymous referee for many helpful comments.

References

- [1] N. Alon, E. Fischer, I. Newman and A. Shapira, A combinatorial characterization of the testable graph properties: it's all about regularity, *SIAM J. on Computing*, 39 (2009), 143–167
- [2] N. Alon, P. D. Seymour and R. Thomas, A separator theorem for non-planar graphs, *J. Amer. Math. Soc.* 3 (1990), 801–808.
- [3] N. Alon and A. Shapira, A characterization of the (natural) graph properties testable with one-sided error, *SIAM J. on Computing* 37 (2008), 1703–1727
- [4] N. Alon and A. Shapira, Homomorphisms in graph property testing, in *Topics in discrete mathematics*, 281–313, Springer, Berlin, 2006.
- [5] N. Alon and A. Shapira, A separation theorem in property testing, *Combinatorica*, 28 (2008), 261–281
- [6] M. Blum, M. Luby and R. Rubinfeld, Self-testing/correcting with applications to numerical problems, *JCSS* 47 (1993), 549–595.
- [7] H. L. Bodlaender, Discovering treewidth, *Lecture Notes in Computer Science*, 3381 (2005), 1–16.
- [8] A. Bogdanov, K. Obata, and L. Trevisan, A lower bound for testing 3-colorability in bounded-degree graphs, *FOCS 2002*, 93–102.

- [9] C. Borgs, J. Chayes, L. Lovász, V. T. Sos, B. Szegedy, and K. Vesztergombi, Graph limits and parameter testing, *STOC 2006*, 261–270.
- [10] A. Czumaj, A. Shapira, and C. Sohler, Testing hereditary properties of non-expanding bounded-degree graphs, *SIAM J. on Computing*, 38 (2009), 2499–2510.
- [11] A. Czumaj and C. Sohler, Sublinear-time algorithms, *Bulletin of the EATCS*, 89 (2006) 23–47.
- [12] A. Czumaj and C. Sohler, On testable properties in bounded degree graphs, *Proc. of SODA 2007*, 494–501.
- [13] E. Demaine, M. Hajiaghayi and K. Kawarabayashi, Algorithmic graph minor theory: decomposition, approximation, and coloring, *Proc. of FOCS 2005*, 637-646.
- [14] R. Diestel, **Graph Theory** (Third Edition), Springer, Heidenberg, 2005.
- [15] C. Domingo and J. Shawe-Taylor, The Complexity of Learning Minor-Closed Graph Classes, *Proc. of Algorithmic Learning Theory 1995*, 249–260.
- [16] R. Downey and M. Fellows, **Parameterized Complexity**, Springer, 1999.
- [17] G. Elek, The combinatorial cost, *L’Enseignement Mathématique*, 53 (2007), 225-236.
- [18] G. Elek, Parameter testing with bounded degree graphs of subexponential growth, *Random Structures and Algorithms*, to appear.
- [19] E. Fischer, The art of uninformed decisions: A primer to property testing, *The Computational Complexity Column of The Bulletin of the European Association for Theoretical Computer Science* 75 (2001), 97–126.
- [20] O. Goldreich, S. Goldwasser and D. Ron, Property testing and its connection to learning and approximation, *J. of the ACM* 45(4): 653–750 (1998).
- [21] O. Goldreich and D. Ron, Property Testing in Bounded-Degree Graphs, *Algorithmica* 32 (2002), 302–343. Also, *Proc. of STOC’07*.
- [22] O. Goldreich and D. Ron, A sublinear bipartiteness tester for bounded degree graphs, *Combinatorica*, 19 (1999), 335–373.
- [23] O. Goldreich and D. Ron, On testing expansion in bounded-degree graphs, *ECCC report TR00-020*, 2000.
- [24] O. Goldreich and L. Trevisan, Three theorems regarding testing graph properties, *Random Structures and Algorithms*, 23(1):23–57, 2003.

- [25] A. Hassidim, J. A. Kelner, H. N. Nguyen and K. Onak, Local Graph Partitions for Approximation and Testing, Proc. of FOCS 2009, to appear.
- [26] J. E. Hopcroft and R. E. Tarjan, Efficient planarity testing, J. of the ACM, 21 (1974), 549–568.
- [27] S. Kale and C. Seshadhri, Testing expansion in bounded-degree graphs, Proc. of ICALP 2008, 527–538.
- [28] A. Kostochka, The minimum Hadwiger number for graphs with a given mean degree of vertices, Metody Diskret. Analliz. 38 (1982), 37–58 [in Russian].
- [29] A. Kostochka, A lower bound for the Hadwiger number of graphs by their average degree, Combinatorica, 4 (1984), 307–316.
- [30] K. Kuratowski, Sur le problème des courbes gauches en topologie, Fund. Math. 15: 271–283.
- [31] R. J. Lipton and R. E. Tarjan, A separator theorem for planar graphs, SIAM J. Appl. Math., 36 (1979), 177–189.
- [32] R. J. Lipton and R. E. Tarjan, Applications of a planar separator theorem, SIAM J. on Computing 9 (1980), 615–627.
- [33] L. Lovász, Graph minor theory, Bull. Amer. Math. Soc. 43 (2006), no. 1, 75–86.
- [34] A. Nachmias and A. Shapira, Testing the expansion of a graph, ECCO report TR07-118, 2007.
- [35] M. Parnas and D. Ron, Testing the diameter of graphs, Random structures and algorithms, 20 (2002), 165–183.
- [36] N. Robertson and P.D. Seymour: Graph minors XIII. The disjoint paths problem, J. Combin. Theory Ser. B 63 (1995), 65–110.
- [37] N. Robertson and P. D. Seymour, Graph minors. XX. Wagner’s Conjecture, J. Combin. Theory Ser. B 92 (2004), 325–357.
- [38] R. Rubinfeld, Sublinear time algorithms, in *International Congress of Mathematicians. Vol. III*, 1095–1110, Eur. Math. Soc., Zürich, 2006.
- [39] D. Ron, Property testing, in: P. M. Pardalos, S. Rajasekaran, J. Reif and J. D. P. Rolim, editors, *Handbook of Randomized Computing*, Vol. II, Kluwer Academic Publishers, 2001, 597–649.
- [40] R. Rubinfeld and M. Sudan, Robust characterization of polynomials with applications to program testing, SIAM J. on Computing, 25 (1996), 252–271.

- [41] O. Schramm, Graph sequences with hyperfinite limits are hyperfinite, *Electron. Res. Announc. Math. Sci.* Vol. 15 (2008) 17-23.
- [42] E. Szemerédi, Regular partitions of graphs, In: *Proc. Colloque Inter. CNRS* (J. C. Bermond, J. C. Fournier, M. Las Vergnas and D. Sotteau, eds.), 1978, 399–401.
- [43] A. Thomason, An extremal function for contractions of graphs, *Math. Proc. Cambridge Philos. Soc.* 95 (1984), 261–265.
- [44] A. Thomason, The extremal function for complete minors, *J. of Combinatorial Theory Series B*, 81 (2001), 318–338.
- [45] K. Wagner, Über eine eigenschaft der ebenen komplexe, *Math. Ann.* 114 (1937), 570–590.