Every Monotone Graph Property is Testable[*]

Noga Alon [†]   Asaf Shapira [‡]

Abstract

A graph property is called *monotone* if it is closed under removal of edges and vertices. Many monotone graph properties are some of the most well-studied properties in graph theory, and the abstract family of all monotone graph properties was also extensively studied. Our main result in this paper is that any monotone graph property can be tested with one-sided error, and with query complexity depending only on ϵ. This result unifies several previous results in the area of property testing, and also implies the testability of well-studied graph properties that were previously not known to be testable. At the heart of the proof is an application of a variant of Szemerédi's Regularity Lemma. The main ideas behind this application may be useful in characterizing all testable graph properties, and in generally studying graph property testing.

As a byproduct of our techniques we also obtain additional results in graph theory and property testing, which are of independent interest. One of these results is that the query complexity of testing testable graph properties with one-sided error may be arbitrarily large. Another result, which significantly extends previous results in extremal graph-theory, is that for any monotone graph property $\mathcal{P}$, any graph that is ϵ-far from satisfying $\mathcal{P}$, contains a subgraph of size depending on ϵ only, which does not satisfy $\mathcal{P}$. Finally, we prove the following compactness statement: If a graph G is ϵ-far from satisfying a (possibly infinite) set of monotone graph properties $\mathcal{P}$, then it is at least $\delta_{\mathcal{P}}(\epsilon)$-far from satisfying one of the properties.

1 Introduction

1.1 Definitions and Background

All graphs considered here are finite, undirected, and have neither loops nor parallel edges. Let $\mathcal{P}$ be a property of graphs, namely, a family of graphs closed under isomorphism. All graph properties discussed in this paper are assumed to be decidable, that is, we disregard properties for which it is not possible to tell whether a given graph satisfies them. A graph G with n vertices is said to be *ϵ-far from satisfying* $\mathcal{P}$ if one must add or delete at least ϵn^2 edges in order to turn G into a graph satisfying $\mathcal{P}$. A *tester* for $\mathcal{P}$ is a randomized algorithm which, given the quantity n and the

ability to query whether a desired pair of vertices of an input graph G with n vertices are adjacent or not, distinguishes with high probability (say, $2/3$), between the case of G satisfying $\mathcal{P}$ and the case of G being ϵ-far from satisfying $\mathcal{P}$. One of the striking results in the area of property-testing is that many natural graph properties have a tester, whose total number of queries is bounded only by a function of ϵ, which is independent of the size of the input graph. A property having such a tester is called *testable*. Note, that if the number of queries performed by the tester is bounded by a function of ϵ only, then so is its running time. A tester is said to have *one-sided error* if whenever G satisfies $\mathcal{P}$, the algorithm declares that this is the case with probability 1. Throughout the paper, we assume that a tester first samples a set of vertices S, queries all the pairs $(i,j) \in S$, and then accepts or rejects by considering the graph spanned by the set. As observed in [3] and formally proved in [23], this can be assumed with no loss of generality, as this assumption at most squares the query complexity (and we will not care about such factors in this paper).

The general notion of property testing was first formulated by Rubinfeld and Sudan [33], who were motivated mainly by its connection to the study of program checking. The study of the notion of testability for combinatorial structures, and mainly for labelled graphs, was introduced in the seminal paper of Goldreich, Goldwasser and Ron [22], who showed that several natural graph properties are testable. In the wake of [22], many other graph properties were shown to be testable, while others were shown to be non-testable. See [18], [32] and [31] for additional results and references on graph property-testing as well as on testing properties of other combinatorial structures.

1.2 Related Work

The most interesting results in property-testing are those that show that large families of problems are testable. The main result of [22] states that a certain abstract graph partition problem, which includes as a special case k-colorability, having a large cut and having a large clique, is testable. The authors of [23] gave a characterization of the partition problems discussed in [22] that are testable with one-sided error. In [3], a logical characterization of a family of testable graph properties was obtained. According to this characterization, every first order graph-property of type $\exists\forall$ is testable, while there are first-order graph properties of type $\forall\exists$ that are not testable. These results were extended in [17].

There are also several general testability and non-testability results in other areas besides testing graph properties. In [4] it is proved that every regular language is testable. This result was extended to any read-once branching program in [29]. On the other hand, it was proved in [20], that there are read-twice branching programs that are not-testable. The main result of [6] states that any constraint satisfaction problem is testable.

With this abundance of general testability results, a natural question is what makes a combinatorial property testable. As graphs are the most well studied combinatorial structures in the theory of computation, it is natural to consider the problem of characterizing the testable graph properties, as the most important open problem in the area of property testing. Regretfully, though, finding such a characterization seems to be a very challenging endeavor, which is still open. Therefore, a natural line of research is to find large families of testable graph properties.

1.3 The Main New Result

Our main goal in this paper is to show that all the graph properties that belong to a large, natural and well studied family of graph properties are testable. In fact, we even show that these properties are testable with one-sided error. A graph-property $\mathcal{P}$ is said to be *monotone* if it is closed under removal of edges and vertices. In other words, if a graph G does not satisfy $\mathcal{P}$, then any graph that contains G as a (not necessarily induced) subgraph does not satisfy $\mathcal{P}$ as well. Various monotone graph properties were extensively studied in graph theory. As examples of monotone properties one can consider the property of having a homomorphism to a fixed graph H (which includes as a special case the property of being k-colorable, see Definition 2.2), and the property of not containing a (not necessarily induced) copy of some fixed graph H. Another set of well studied monotone properties are those defined by having a *fractional chromatic number*, *vector chromatic number* and *Lovász theta function* bounded by some constant c, which need not be an integer (see [25] and [27]). Another monotone property is being $(k, \mathcal{H})$-*Ramsey*: For a (possibly infinite) family of graphs $\mathcal{H}$, a graph is said to be $(k, \mathcal{H})$-Ramsey if one can color its edges using k colors, such that no color class contains a copy of a graph $H \in \mathcal{H}$. This property is the main focus of Ramsey-Theory, see [24] and its references. As another example, one can consider the property of being $(k, \mathcal{H}, f)$-Multicolorable; For a (possibly infinite) family of graphs $\mathcal{H}$ and a function f from $\mathcal{H}$ to the positive integers, a graph is said to be $(k, \mathcal{H}, f)$-Multicolorable if one can color its edges using k colors, such that every copy of a graph $H \in \mathcal{H}$ receives at least $f(H)$ colors. See [16], [14] and their references for a discussion of some special cases. The abstract family of monotone graph properties has also been extensively studied in graph theory. See [21], [13], [11] and their references. Our main result is the following:

Theorem 1 (Main Result) *Every monotone graph property is testable with one-sided error.*

We stress that we actually prove a slightly weaker statement than the one given above, as the monotone property has to satisfy some technical conditions (which cannot be avoided). However, as the cases where the actual result is weaker than what is stated in Theorem 1 deal with extremely unnatural properties, and even in these cases the actual result is roughly the same, we postpone the precise statement to Section 5 (see Theorem 6). Another important note is that in [23], Goldreich and Trevisan define a monotone graph property to be one that is closed under removal of edges, and not necessarily under removal of vertices. They show that there are such properties that are not testable even with two sided error. In fact, their result is stronger as the property they define belongs to NP and requires query complexity $\Omega(n^2)$. This means that Theorem 1 cannot be extended, in a strong sense, to properties that are only closed under removal of edges.

As we have mentioned above, having a homomorphism to a fixed graph H, k-colorability and the property of not containing a copy of a fixed graph H, are monotone properties, and are thus testable with one-sided error by Theorem 1. These properties were known to be testable before, and as Theorem 1 applies to general monotone properties, the bounds it supplies for these properties are inferior compared to the ones proved by the ad-hoc arguments (see [5], [22], [23] and [7]). In Theorem 4 we prove that this is unavoidable. The main importance of Theorem 1 thus lies in its generality. However, as described in the beginning of this subsection, there are additional natural and well-studied monotone graph properties that prior to this work were not known to be testable, and we may thus use Theorem 1 to conclude that these properties are testable with one-sided error. We also believe that Theorem 1 and its proof may be an important step towards a combinatorial characterization of the graph properties that are testable with one-sided error. Another important

aspect of Theorem 1 is that it can be used to prove general results on graph property testing. Two examples are Theorems 4 and 5, which we describe in the next subsection. Another result appears in a related subsequent paper [8] and is discussed in Section 5. We believe that Theorem 1 will be useful for proving other consequences as well. See Section 7 for more details and possible natural lines of research suggested by the results of this paper.

1.4   Techniques and Additional Results

The first technical ingredient in the proof of Theorem 1 is the proof of an (almost) equivalent formulation of it. For a (possibly infinite) family of graphs $\mathcal{F}$ we say that a graph is $\mathcal{F}$-free if it contains no member from $\mathcal{F}$ as a (not necessarily induced) subgraph. Clearly, being $\mathcal{F}$-free is a monotone property. It is well known (see e.g. [2]) that for any *finite* family of graphs $\mathcal{F}$, the property of being $\mathcal{F}$-free is testable. This follows from a standard application of Szemerédi's Regularity Lemma. As we discuss in Section 2, this lemma is inadequate for obtaining a similar result for infinite families of graphs. The main technical step in the proof of Theorem 1 is the following theorem, which is the main technical contribution of this paper.

Theorem 2 *For every (possibly infinite) family of graphs $\mathcal{F}$, there are functions $N_{\mathcal{F}}(\epsilon)$ and $Q_{\mathcal{F}}(\epsilon)$ with the following properties: If G is a graph on $n \geq N_{\mathcal{F}}(\epsilon)$ vertices which is ϵ-far from being $\mathcal{F}$-free, then a random subset of $Q_{\mathcal{F}}(\epsilon)$ vertices of G spans a member of $\mathcal{F}$ with probability at least 2/3.*

Note that Theorem 2 immediately implies that for every family of graphs $\mathcal{F}$, the property of being $\mathcal{F}$-free is testable. In order to prove Theorem 2 we apply a strong version of the regularity lemma, proved by Alon, Fischer, Krivelevich and Szegedy [3]. We believe that our application of this lemma may be useful for attacking other problems. As a byproduct of our argument we obtain the following graph theoretic result.

Theorem 3 *For every monotone graph property $\mathcal{P}$, there is a function $W_{\mathcal{P}}(\epsilon)$ with the following property: If G is ϵ-far from satisfying $\mathcal{P}$, then G contains a subgraph of size at most $W_{\mathcal{P}}(\epsilon)$, which does not satisfy $\mathcal{P}$.*

The above theorem significantly extends a result of Rödl and Duke [30], conjectured by Erdős, which asserts that the above statement holds for the k-colorability property. Theorem 3 applies to any monotone property, and in particular to all the properties discussed in the beginning of the previous subsection.

As will become evident from the proof of Theorem 1 (which is based on Theorem 2), the upper bounds for testing a monotone property depend on the property being tested. In other words, what we prove is that for every property $\mathcal{P}$, there is a function $Q_{\mathcal{P}}(\epsilon)$ such that $\mathcal{P}$ can be tested with query complexity $Q_{\mathcal{P}}(\epsilon)$. A natural question one may ask, is if the dependency on the specific property being tested can be removed. We rule out this possibility by proving the following.

Theorem 4 *For any function $Q : (0, 1) \mapsto \mathbb{N}$, there is a monotone graph property $\mathcal{P}$, which has no one-sided error property-tester with query-complexity bounded by $o(Q(\epsilon))$.*

Prior to this work, the best lower bound proved for testing a testable graph property with one-sided error was obtained in [1], where it is shown that for every non-bipartite graph H, the query

complexity of testing whether a graph does not contain a copy of H is at least $(1/\epsilon)^{\Omega(\log 1/\epsilon)}$. The fact that for every H this property is testable with one-sided error, follows from [2] and [3], and also as a special case from Theorem 1. As by Theorem 1 every monotone graph property is testable with one-sided error, Theorem 4 establishes that the one-sided error query complexity of testing testable graph properties, even those that are testable with one-sided error, may be *arbitrarily large*.

Our next result can be considered a compactness-type result in property testing. Suppose $\mathcal{P}_1, \ldots, \mathcal{P}_k$ are k graph properties that are closed under removal of edges. It is clear that if a graph G is ϵ-far from satisfying these k properties then it is at least ϵ/k-far from satisfying at least one of them. However, it is not clear that there is a fixed $\delta > 0$ such that even if $k \to \infty$, G must be δ-far from satisfying one of these properties. By using Theorem 2 we can prove that if these properties are monotone then such an δ exists. We also show that in general there is no such δ.

Theorem 5 *For any (possibly infinite) set of monotone graph properties $\mathcal{P} = \{\mathcal{P}_1, \mathcal{P}_2, \ldots\}$, there is a function $\delta_{\mathcal{P}} : (0, 1) \mapsto (0, 1)$ with the following property: If a graph G is ϵ-far from satisfying all the properties of $\mathcal{P}$, then for some i, the graph G is $\delta_{\mathcal{P}}(\epsilon)$-far from satisfying $\mathcal{P}_i$. Furthermore, there are properties $\mathcal{P} = \{\mathcal{P}_1, \mathcal{P}_2, \ldots\}$, which are closed under removal of edges, for which no such $\delta_{\mathcal{P}}$ exists.*

1.5 Recent results

By applying the techniques of this paper along with several additional ideas we have managed to extend Theorem 1 by showing that any hereditary graph property is testable with one-sided error (a graph property is hereditary if it is closed under removal of vertices, and not necessarily under removal of edges). Besides implying that many additional graph properties are testable, we can also use this result to obtain a precise characterization of the graph properties, which can be tested with one-sided error by testers with a certain natural restriction (essentially all the testers that have been designed thus far in the literature satisfy this restriction). These results, which appear in a subsequent paper [9], demonstrate the relevance of the techniques developed in this paper to the problem of characterizing the testable graph properties. Also, in joint work with Benny Sudakov [10], we have obtained approximation algorithms for the edit distance of a given graph from satisfying an arbitrary monotone graph property. We also obtained nearly matching hardness of approximation results. Some of the results of [10] also apply the main technique developed in this paper. In a recent paper, Avart et al. [12] extended our main result to the case of 3-uniform hypergraphs, which applies the main ideas of this paper along with the hypergraph regularity lemma. Finally, Lovász and Szegedy, [28] recently gave an alterative proof of the main result of [9] on testing hereditary properties using the notion of convergent graph sequences.

1.6 Organization

The rest of the paper is organized as follows. In Section 2 we introduce the basic notions of regularity and state the regularity lemmas that we use and some of their standard consequences. We also (do our best to) explain why the standard regularity lemma and its applications seem inadequate for proving Theorem 2. In Section 3 we give a high level description of the proof of Theorem 2 as well as the main ideas behind it. The full proof of Theorem 2 appears in Section 4. In Section 5 we give the precise statement of Theorem 1 and use Theorem 2 in order to prove it. In Section 7, we describe several possible extensions and open problems that this paper suggests. The proofs of

Theorems 3 and 5 appear in Section 4 and the proof of Theorem 4 appears in Section 6. Throughout the paper, whenever we relate, for example, to a function $f_{3.1}$, we mean the function f defined in Lemma/Claim/Theorem 3.1.

2   Regularity Lemmas: Definitions, Statements and Applications

In this section we discuss the basic notions of regularity, some of the basic applications of regular partitions and state the regularity lemmas that we use in the proof of Theorem 2. For a comprehensive survey on the regularity lemma the reader is refereed to [26]. We start with some basic definitions. For every two nonempty disjoint vertex sets A and B of a graph G, we define $e(A, B)$ to be the number of edges of G between A and B. The *edge density* of the pair is defined by $d(A, B) = e(A, B)/|A||B|$.

Definition 2.1 (γ-regular pair) *A pair (A, B) is γ-regular, if for any two subsets $A' \subseteq A$ and $B' \subseteq B$, satisfying $|A'| \geq \gamma|A|$ and $|B'| \geq \gamma|B|$, the inequality $|d(A', B') - d(A, B)| \leq \gamma$ holds.*

Note that a sufficiently large random bipartite graph, where each edge is chosen independently with probability d, is very likely to be a γ-regular pair with density roughly d, for any $\gamma > 0$. Thus, in some sense, the smaller γ is, the closer a γ-regular pair is to looking like a random bipartite graph. For this reason, the reader who is unfamiliar with the regularity lemma and its applications, should try and compare the statements given in this section to analogous statements about random graphs. Throughout the paper we will make an extensive use of the notion of graph homomorphism, which we turn to formally define.

Definition 2.2 (Homomorphism) *A homomorphism from a graph F to a graph K, is a mapping $\varphi : V(F) \mapsto V(K)$ that maps edges to edges, namely $(v, u) \in E(F)$ implies $(\varphi(v), \varphi(u)) \in E(K)$.*

Observe, that a graph F has a homomorphism into the complete graph of size k if and only if F is k-colorable. In what follows, $F \mapsto K$ denotes the fact that there is a homomorphism from F to K. Let F be a graph on f vertices and K a graph on k vertices, and suppose $F \mapsto K$. Let G be a graph obtained by taking a copy of K, replacing every vertex with a sufficiently large independent set, and every edge with a random bipartite graph of edge density d. It is easy to show that with high probability, G contains many copies of F. The following lemma shows that in order to infer that G contains many copies of F, it is enough to replace every edge with a "regular enough" pair. Intuitively, the larger f and k are, and the sparser the regular pairs are, the more regular we need each pair to be, because we need the graph to be "closer" to a random graph. This is formulated in Lemma 2.3 below. Several versions of this lemma were previously proved in papers using the regularity lemma. See, e.g., [26]. The reader should think of the mapping φ in the statement of the lemma as defining the homomorphism from F to the (implicit) graph K.

Lemma 2.3 *For every real $0 < \eta < 1$, and integers $k, f \geq 1$ there exist $\gamma = \gamma_{2.3}(\eta, k, f)$, $\delta = \delta_{2.3}(\eta, k, f)$ and $M = M_{2.3}(\eta, k, f)$ with the following property. Let F be any graph on f vertices, and let $U_1, \ldots, U_k$ be k pairwise disjoint sets of vertices in a graph G, where $|U_1| = \ldots = |U_k| = m \geq M$. Suppose there is a mapping $\varphi : V(F) \mapsto \{1, \ldots, k\}$ such that the following holds: If (i, j) is an edge of F then $(U_{\varphi(i)}, U_{\varphi(j)})$ is γ-regular with density at least η. Then, the sets $U_1, \ldots, U_k$ span at least δm^f copies of F.*

Comment 2.4 *Note, that the functions $\gamma_{2.3}(\eta, k, f)$ and $\delta_{2.3}(\eta, k, f)$ may and will be assumed to be monotone non-increasing in k and f. Similarly, we will assume that the function $M_{2.3}(\eta, k, f)$ is monotone non-decreasing in k and f. Also, for ease of future definitions (in particular the one given in (4)) we set $\gamma_{2.3}(\eta, k, 0) = \delta_{2.3}(\eta, k, 0) = M_{2.3}(\eta, k, 0) = 1$ for any $k \geq 1$ and $0 < \eta < 1$.*

A partition $\mathcal{A} = \{V_i \mid 1 \leq i \leq k\}$ of the vertex set of a graph is called an *equipartition* if $|V_i|$ and $|V_j|$ differ by no more than 1 for all $1 \leq i < j \leq k$ (so in particular each V_i has one of two possible sizes). When we refer to the size of such an equipartition, we mean the number of partition classes of the equipartition (k above). The Regularity Lemma of Szemerédi can be formulated as follows.

Lemma 2.5 ([34]) *For every m and $\gamma > 0$ there exists a number $T = T_{2.5}(m, \gamma)$ with the following property: Any graph G on $n \geq T$ vertices, has an equipartition $\mathcal{A} = \{V_i \mid 1 \leq i \leq k\}$ of $V(G)$ with $m \leq k \leq T$, for which all pairs (V_i, V_j), but at most $\gamma\binom{k}{2}$ of them, are γ-regular.*

The original formulation of the lemma allows also for an exceptional set with up to γn vertices outside of this equipartition, but one can first apply the original formulation with a somewhat smaller parameter instead of γ and then evenly distribute the exceptional vertices among the sets of the partition to obtain this formulation. $T_{2.5}(m, \gamma)$ may and is assumed to be monotone nondecreasing in m and monotone non-increasing in γ.

A standard application of Lemmas 2.3 and 2.5 shows that for any *finite* set of graphs $\mathcal{F}$, the property of not containing a member of $\mathcal{F}$, that is being $\mathcal{F}$-free, is testable. We first use Lemma 2.3 by setting f and k to be the size of the largest graph in $\mathcal{F}$ and letting $\eta = \epsilon$. Lemma 2.3 gives a $\gamma_{2.3}$, which tells us how regular an equipartition should be (that is, how small should γ be) in order to find many copies of a member of $\mathcal{F}$ in it, assuming the input graph is ϵ-far from being $\mathcal{F}$-free. We then apply Lemma 2.5, with $\gamma = \gamma_{2.3}$. The main difficulty with applying this strategy when $\mathcal{F}$ is infinite is that we do not know a priori the size of the member of $\mathcal{F}$ that we will eventually find in the equipartition that Lemma 2.5 returns. After finding $F \in \mathcal{F}$ in an equipartition, we may find out that F is too large for Lemma 2.3 to be applied, because Lemma 2.5 was not used with a small enough γ. One may then try to find a new equipartition based on the size of F. However, that requires using a smaller γ, and thus the new equipartition may be larger (that is, contain more partition classes), and thus contain only larger members of $\mathcal{F}$. Hence, even the new γ is not good enough in order to apply Lemma 2.3. This leads to a circular definition of constants, which seems unbreakable. Our main tool in the proof of Theorem 2 is Lemma 2.7 below, proved in [3] for a different reason, which enables us to break this circular chain of definitions. This lemma can be considered a variant of the standard regularity lemma, where one can use a function that defines γ as a function of the size of the equipartition[1], rather than having to use a fixed γ as in Lemma 2.5. To state the Lemma we need the following definition.

Definition 2.6 (The function $W_{\mathcal{E}, m}$) *Let $\mathcal{E}(r) : \mathbb{N} \mapsto (0, 1)$ be an arbitrary monotone non-increasing function. Let also m be an arbitrary positive integer. We define the function $W_{\mathcal{E}, m} : \mathbb{N} \mapsto (0, 1)$ inductively as follows: $W_{\mathcal{E}, m}(1) = T_{2.5}(m, \mathcal{E}(0))$. For any integer $i > 1$ put $R = W_{\mathcal{E}, m}(i - 1)$ and define*

$$W_{\mathcal{E}, m}(i) = T_{2.5}(R, \mathcal{E}(R)/R^2). \tag{1}$$

---

[1] This is a simplification of the actual statement, see item (3) in the statement of Lemma 2.7

Lemma 2.7 ([3]) *For every integer m and monotone non-increasing function $\mathcal{E}(r) : \mathbb{N} \mapsto (0,1)$ define*

$$S = S_{2.7}(m, \mathcal{E}) = W_{\mathcal{E},m}(100/\mathcal{E}(0)^4).$$

For any graph G on $n \geq S$ vertices, there exist an equipartition $\mathcal{A} = \{V_i \mid 1 \leq i \leq k\}$ of $V(G)$ and an induced subgraph U of G, with an equipartition $\mathcal{B} = \{U_i \mid 1 \leq i \leq k\}$ of the vertices of U, that satisfy:

1. *$m \leq k \leq S$.*

2. *$U_i \subseteq V_i$ for all $i \geq 1$, and $|U_i| \geq n/S$.*

3. *In the equipartition $\mathcal{B}$, **all** pairs are $\mathcal{E}(k)$-regular.*

4. *All but at most $\mathcal{E}(0)\binom{k}{2}$ of the pairs $1 \leq i < j \leq k$ are such that $|d(V_i, V_j) - d(U_i, U_j)| < \mathcal{E}(0)$.*

Comment 2.8 *For technical reasons (see the proof in [3]), Lemma 2.7 requires that for any $r > 0$ the function $\mathcal{E}(r)$ will satisfy*

$$\mathcal{E}(r) \leq \min\{\mathcal{E}(0)/4, 1/4r^2\}. \tag{2}$$

One of the difficulties in the proof of Theorem 2, is in showing that all the constants that are used in the course of the proof can be upper bounded by functions depending on ϵ only. The following observation will thus be useful.

Proposition 2.9 *If m is bounded by a function of ϵ only and $\mathcal{E}(r)$ satisfies (2), then the integer $S = S_{2.7}(m, \mathcal{E})$ can be upper bounded by a function of ϵ only[2].*

The dependency of the function $T_{2.5}(m, \gamma)$ on γ is a tower of exponents of height polynomial in $1/\gamma$ (see the proof in [26]). Thus, even for moderate functions $\mathcal{E}$ the integer S has a huge dependency on ϵ, which is a tower of towers of exponents of height polynomial in $1/\epsilon$.

3 Overview of the Proof of Theorem 2

Though we believe that the proof of Theorem 2 is not harder than several other proofs applying the regularity lemma, we could not avoid the usage of a hefty number of constants that may hide the main ideas of the proof. We thus give in this section a general overview of the proof, and the way we overcome the difficulties described in Section 2. The complete proof is given in Section 4.

For an equipartition of a graph G, let the *regularity graph* of G, denoted $R = R(G)$, be the following graph: We first use Lemma 2.5 in order to obtain the equipartition satisfying the assertions of the lemma. Let k be the size of the equipartition. Then, R is a graph on k vertices, where vertices i and j are connected if and only if (V_i, V_j) is a dense regular pair (with the appropriate parameters). In some sense, the regularity graph is an approximation of the original graph, up to γn^2 modifications. One of the main (implicit) implications of the regularity lemma is the following: Suppose we consider two graphs to be *similar* if their regularity graphs are identical. It thus follows from Lemma 2.5 that for every $\gamma > 0$, the number of graphs that are pairwise non-similar is bounded by a function of γ only

---

[2]In our application of Lemma 2.7 the function $\mathcal{E}$ will (implicitly) depend on ϵ. For example, it will be convenient to set $\mathcal{E}(0) = \epsilon$. However, note that even in this case $S_{2.7}(m, \mathcal{E})$ can be upper bounded by a function of ϵ only.

$(2^{\binom{T}{2}}$, where $T = T_{2.5}(1/\gamma, \gamma))$. Namely, up to γn^2 modifications, all the graphs can be approximated using a set of equipartitions of size bounded by a function of γ only. The reader is referred to [15] where this interpretation of the regularity lemma is also (implicitly) used. This leads us to the key definitions of the proof of Theorem 2. The reader should think of the graphs R considered below as the set of regularity graphs discussed above, and the parameter r as representing the size of R.

Definition 3.1 (The family $\mathcal{F}_r$) *For any (possibly infinite) family of graphs $\mathcal{F}$, and any integer r let $\mathcal{F}_r$ be the following set of graphs: A graph R belongs to $\mathcal{F}_r$ if it has at most r vertices and there is at least one $F \in \mathcal{F}$ such that $F \mapsto R$.*

Practicing definitions, observe that if $\mathcal{F}$ is the family of odd cycles, then $\mathcal{F}_r$ is precisely the family of non-bipartite graphs of size at most r. In the proof of Theorem 2, the set $\mathcal{F}_r$, defined above, will represent a subset of the regularity graphs of size at most r. Namely, those R for which there is at least one $F \in \mathcal{F}$ such that $F \mapsto R$. As r will be a function of ϵ only, and thus finite, we can take the maximum over all the graphs $R \in \mathcal{F}_r$, of the size of the smallest $F \in \mathcal{F}$ such that $F \mapsto R$. We thus define

Definition 3.2 (The function $\Psi_{\mathcal{F}}$) *For any family of graphs $\mathcal{F}$ and integer r for which $\mathcal{F}_r \neq \emptyset$, define*

$$\Psi_{\mathcal{F}}(r) = \max_{R \in \mathcal{F}_r} \quad \min_{\{F \in \mathcal{F} : F \mapsto R\}} |V(F)|. \tag{3}$$

Define $\Psi_{\mathcal{F}}(r) = 0$ if $\mathcal{F}_r = \emptyset$. Therefore, $\Psi_{\mathcal{F}}(r)$ is monotone non-decreasing in r.

Practicing definitions again, note that if $\mathcal{F}$ is the family of odd cycles, then $\Psi_{\mathcal{F}}(r) = r$ when r is odd, and $\Psi_{\mathcal{F}}(r) = r - 1$ when r is even. The "right" way to think of the function $\Psi_{\mathcal{F}}$ is the following: Let R be a graph of size at most r and suppose we are guaranteed that there is a graph $F' \in \mathcal{F}$ such that $F' \mapsto R$ (thus $R \in \mathcal{F}_r$). Then by this information only and *without* having to know the structure of R itself, the definition of $\Psi_{\mathcal{F}}$ implies that there is a graph $F \in \mathcal{F}$ of size at most $\Psi_{\mathcal{F}}(r)$, such that $F \mapsto R$.

The function $\Psi_{\mathcal{F}}$ has a critical role in the proof of Theorem 2. The first usage of this function is that as by Lemma 2.5 we can upper bound the size of the regularity graph R, we can also upper bound the size of the smallest graph $F \in \mathcal{F}$ for which $F \mapsto R$. A second important property of $\Psi_{\mathcal{F}}$ is discussed in Section 5. A natural question one may ask is whether there is a function Ψ that can upper bound $\Psi_{\mathcal{F}}$ for all families $\mathcal{F}$. As it turns out, this is impossible, namely the dependency on the specific family $\mathcal{F}$ is unavoidable. See the discussion following the proof of Theorem 4 in Section 6. As we have mentioned in the previous section, the main difficulty that prevents one from proving Theorem 2 using Lemma 2.3 is that one does not know a priori the size of the graph that one may expect to find in the equipartition. This leads us to define the following function where $0 < \epsilon < 1$ is an arbitrary real.

$$\mathcal{E}'(r) = \begin{cases} \epsilon/8, & r = 0 \\ \gamma_{2.3}(\epsilon/8, \ r, \ \Psi_{\mathcal{F}}(r)), & r \geq 1 \end{cases} \tag{4}$$

In simple words, given r, which will represent the size of the equipartition and thus also the size of the regularity graph which it defines, $\mathcal{E}'(r)$ returns "how regular" this equipartition should be in order to allow one to find many copies of the *largest* graph one may possibly have to work with. Note,

that we obtain the upper bound on the size of this largest possible graph, by invoking $\Psi_{\mathcal{F}}(r)$. As for different families of graphs $\mathcal{F}$, the function $\Psi_{\mathcal{F}}(r)$ may behave differently, $\mathcal{E}'(r)$ may also behave differently for different families $\mathcal{F}$, as it is defined in terms of $\Psi_{\mathcal{F}}(r)$. However, and this is one of the key points of the proof, as we are fixing the family of graphs $\mathcal{F}$, the function $\mathcal{E}'(r)$ depends only on r.

Given the above definitions we apply Lemma 2.7 with a slight modification of $\mathcal{E}'(r)$ in order to obtain an equipartition of G. We then throw away edges that reside inside the sets V_i and between (V_i, V_j), whose edge density differs significantly from that of (U_i, U_j). We then argue that we thus throw away less than ϵn^2 edges. As G is by assumption ϵ-far from not containing a member of $\mathcal{F}$, the new graph still contains a copy of $F \in \mathcal{F}$. By the definition of the new graph, it thus means that there is a (natural) homomorphism from F to the regularity graph of G. We then arrive at the main step of the proof, where we use the key property of Lemma 2.7, item (3), and the definition of $\mathcal{E}'(r)$ to get that the sets U_i are regular enough to let us use Lemma 2.3 on them and to infer that they span many copies of F. It thus follows, that a large enough sample of vertices spans a copy of F with high probability. The complete details appear in Section 4.

4   Proofs of Theorems 2, 3 and 5

We start with the proof of Theorem 2. We assume the reader is familiar with the overview of its proof given in Section 3.

Proof of Theorem 2: Fix any family of graphs $\mathcal{F}$. Our goal is to show the existence of functions $N_{\mathcal{F}}(\epsilon)$ and $Q_{\mathcal{F}}(\epsilon)$ with the following properties: If a graph G on $n \geq N_{\mathcal{F}}(\epsilon)$ vertices is ϵ-far from being $\mathcal{F}$-free, then a random subset of $Q_{\mathcal{F}}(\epsilon)$ vertices of $V(G)$ spans a member of $\mathcal{F}$ with probability at least $2/3$. For the rest of the proof, let $\mathcal{E}'(r) : N \mapsto (0,1)$ be as defined in (4). In order to apply Lemma 2.7, we need to define a function $\mathcal{E}$, based on $\mathcal{E}'$, which will satisfy the technical condition (2) in Comment 2.8. We thus set $\mathcal{E}(0) = \mathcal{E}'(0)$ ($= \epsilon/8$) and define for any $r > 0$,

$$\mathcal{E}(r) = \min\{\mathcal{E}'(r),\ \mathcal{E}(0)/4,\ 1/4r^2\}. \tag{5}$$

For the rest of the proof set

$$S(\epsilon) = S_{2.7}(8/\epsilon, \mathcal{E}).$$

We may indeed define $S(\epsilon)$ using $\mathcal{E}$ as it satisfies (2). Furthermore, as we define $S(\epsilon)$ using $m = 8/\epsilon$ we get by Proposition 2.9 that $S(\epsilon)$ is indeed a function of ϵ only. We now set

$$N = N_{\mathcal{F}}(\epsilon) = S(\epsilon) \cdot M_{2.3}(\epsilon/8, S(\epsilon), \Psi_{\mathcal{F}}(S(\epsilon))) \tag{6}$$

to be an integer bounded by a function of ϵ as well. We postpone the definition of $Q_{\mathcal{F}}(\epsilon)$ till the end of the proof.

Given a graph G on n vertices, with $n \geq N \geq S(\epsilon)$, we can use Lemma 2.7 with $m = 8/\epsilon$ and $\mathcal{E}(r)$ as defined in (5), in order to obtain an equipartition of $V(G)$ into $8/\epsilon \leq k \leq S(\epsilon)$ clusters $V_1, \ldots, V_k$ (this is possible by item (1) in Lemma 2.7). By item (2) of Lemma 2.7, for every $1 \leq i \leq k$ we have sets $U_i \subseteq V_i$ each of size at least $n/S(\epsilon)$. Remove from G the following edges according to the following order:

1. Any edge (u, v) for which both u and v belong to the same cluster V_i. As each of the clusters contains at most $n/k + 1$ vertices, the total number of edges removed is at most $k(n/k)^2$. As $k \geq 8/\epsilon$ we have $k(n/k)^2 < \frac{\epsilon}{8}n^2$.

2. If for some $i < j$ we have $|d(V_i, V_j) - d(U_i, U_j)| > \frac{\epsilon}{8} = \mathcal{E}(0)$, remove all the edges connecting vertices that belong to V_i to vertices that belong to V_j. By item (4) of Lemma 2.7, there are at most $\frac{\epsilon}{8}k^2$ such pairs i, j. As V_i and V_j contain at most $(n/k + 1)$ vertices, we remove at most $\frac{\epsilon}{8}k^2 \cdot (n/k + 1)^2 \leq \frac{\epsilon}{7}n^2$ edges in this step.

3. If for some $i < j$ we have $d(U_i, U_j) < \frac{\epsilon}{8}$, remove all the edges connecting vertices that belong to V_i to vertices that belong to V_j. As we have already removed in the previous step all the edges between pairs (V_i, V_j) for which $|d(V_i, V_j) - d(U_i, U_j)| > \frac{\epsilon}{8}$, we may conclude that if $d(U_i, U_j) < \frac{\epsilon}{8}$ then we also have $d(V_i, V_j) < \frac{\epsilon}{8} + \mathcal{E}(0) = \frac{\epsilon}{4}$. As V_i and V_j contain at most $(n/k + 1)$ vertices, we thus remove at most $k^2 \cdot \frac{\epsilon}{4}(n/k + 1)^2 \leq \frac{\epsilon}{3}n^2$ edges.

Call the graph obtained after removing the above edges G', and observe that G' is obtained from G by removing less than ϵn^2 edges. By item (3) of Lemma 2.7, in G all the pairs (U_i, U_j) are $\mathcal{E}(k)$-regular. Thus, by the third step of obtaining G' we get the following property:

Proposition 4.1 *If $v_i \in V_i$ is connected to $v_j \in V_j$ in G', then (U_i, U_j) is a $\mathcal{E}(k)$-regular pair with density at least $\frac{\epsilon}{8}$ in G.*

Consider a graph R on k vertices $r_1, \ldots, r_k$, where vertices r_i and r_j are connected if and only if (U_i, U_j) is an $\mathcal{E}(k)$-regular pair in G with density at least $\frac{\epsilon}{8}$. This is the regularity graph, which we have mentioned in Section 3, of the graph induced by the sets $U_1, \ldots, U_k$. As G is by assumption ϵ-far from being $\mathcal{F}$-free, and G' is obtained from G by removing less than ϵn^2 edges, G' must contain a copy of a graph $F' \in \mathcal{F}$. Let R_i contain all the vertices of F' that belong to cluster V_i and note that by Proposition 4.1, there is a natural homomorphism $\varphi : V(F') \mapsto V(R)$ which maps all the vertices of $R_i \subseteq V(F')$ to r_i. As $|V(R)| = k$ and F' is a graph in $\mathcal{F}$ such that $F' \mapsto R$, we conclude that $R \in \mathcal{F}_k$ (recall Definition 3.1). Therefore, there is a graph $F \in \mathcal{F}$ of size at most $\Psi_{\mathcal{F}}(k)$ such that $V(F) \mapsto V(R)$ (recall Definition 3.2). Let $\varphi : V(F) \mapsto V(R)$ be the homomorphism mapping the vertices of F to the vertices of R. By definition, we have that whenever (i, j) is an edge of F their image $(\varphi(i), \varphi(j))$ is an edge of R. Furthermore, by definition of R we know that if $(\varphi(i), \varphi(j))$ is an edge of R then $(U_{\varphi(i)}, U_{\varphi(j)})$ is an $\mathcal{E}(k)$-regular pair with density at least $\frac{\epsilon}{8}$.

We have thus arrived at the following situation: We have k clusters of vertices $U_1, \ldots, U_k$ of the same size. We also have a graph F of size at most $\Psi_{\mathcal{F}}(k)$, and a mapping $\varphi : V(H) \mapsto \{1, \ldots, k\}$ that satisfies the condition; if $(i, j) \in E(F)$ then $(U_{\varphi(i)}, U_{\varphi(j)})$ is an $\mathcal{E}(k)$-regular pair with density $\epsilon/8$. This, together with the definition of $\mathcal{E}(k)$, implies that we can use Lemma 2.3 on the graph U spanned by $U_1, \ldots, U_k$. Let $f \leq \Psi_{\mathcal{F}}(k)$ denote the size of F. Item (4) in Lemma 2.7 states that each U_i contains at least $n/S(\epsilon)$ vertices. Also, by (6), and by the monotonicity properties of $M_{2.3}$ discussed in Comment 2.4, we have for any $1 \leq i \leq k$

$$|U_i| \geq n/S(\epsilon) \geq M_{2.3}(\epsilon/8, S(\epsilon), \Psi_{\mathcal{F}}(S(\epsilon))) \geq M_{2.3}(\epsilon/8, k, \Psi_{\mathcal{F}}(k)).$$

Therefore, we may apply Lemma 2.3 on the sets $U_1, \ldots, U_k$ to conclude that U spans at least

$$\delta \prod_{i=1}^{f} |U_i| \geq \delta(n/S(\epsilon))^f \geq \delta n^f / S(\epsilon)^{\Psi_{\mathcal{F}}(k)} \geq \delta n^f / S(\epsilon)^{\Psi_{\mathcal{F}}(S(\epsilon))} \tag{7}$$

copies of F, where $\delta = \delta_{2.3}(\epsilon/8, k, \Psi_{\mathcal{F}}(k))$. By Comment 2.4, the function $\delta_{2.3}(\eta, k, f)$ is monotone non-increasing in k and f. Also, $\Psi_{\mathcal{F}}(k)$ is monotone nondecreasing in k. Hence, as $k \leq S(\epsilon)$ we have that $\delta \geq \delta_{2.3}(\epsilon/8, S(\epsilon), \Psi_{\mathcal{F}}(S(\epsilon))))$, and in particular $1/\delta$ is upper bounded by a function of ϵ only. As U is a subgraph of G, we may conclude that G contains at least as many copies of F as (7). Thus, if we independently sample $2S(\epsilon)^{\Psi_{\mathcal{F}}(S(\epsilon))}/\delta$ sets of $\Psi_{\mathcal{F}}(S(\epsilon))$ ($\geq f$) vertices (which is a total of $2\Psi_{\mathcal{F}}(S(\epsilon)) \cdot S(\epsilon)^{\Psi_{\mathcal{F}}(S(\epsilon))}/\delta$ vertices) we have probability at least $2/3$ of finding a copy of $F \in \mathcal{F}$.

We can now give the formal definition of $Q_{\mathcal{F}}(\epsilon)$. Given a family of graphs $\mathcal{F}$ let $\Psi_{\mathcal{F}}(r)$ be the function from Definition 3.2. We note that the only place where $Q_{\mathcal{F}}(\epsilon)$ depends on $\mathcal{F}$ is in the function $\Psi_{\mathcal{F}}(r)$. Using $\Psi_{\mathcal{F}}(r)$ define the function $\mathcal{E}(r)$ as in (5). Given $\epsilon > 0$ define the function $W_{\mathcal{E}, 8/\epsilon}$ as in Definition 2.6 and put $S(\epsilon) = W_{\mathcal{E}, 8/\epsilon}(100/(\epsilon/8)^4)$. Finally, we can set

$$Q_{\mathcal{F}}(\epsilon) = \frac{2\Psi_{\mathcal{F}}(S(\epsilon)) \cdot S(\epsilon)^{\Psi_{\mathcal{F}}(S(\epsilon))}}{\delta_{2.3}(\epsilon/8, S(\epsilon), \Psi_{\mathcal{F}}(S(\epsilon)))} \tag{8}$$

to be a function of ϵ only. This completes the proof of the theorem. ∎

From the definition of $\mathcal{E}'(r)$ in (4) it is clear that if the function $\Psi_{\mathcal{F}}(r)$ is recursive, then so is $\mathcal{E}'(r)$ and therefore also $\mathcal{E}(r)$ (for this we also need the fact that $\gamma_{2.3}(\eta, k, f)$ is recursive, which follows from the standard proofs of Lemma 2.3, see [26]). In this case the function $W_{\mathcal{E}, m}(i)$ is also recursive (see Definition 2.6), and therefore also the function $S_{2.7}(8/\epsilon, \mathcal{E})$. Finally, this means that the integer $S(\epsilon)$, used in the above proof, can also be computed. Now, given $S(\epsilon)$ and the fact that $\Psi_{\mathcal{F}}(r)$ is recursive, one can use (6) and (8) as well as the fact that $\delta_{2.3}(\eta, k, f)$ and $M_{2.3}(\eta, k, f)$ are recursive (see the proof in [26]) in order to compute $N_{\mathcal{F}}(\epsilon)$ and $Q_{\mathcal{F}}(\epsilon)$.

We finish this section with the proofs of Theorems 3 and 5.

Proof of Theorem 3: We claim that we can set $W_{\mathcal{P}}(\epsilon) = \max\{N_{\mathcal{F}}(\epsilon), Q_{\mathcal{F}}(\epsilon)\}$ with $\mathcal{F} = \mathcal{F}_{\mathcal{P}}$ as in the proof of Theorem 1, and $N_{\mathcal{F}}(\epsilon)$, $Q_{\mathcal{F}}(\epsilon)$ the functions from Theorem 2. Indeed, If G is ϵ-far from satisfying $\mathcal{P}$, and G has less than $N_{\mathcal{F}}(\epsilon)$ vertices, we can take G itself to be a subgraph of G not satisfying $\mathcal{P}$. Suppose now that G has more than $N_{\mathcal{F}}(\epsilon)$ vertices. As G is also ϵ-far from being $\mathcal{F}$-free, we get from Theorem 2 that G contains a subgraph (in fact, many) of size $Q_{\mathcal{F}}(\epsilon)$, which is not $\mathcal{F}$-free and therefore, does not satisfy $\mathcal{P}$. ∎

Proof of Theorem 5: For each of the monotone properties $\mathcal{P}_i$, let $\mathcal{F}_i$ be the family of graphs, which do not satisfy $\mathcal{P}_i$, and let $\mathcal{F} = \mathcal{F}_1 \bigcup \mathcal{F}_2 \bigcup \mathcal{F}_3 \bigcup \ldots$. Clearly, a graph G satisfies all the properties of $\mathcal{P}$ if and only if it is $\mathcal{F}$-free. Consider a graph G, which is ϵ-far from satisfying all the properties of $\mathcal{P}$. In this case G is also ϵ-far from being $\mathcal{F}$-free. The proof of Theorem 2 establishes that there is a graph $F \in \mathcal{F}$ of size at most $f = f_{\mathcal{F}}(\epsilon)$ such that G contains $\delta_{\mathcal{F}}(\epsilon)n^f$ copies of F. Note, that removing an edge from G destroys at most $\binom{n}{f-2} \leq n^{f-2}$ copies of F. Thus, one must remove at least $\delta_{\mathcal{F}}(\epsilon)n^2$ edges from G in order to make it F-free. Let i be such that $F \in \mathcal{F}_i$. We may now infer that G is $\delta_{\mathcal{F}}(\epsilon)$-far from satisfying $\mathcal{P}_i$. Finally, note that as $\mathcal{F}$ is determined by $\mathcal{P}$, we can also say that G is $\delta_{\mathcal{P}}(\epsilon)$-far from satisfying $\mathcal{P}_i$.

To show that in case the properties $\mathcal{P}_i$ are just closed under removal of edges the above does not hold, consider the following: For any integer n, let $H_1, H_2, \ldots$ be some ordering of the graphs on n vertices, which contain precisely $n^{3/2}$ edges. A graph of size n is said to satisfy property $\mathcal{P}_i$ if it contains no copy of H_i. Clearly, any property $\mathcal{P}_i$ is closed under removal of edges, but not

necessarily under removal of vertices. Observe, that any graph with at least $n^{3/2}$ edges does not satisfy one of the properties $\mathcal{P}_i$. Therefore, any graph G of size n, which contains $2\epsilon n^2$ edges is ϵ-far from satisfying all the properties $\mathcal{P}_i$. We claim that any such G is not $\frac{\log n}{\sqrt{n}}$-far from satisfying any one of these properties. To this end, it is enough to show that for any graph H_i, we can remove at most $n^{3/2} \log n$ edges from G and thus make it H_i-free. To see this, note that as G and H_i are both of size n, G spans at most $n!$ copies of H_i. As H_i contains $n^{3/2}$ edges a randomly chosen edge of G is spanned by H_i with probability at least $n^{3/2}/\binom{n}{2} > 1/\sqrt{n}$. Thus, if we remove from G a set of $n^{3/2} \log n$ edges, were each edge is randomly and uniformly chosen from the edges of G (with repetitions), the probability that none of the edges of one of the copies of H_i in G were removed is at most $(1 - 1/\sqrt{n})^{n^{3/2} \log n} < 1/n!$. By the union bound, the probability that for *some* copy of H_i in G, none of its edges were removed is strictly smaller than 1. Thus, there exists a choice of $n^{3/2} \log n$ edges, whose removal from G makes it H_i-free. ∎

5   Proof of Theorem 1

For a monotone graph property $\mathcal{P}$, define $\mathcal{F} = \mathcal{F}_\mathcal{P}$ to be the set of graphs which are minimal with respect to not satisfying property $\mathcal{P}$. In other words, a graph F belongs to $\mathcal{F}$ if it does not satisfy $\mathcal{P}$, but any graph obtained from F by removing an edge or a vertex, satisfies $\mathcal{P}$. Thus, for example, if $\mathcal{P}$ is the property of being 2-colorable, then $\mathcal{F}$ is the set of odd-cycles. Clearly, a graph satisfies $\mathcal{P}$ if and only it contains no member of $\mathcal{F}$ as a (not necessarily induced) subgraph.

As we have mentioned in Section 1, we will prove a slightly different version of Theorem 1. In order to precisely restate Theorem 1 we need two definitions. Note, that in defining a tester in Section 1, we did not mention whether the error parameter ϵ is given as part of the input, or whether the tester is designed to distinguish between graphs that satisfy $\mathcal{P}$ from those that are ϵ-far from satisfying it, when ϵ is a known fixed constant. In fact, the literature about property testing is not clear about this issue as in some papers ϵ is assumed to be a part of the input while in others it is not. We define a property to be *uniformly* testable if there is a tester for it that receives ϵ as part of the input. We define a property to be *non-uniformly* testable if for every fixed ϵ, there is a tester that can distinguish between graphs that satisfy $\mathcal{P}$ from those ϵ-far from satisfying it (which may not work properly for other values of ϵ). We are now ready to restate Theorem 1.

Theorem 6 (Theorem 1 restated): *Every monotone graph property $\mathcal{P}$ is non-uniformly testable with one-sided error. Moreover, if the function $\Psi_\mathcal{F}$ is recursive (where $\mathcal{F} = \mathcal{F}_\mathcal{P}$) then $\mathcal{P}$ is also uniformly testable with one-sided error.*

We stress that all reasonable graph properties $\mathcal{P}$, in particular those that were discussed in Section 1, are such that $\Psi_\mathcal{F}$ is recursive (a function is recursive if there is an algorithm that computes it in finite time). In particular, all the monotone properties mentioned in Section 1 are uniformly testable with one-sided error. We thus bother to define uniformly and non-uniformly testing as well as discuss $\Psi_\mathcal{F}$ because it has the following interesting property: Not only is it sufficient to require $\Psi_\mathcal{F}$ to be recursive in order to infer that $\mathcal{P}$ can be tested uniformly with one-sided error, but this is also *necessary*. In other words, the recursiveness of $\Psi_\mathcal{F}$ determines whether $\mathcal{P}$ can be tested uniformly[3]. This is somewhat surprising as $\Psi_\mathcal{F}$ has little to do with property testing. Using this

---

[3]This is in fact a simplification of the actual result that we can show. See [8] for the precise statement.

necessary condition, it is possible to show that there are graph properties that can be *non-uniformly* tested with one-sided error, but cannot be *uniformly* tested, even with two-sided error. In fact, there are such graph properties, which are monotone and belong to *coNP*. The proofs of the necessity of $\Psi_\mathcal{F}$ being recursive in order to obtain a uniform tester, as well as the existence of a property that cannot be tested uniformly are rather involved and significantly deviate from the main topic of this paper. Hence, we refrain from describing them here. These results will appear in a subsequent paper [8].

Proof of Theorem 6: Let $\mathcal{F} = \mathcal{F}_\mathcal{P}$ be as defined above, and let $N_\mathcal{F}(\epsilon)$ and $Q_\mathcal{F}(\epsilon)$ be the functions of Theorem 2. As satisfying $\mathcal{P}$ is equivalent to being $\mathcal{F}$-free, we focus on testing the property of being $\mathcal{F}$-free. We first show that every monotone property is non-uniformly testable. In this case we may design a tester for every given error parameter ϵ (but one that can handle *any* graph as an input). In this case, for every fixed ϵ, the tester knows the values of $N_\mathcal{F}(\epsilon)$ and $Q_\mathcal{F}(\epsilon)$ in advance (i.e. they are part of its description). If the size of the input graph is less than $N_\mathcal{F}(\epsilon)$, the algorithm queries about all edges of the graph and accepts if and only if the graph is $\mathcal{F}$-free (obviously, in this case the algorithm always answers correctly). If the size of the input graph is larger than $N_\mathcal{F}(\epsilon)$, it samples $Q_\mathcal{F}(\epsilon)$ random vertices and accepts if and only if the graph spanned by this set of vertices is $\mathcal{F}$-free. Clearly, if G is $\mathcal{F}$-free the algorithm declares that this is the case with probability 1. On the other hand, if it is ϵ-far from being $\mathcal{F}$-free then by Theorem 2 the sample of size $Q_\mathcal{F}(\epsilon)$ will contain $F \in \mathcal{F}$ with probability at least 2/3, and thus the algorithm will reject the input with this probability. In any case, the query complexity, which is $\max\{N_\mathcal{F}(\epsilon), Q_\mathcal{F}(\epsilon)\}$, is bounded by a function of ϵ only.

We now turn to uniform testers. In this case, we can imitate the proof of the case where ϵ is given in advance, which was described above. The only technical obstacle that may prevent us from carrying out the same testing algorithm, is that the algorithm should be able to compute $N_\mathcal{F}(\epsilon)$ and $Q_\mathcal{F}(\epsilon)$. As the details of the proof of Theorem 2 reveal (see the discussion following the proof of Theorem 2 in Section 4), the only step in computing $N_\mathcal{F}(\epsilon)$ and $Q_\mathcal{F}(\epsilon)$, which is not well defined (i.e. that depends on $\mathcal{F}$) is the computation of the function $\Psi_\mathcal{F}(r)$ (see Definition 3.2). In other words, if $\Psi_\mathcal{F}$ is recursive, then so are $N_\mathcal{F}(\epsilon)$ and $Q_\mathcal{F}(\epsilon)$. We thus get that if $\Psi_\mathcal{F}$ is recursive, we can uniformly test the property of being $\mathcal{F}$-free. ∎

6 Proof of Theorem 4

In this section we describe the proof of Theorem 4. We remind the reader that we denote by $F \mapsto K$ the fact that there is a homomorphism from F to K (see Definition 2.2). In what follows, an s-blowup of a graph K is the graph obtained from K by replacing every vertex $v_i \in V(K)$ with an independent set I_i, of size s, and replacing every edge $(v_i, v_j) \in E(K)$ with a complete bipartite graph whose partition classes are I_i and I_j. It is easy to see that a blowup of K is far from being K-free (K-free is the property of not containing a copy of K). It is also easy to see that if $F \mapsto K$, then a blowup of K is far from being F-free (see [1] Lemma 3.3). However, in this case the farness of the blowup from being F-free is a function of the size of F. As it turns out, for the proof of Theorem 4 we need a stronger assertion where the farness is only a function of k. This is given in Lemma 6.1 below, which is proved in [8].

Lemma 6.1 ([8]) *Let F be a graph on f vertices with at least one edge, let K be a graph on k vertices, and suppose $F \mapsto K$ (thus, $k \geq 2$). Then, for every sufficiently large $n \geq n(f)$, an n/k-blowup of K, is $\frac{1}{2k^2}$-far from being F-free.*

For the proof of Theorem 4 we also need the following simple observation. Recall that by the results of [3] and [23] we may assume that a tester first picks as set of vertices S, and then accepts or rejects according to the graph induced by S. We also comment on the following subtle point: when defining a testable graph property we require the query complexity to be *upper bounded* by a function of ϵ. This, however, does not imply that it is *independent* of n. Therefore, when considering arbitrary testers we have to consider testers whose query complexity and decisions may depend on the size of the input, which we denote by n.

Claim 6.2 *Let $\mathcal{F}$ be a family of graphs, such that no $F \in \mathcal{F}$ has isolated vertices and let T be a one-sided error tester for the property of being $\mathcal{F}$-free with query complexity $Q(\epsilon, n)$. If for some $\epsilon_0 > 0$ and n, after T samples a set of vertices S of size $Q(\epsilon_0, n)$, the graph induced by S is $\mathcal{F}$-free, then T must accept the input.*

Proof: Fix any n and $\epsilon_0 > 0$ and suppose that when we execute T on a graph G of size n with $\epsilon = \epsilon_0$, and the sample of vertices S spans a graph H that is $\mathcal{F}$-free, the algorithm still rejects the input. Consider a graph G' of size n, which contains a graph isomorphic to H on some of its $|S|$ vertices. Other than the edges of the copy of H, the graph G' has no other edges. As $\mathcal{F}$ contains no graph with isolated vertices and H is $\mathcal{F}$-free we infer that G' is $\mathcal{F}$-free. Suppose we execute T on G' with $\epsilon = \epsilon_0$. As G and G' are of the same size n, when given G' as input the algorithm samples a set of vertices of size $Q(\epsilon_0, n) = |V(H)|$. As we assume that when given H the algorithm rejects, we get that there is a non-zero probability that T will reject G', contradicting the assumption that it has one-sided error. $\blacksquare$

As our goal is to prove a lower bound on the query complexity we may and will assume that Q is monotone non-increasing (hence, monotone non-decreasing in $1/\epsilon$). For every such function Q we will define a property $\mathcal{P} = \mathcal{P}(Q)$ needed in order to prove Theorem 4. These properties can be thought of as *sparse bipartiteness* as they will be defined in terms of not containing a certain subset of the set of odd-cycles.

Let $Q : (0,1) \mapsto \mathbb{N}$ be an arbitrary monotone non-increasing function. For such a function, let Q^i be the following i times iterated version of Q. We put $Q^1(x) = Q(x)$ and for any $i \geq 1$ define

$$Q^{i+1}(x) = 2Q\left(\frac{1}{2(Q^i(x) + 2)^2}\right) + 1. \tag{9}$$

Define $I(Q) = \{Q^i(1/2) : i \in N\}$ and note that $I(Q)$ contains only odd integers. For a function as above, let $C(Q) = \{C_i : i \in I(Q)\}$, that is $C(Q)$ is the set of odd cycles whose lengths are the integers of the set $I(Q)$. Finally, let $\mathcal{P} = \mathcal{P}(Q)$ denote the property of not containing any of the odd-cycles of $C(Q)$ as a (not necessarily induced) subgraph.

Proof of Theorem 4: Given a monotone non-increasing function Q, let $\mathcal{P} = \mathcal{P}(Q)$ be the property defined above. We show that for any positive integer k for which $k - 2 \in I(Q)$, any one-sided error tester that distinguishes between graphs of sufficiently large n that satisfy $\mathcal{P}$ from those that are

$\frac{1}{2k^2}$-far from satisfying it, has query complexity at least $Q(1/2k^2)$. As Q is by assumption monotone non-increasing, $I(Q)$ contains infinitely many integers. Hence, for infinitely many values of ϵ, and for all large enough n, the query complexity of such a one-sided error tester is at least $Q(\epsilon)$. Note also that the set of these ϵ's approaches zero.

Fix any integer k for which $k-2 \in I(Q)$ and assume $k-2 = Q^i(1/2)$. As $I(Q)$ contains only odd integers, k is also odd. Define $\ell = Q^{i+1}(1/2)$ and recall that by (9), we have $\ell = 2Q(1/2k^2) + 1$. As it is clear that there is a homomorphism from C_ℓ to C_k, we get by Lemma 6.1 that for any $n \geq N(\ell)$, an n/k-blowup of C_k is $\frac{1}{2k^2}$-far from being C_ℓ-free. Denote such a blowup by G. As by definition $C_\ell \in C(Q)$, the graph G is also $\frac{1}{2k^2}$-far from satisfying $\mathcal{P}$. Also, as $k-2$ is odd, G contains no copy of C_{k-2}. In particular, G contains no member of $C(Q)$ of length less than ℓ. As property $\mathcal{P}$ is determined in terms of not containing a certain set of odd cycles, none of which has isolated vertices, we get from Claim 6.2 that a one-sided error tester must find a copy of a graph not satisfying $\mathcal{P}$, in order to determine that it does not satisfy $\mathcal{P}$. Therefore, for any $n \geq N(\ell)$ the query complexity of any tester for distinguishing between graphs of size n satisfying $\mathcal{P}$ from graphs of size n that are $\frac{1}{2k^2}$-far from satisfying it, is at least ℓ. As $\ell = 2Q(1/2k^2) + 1 \geq Q(1/2k^2)$ the proof is complete. $\blacksquare$

An immediate consequence of Theorem 4 is that there is no function $Q(\epsilon)$ that upper bounds the query complexity $Q_{\mathcal{F}}(\epsilon)$, of testing the property of being $\mathcal{F}$-free for all families of graphs, $\mathcal{F}$. In other words, the dependence on the specific family of graph is unavoidable. By the same reasoning, the dependence on $\mathcal{P}$ in Theorem 3 is also unavoidable. As we have commented after the proof of Theorem 2 in Section 4, the only dependence of the function $Q_{\mathcal{F}}(\epsilon)$ defined in the proof of Theorem 1 (see (8)), on $\mathcal{P}$ is due to the function $\Psi_{\mathcal{F}}$ from Definition 3.2 (where $\mathcal{F} = \mathcal{F}_{\mathcal{P}}$ is the set of minimal graphs with respect to not satisfying $\mathcal{F}$). This implies that the function $\Psi_{\mathcal{F}}$ must depend on $\mathcal{F}$ and thus also on $\mathcal{P}$, as otherwise we could obtain an upper bound on $Q_{\mathcal{F}}(\epsilon)$ which would apply to all families of graphs, thus contradicting Theorem 4. We conjecture that Theorem 4 can be extended to two-sided error testers, see Section 7.

As we have commented at the beginning of this section, the proof of Theorem 4 heavily relies on the fact that the farness of the graph considered in Lemma 6.1 from being F-free is only a function of k. From the proof of Theorem 4 it should indeed be clear that if this farness had been a function of the size of F, then the length of each cycle of the family would have depended on its own size, which would result in a cycle of definitions.

7 Concluding Remarks and Open Problems

- Besides proving that a large family of graph properties are all testable, and that specific properties that were previously not known to be testable are in fact testable, another important aspect of Theorem 1 is that it can be used to prove general results on testing graph properties. Two such results are Theorems 4 and 5. Another result, discussed in Section 5, is that there are graph properties that can be non-uniformly tested, but cannot be uniformly tested [8]. We believe that Theorem 1 will be useful for proving other results as well.

- Our main result gives that the natural family of monotone graph properties are all testable with one sided error. This gives rise to several questions. For example, one can study the relation between testing with one-sided and two-sided error by considering how large can be the gap between the query complexity of testing a monotone graph property with one-sided and

two-sided error. Specifically, it will be interesting to investigate, whether there is a monotone property, for which there is a super-polynomial gap between the two tasks. It will also be interesting to strengthen Theorem 4 by proving that for any function $Q : (0, 1) \mapsto \mathbb{N}$, there is a monotone graph property that cannot be tested with $o(Q(\epsilon))$ queries, even with two-sided error. Currently, the best lower bound on the *two-sided* error query complexity of a monotone graph property is a $(1/\epsilon)^{\Omega(\log 1/\epsilon)}$ lower bound for testing the property of not containing a copy of a graph H, for any non-bipartite H [7].

- A particularly interesting problem to study regarding the family of monotone graph properties is to obtain a characterization of the monotone properties, which are testable with $poly(1/\epsilon)$ queries. For some properties, such as k-colorability, it is known that $poly(1/\epsilon)$ queries suffice (see [22] and [5]). For others, such as being H-free for any non-bipartite H, it is known that $poly(1/\epsilon)$ are not sufficient (see [1] and [7]).

  Even a special case of this problem seems hard to resolve. While it is known that the property of not containing an odd cycle, namely being bipartite, can be tested with $\tilde{O}(1/\epsilon)$ queries (see [5]), Theorem 4 establishes that testing the property of being $\mathcal{F}$-free, where $\mathcal{F}$ is a subset of the family of odd-cycles, may be arbitrarily hard (at least with one-sided error). It is interesting to check if one can at least characterize the families of odd-cycles $\mathcal{F}$, for which one can test the property of being $\mathcal{F}$-free with $poly(1/\epsilon)$ queries.

- As was mentioned in the introduction, a result of Goldreich and Trevisan [23] rules out the possibility of extending Theorem 2 to graph properties that are only closed under removal of edges. It seems interesting to bridge the gap between their result and the main result of this paper by characterizing the testable graph properties that are closed under edge removal.

- Two graph properties $\mathcal{P}_1$ and $\mathcal{P}_2$ are defined in [3] to be *indistinguishable* if for every $\epsilon > 0$ and large enough n, any graph on n vertices satisfying one property is never ϵ-far from satisfying the other. It is shown in [3] that in this case, $\mathcal{P}_1$ is testable if and only if $\mathcal{P}_2$ is testable. It is first proved in [3] that ceratin colorability properties are testable with one-sided error. It is then shown that every first order graph property of type $\exists\forall$ is indistinguishable from some colorability property, thus obtaining that these properties are also testable. It would be interesting to characterize (either combinatorially, logically or by other means) the graph properties that are indistinguishable from some monotone property. By Theorem 1, this will immediately imply that these properties are testable, possibly with two-sided error.

- Fischer and Newman [19] have recently shown that if a graph property $\mathcal{P}$ is testable, then it is also estimable, that is, it is possible to estimate how far is a given graph from satisfying $\mathcal{P}$, within an error $\delta > 0$ in time depending only on δ. Combining Theorem 1 and the result of [19] gives that any monotone property is estimable. We further note, that this result (in fact, a stronger one) follows directly from the main result of [10], which was obtained independently of [19].

- The proof of Theorem 5 gives weak lower bounds for the function $\delta_{\mathcal{P}}(\epsilon)$. It may be interesting to check if this dependency can be linear or polynomial for some natural families $\mathcal{P}$.

References

[1] N. Alon, Testing subgraphs in large graphs, Proc. 42^{nd} IEEE FOCS, IEEE (2001), 434-441. Also: Random Structures and Algorithms 21 (2002), 359-370.

[2] N. Alon, R. A. Duke, H. Lefmann, V. Rödl and R. Yuster, The algorithmic aspects of the Regularity Lemma, Proc. of 33^{rd} IEEE FOCS, Pittsburgh, IEEE (1992), 473-481. Also: J. of Algorithms 16 (1994), 80-109.

[3] N. Alon, E. Fischer, M. Krivelevich and M. Szegedy, Efficient testing of large graphs, Proc. of 40^{th} FOCS, New York, NY, IEEE (1999), 656–666. Also: Combinatorica 20 (2000), 451-476.

[4] N. Alon, M. Krivelevich, I. Newman and M. Szegedy, Regular languages are testable with a constant number of queries, Proc. 40^{th} FOCS, New York, NY, IEEE (1999), 645–655. Also: SIAM J. on Computing 30 (2001), 1842-1862.

[5] N. Alon and M. Krivelevich, Testing k-colorability, SIAM J. Discrete Math., 15 (2002), 211-227.

[6] N. Alon and A. Shapira, Testing satisfiability, Proc. 13^{th} Annual ACM-SIAM SODA, ACM Press (2002), 645-654. Also: Journal of Algorithms, 47 (2003), 87-103.

[7] N. Alon and A. Shapira, Testing subgraphs in directed graphs, Proc. of the 35^{th} Annual Symp. on Theory of Computing (STOC), San Diego, California, 2003, 700–709. Also: JCSS 69 (2004), 354-382.

[8] N. Alon and A. Shapira, A separation theorem in property-testing, manuscript.

[9] N. Alon and A. Shapira, A characterization of the (natural) graph properties testable with one-sided error, Proc. of FOCS 2005, 429-438.

[10] N. Alon, A. Shapira and B. Sudakov, Additive approximation for edge-deletion problems, Proc. of FOCS 2005, 419-428.

[11] N. Alon and J. H. Spencer, **The Probabilistic Method**, Second Edition, Wiley, New York, 2000.

[12] C. Avart, V. Rödl, and M. Schacht, Every monotone 3-graph property is testable, submitted.

[13] J. Balogh, B. Bollobás and D. Weinreich, Measures on monotone properties of graphs, Discrete Applied Mathematics, to appear.

[14] D. Eichhorn and D. Mubayi, Edge-coloring cliques with many colors on subcliques, Combinatorica 20 (2000), 441-444.

[15] P. Erdős, P. Frankl and V. Rödl, The asymptotic number of graphs not containing a fixed subgraph and a problem for hypergraphs having no exponent, Graphs Combin. 2 (1986) 113-121.

[16] P. Erdős and A. Gyárfás, A variant of the classical Ramsey problem, Combinatorica 17 (1997), 459-467.

[17] E. Fischer, Testing graphs for colorability properties, Proc. of the 12^{th} SODA (2001), 873-882.

[18] 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.

[19] E. Fischer and I. Newman, Testing versus estimation of graph properties, Proc. of 37^{th} STOC, 2005, 138-146.

[20] E. Fischer, I. Newman and J. Sgall, Functions that have read-twice constant width branching programs are not necessarily testable, Random Structures and Algorithms, 24 (2004), 175-193.

[21] E. Friedgut and G. Kalai, Every monotone graph property has a sharp threshold. Proc. Amer. Math. Soc. 124 (1996), 2993-3002.

[22] O. Goldreich, S. Goldwasser and D. Ron, Property testing and its connection to learning and approximation, Proc. of 37^{th} Annual IEEE FOCS, (1996), 339–348. Also, JACM 45(4): 653-750 (1998).

[23] O. Goldreich and L. Trevisan, Three theorems regarding testing graph properties, Proc. 42^{nd} IEEE FOCS, IEEE (2001), 460-469. Also, Random Structures and Algorithms, 23(1):23-57, 2003.

[24] R. L. Graham, B. L. Rothschild and J. H. Spencer, **Ramsey Theory**, Second Edition, Wiley, New York, 1990.

[25] D. Karger, R. Motwani and M. Sudan, Approximate graph coloring by semidefinite programming, JACM 45(2), 1998, 246-265.

[26] J. Komlós and M. Simonovits, Szemerédi's Regularity Lemma and its applications in graph theory. In: *Combinatorics, Paul Erdös is Eighty*, Vol II (D. Miklós, V. T. Sós, T. Szönyi eds.), János Bolyai Math. Soc., Budapest (1996), 295–352.

[27] L. Lovász, On the shannon capacity of a graph, IEEE Transactions on Information Theory 25(1), 1979, 1-7.

[28] L. Lovász and B. Szegedy, Graph limits and testing hereditary graph properties, manuscript, 2005.

[29] I. Newman, Testing of functions that have small width branching programs, Proc. of 41^{th} FOCS (2000), 251-258.

[30] V. Rödl and R. Duke, On graphs with small subgraphs of large chromatic number, *Graphs and Combinatorics* 1 (1985), 91–96.

[31] 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.

[32] R. Rubinfeld, Sublinear time algorithms, to appear.

[33] R. Rubinfeld and M. Sudan, Robust characterization of polynomials with applications to program testing, *SIAM J. on Computing* 25 (1996), 252–271.

[34] 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.