Graphs with tiny vector chromatic numbers and huge chromatic numbers

Extended Abstract

Uriel Feige Michael Langberg Gideon Schechtman
Department of Computer Science and Applied Mathematics
Weizmann Institute of Science, Rehovot 76100
{feige,mikel,gideon}@wisdom.weizmann.ac.il

Abstract

Karger, Motwani and Sudan (JACM 1998) introduced the notion of a vector coloring of a graph. In particular they show that every k-colorable graph is also vector k-colorable, and that for constant k, graphs that are vector k-colorable can be colored by roughly $\Delta^{1-2/k}$ colors. Here Δ is the maximum degree in the graph. Their results play a major role in the best approximation algorithms for coloring and for maximal independent set.

We show that for every positive integer k there are graphs that are vector k-colorable but do not have independent sets significantly larger than $n/\Delta^{1-2/k}$ (and hence cannot be colored with significantly less that $\Delta^{1-2/k}$ colors). For $k = O(\log n/\log\log n)$ we show vector k-colorable graphs that do not have independent sets of size $(\log n)^c$, for some constant c. This shows that the vector chromatic number does not approximate the chromatic number within factors better than $n/polylogn$.

As part of our proof, we analyze "property testing" algorithms that distinguish between graphs that have an independent set of size n/k, and graphs that are "far" from having such an independent set. Our bounds on the sample size improve previous bounds of Goldreich, Goldwasser and Ron (JACM 1998) for this problem.

1. Introduction

An independent set in a graph G is a set of vertices that do not induce any edges. The size of the maximum independent set in G is denoted by $\alpha(G)$. For an integer k, a k coloring of G is a function $\sigma : V \rightarrow [1 \ldots k]$ which assigns colors to the vertices of G. A valid k coloring of G is a coloring in which each color class is an independent set. The chromatic number $\chi(G)$ of G is the smallest k for which there exists a valid k coloring of G.

Finding $\alpha(G)$ and $\chi(G)$ are fundamental NP-hard problems, closely related by the inequality $\alpha(G)\chi(G) \geq n$. Given G, the question of estimating the value of $\alpha(G)$ ($\chi(G)$) or finding large independent sets (small colorings) in G have been studied extensively. Let G be a graph of size n. Both $\chi(G)$ and $\alpha(G)$ can be approximated within a ratio of $O(\frac{n(\log\log n)^2}{\log^3 n})$ [10, 6]. It is known that unless NP=ZPP, neither $\alpha(G)$ nor $\chi(G)$ can be approximated within a ratio of $n^{1-\epsilon}$ for any $\epsilon > 0$ [12, 7]. Under stronger complexity assumptions, there is some $0 < \delta < 1$ such that neither problem can be approximated within a ratio of $n/2^{\log^\delta n}$ [14].

The approximation ratios for these problems significantly improve in graphs that have very large independent sets, or very small chromatic numbers. The algorithms achieving the best ratios in these cases [13, 1, 3, 11] are all based on the idea of vector coloring, introduced by Karger, Motwani and Sudan [13]. A vector k-coloring of a graph is an assignment of unit vectors to its vertices, such that for every edge, the inner product of the vectors assigned to its endpoints is at most (in the sense that it can only be more negative) $-1/(k-1)$. Every k-colorable graph is also vector k-colorable (by identifying each color class with one vertex of a perfect $(k-1)$-dimensional simplex centered at the origin). Moreover, unlike the chromatic number, a vector k-coloring (when it exists) can be found in polynomial time using semidefinite programming (up to arbitrarily small error in the inner products). Given a vector k-coloring of a graph, Karger, Motwani and Sudan show how to color a graph with roughly $\Delta^{1-2/k}$ colors, where Δ is the maximum degree in the graph. (In comparison, the technique of inductive coloring might use $\Delta+1$ colors.) In fact, they show how to find an independent set of size roughly $n/\Delta^{1-2/k}$. Combined with other ideas, this leads to coloring algorithms and algorithms for finding independent sets with the best currently known performance guarantees. For example, there is a polynomial time algorithm that colors 3-colorable graphs with roughly $n^{3/14}$ colors [3]. For nonconstant values of k, it is known how to find an independent set of size $\Omega(\log n)$ in a vector $\log n$-colorable graph.

There have also been negative results regarding vector k-colorable graphs. Examples appearing in [13] (and improved by Noga Alon and Mario Szegedy) show vector 3-colorable graphs that do not have independent sets larger than roughly $n^{0.95}$. The case of nonconstant k is addressed in [5] (technically, the results there deal with the Lovasz theta function, which is even a stronger notion than vector coloring), where graphs that are vector $2^{O(\sqrt{\log n})}$-colorable are shown not to have independent sets larger than $2^{O(\sqrt{\log n})}$. In all these negative examples, the vertex sets of the graphs involved can be viewed as a subset of $\{0,1\}^n$, with two vertices connected by an edge if their Hamming distance is larger than some prespecified value.

Our results. In this work we present a different family of graphs with stronger negative properties. For every constant k and every $\epsilon > 0$, we show graphs that are vector k-colorable, with $\alpha(G) \leq n/\Delta^{1-\frac{2}{k}-\epsilon}$. This essentially matches the positive results of [13]. As a function of n rather than Δ, we show vector 3-colorable graphs with $\alpha(G) < n^{0.843}$. Moreover, for $k = O(\log n/\log\log n)$, we show vector k-colorable graphs with $\alpha(G) \leq (\log n)^c$, for some universal c. This shows that the vector coloring number by itself does not approximate the chromatic number within a ratio better than $n/\text{polylog}n$. Another consequence of this (that is touched upon in Remark A.1 of the appendix) is that certain semidefinite programs do not approximate the size of the maximum independent set with a ratio better than $n/\text{polylog}n$.

Theorem 1.1

1. *For every constant $\varepsilon > 0$ and constant $k > 2$, there are infinitely many graphs G that are vector k-colorable and satisfy $\alpha(G) \leq \frac{n}{\Delta^{1-\frac{2}{k}-\varepsilon}}$, where n is the number of vertices in G, Δ is the maximal degree in G, and $\Delta > n^{\delta}$ for some constant $\delta > 0$.*

2. *For some constant c, there are infinitely many graphs G that are vector $O(\frac{\log n}{\log\log n})$-colorable and satisfy $\alpha(G) \leq (\log n)^c$.*

3. *There are infinitely many graphs G which are vector 3-colorable and satisfy $\alpha(G) \leq n^{0.843}$.*

Proof techniques. The graphs that we use are essentially the same graphs that were used in [8] to show integrality gaps for semidefinite programs for max cut. Namely, they are obtained by placing n points at random on a d-dimensional unit sphere, and connecting two points by an edge if the inner product of their respective vectors is below $-1/(k-1)$. Such graphs are necessarily vector k-colorable, as the embedding on the sphere is a vector k-coloring. So the bulk of the work is in proving that they have no large independent set. For this we use a three phase plan, similar to that employed in [8] for max cut. First we consider a continuous version of this graph, where every point on the sphere is a vertex. On this continuous graph (with infinitely many vertices and edges) we use certain symmetrization techniques in order to analyze its properties. Specifically, we prove certain inequalities regarding its expansion. In the second phase, we replace the continuous graph by a so-called dense graph that has finitely many vertices packed very densely on the sphere. We show that the dense graph approximates the continuous graph well, and hence shares the same expansion properties. In the third phase, we consider our finite graph to be a random vertex induced subgraph of the dense graph. Based on the expansion properties of the dense graph, we show that the sparse graph has no large independent set.

Two remarks are in order here. One is that it is very important for our bounds that the final graph does not contain too many vertices compared to the dimension d. A small number of vertices implies low degree, and allows for a more favorable relation between the maximum degree and the chromatic number. For this reason we cannot use the dense graph as is (it has a very large degree), and we subsample it. The other remark is that we do not get an explicit graph as our example, but rather a random graph (or a distribution on graphs). This is to some extent unavoidable, given that there are no known efficient deterministic constructions of Ramsey graphs (graphs in which the size of the maximum independent set and maximum clique are both bounded by some polylog in n). The graphs we construct (when $k = \log n/\log\log n$) are Ramsey graphs, because it can be shown that the maximum clique size is never larger than the vector coloring number.

Property testing. The following problem in property testing is addressed by Goldreich, Goldwasser and Ron [9]. For some value of $\rho < 1$, consider a graph with the following "promise": either it has an independent set of size ρn, or it is far from any such graph, in the sense that any vertex induced subgraph of ρn vertices induces at least ϵn^2 edges. One wants an algorithm that samples as few vertices as possible, looks at the subgraph induced on them, and based on the size of the maximum independent set in it decides correctly (with high probability) which of the two cases above hold. In [9] it is shown that a sample of size proportional to ϵ^{-4} suffices. We are in a somewhat similar situation when we move from the dense graph to our final graph. The dense graph is far from having an independent set of size ρn (where n is the number of vertices in the dense graph). We want to take a small as possible sample (its size will be denoted by s) such that the induced subgraph does not contain an independent set of size ρs. Luckily in our case, we can use a stronger guarantee on the dense graph. We know that even every set of $\rho n/2$ vertices induces at least ϵn^2 edges. In this case we show that s proportional to $1/\epsilon$ suffices. This dramatic improvement over the [9] bound is crucial to the success of our three phase plan. We note that this improvement is not only based on our stronger guarantee on the dense graph G. Even in exactly the same setting of [9], we show that s proportional to ϵ^{-3} suffices. These results are of interest in the context of property testing regardless of there applications to the vector coloring issue.

Strict vector coloring. One may strengthen the notion of vector k-coloring by requiring that for every edge, the inner product of the unit vectors corresponding to its endpoints be exactly $-1/(k-1)$, rather than at most $-1/(k-1)$. This is called a strict vector coloring. This notion is known to be equivalent to the theta function of Lovasz [15, 13]. Every k-colorable graph is also strictly vector k-colorable. As strict vector coloring is a stronger requirement than vector coloring, then potentially, strict vector k-colorable graphs have smaller chromatic numbers than vector k-colorable graphs. So far, there has not been any algorithmic technique that could use this observation to further improve the approximation ratios for chromatic number or for independent set. We remark however that the negative results in the current paper apply only to vector coloring and not to strict vector coloring. It is an open question whether similar negative results are true for strict vector coloring, or equivalently, whether the same gaps (such as $n/\text{polylog} n$) can be shown between the value of the theta function and the size of the maximum independent set. Note that the weaker negative results of [5] and some of the negative results in [13] do apply also to strict vector coloring.

The remainder of the paper is organized as follows. In Section 2 we present the graphs used in the proof of Theorem 1.1 and analyze their properties. In Section 3 we present our results regarding property testing. Finally in Section 4 we prove Theorem 1.1. The semidefinite programs that compute the vector chromatic number and its variants are briefly reviewed in Section A of the Appendix. Due to space limitations, some of our results appear without detailed proof.

2. The graph

In this section we construct and analyze the graphs to be used as a starting point in the proof of Theorem 1.1. Recall our three phase plan. First in Section 2.1 we consider a continuous graph G^c (with infinitely many vertices and edges) and analyze its properties. We focus on analyzing the properties that we will need later on in our proofs. Afterwards, in Section 2.2, we define the so-called dense graph G^d which is a discrete version of G^c, and show that it approximates G^c very well. The graphs used in the proof of Theorem 1.1 are obtained by taking a random sample of the dense graph G^d. In Sections 3 and 4 we show that a sufficiently large random sample of G^d will not have any large independent sets.

2.1. The continuous graph G^c

Let d be a large constant, and let $S^{d-1} = \{v \in R^d \mid \|v\| = 1\}$ be the d dimensional unit sphere. Let $G^c_k = (V^c, E^c)$ be the continuous graph in which: *(a)* The vertex set V^c consists of all the points on the unit sphere S^{d-1}. *(b)* The edge set E^c of G consists of all pairs of vertices whose respective vectors form an angle of at least $\arccos(-1/(k-1))$. As the size of V^c and E^c is infinite (and uncountable), terms such as the *number* of vertices in V^c will be replaced by the continuous analogue *measure*.

We analyze several properties of the graph G^c_k. In our analysis, we will assume that the dimension d is (at least) a very large constant (our proofs rely on such d). Additional constants that will be presented in the remainder of this section are to be viewed as independent of d.

Definition 2.1 (Sphere measure) *Let μ be the normalized $(d-1)$ dimensional natural measure on S^{d-1}, and let μ^2 be the induced measure on $S^{d-1} \times S^{d-1}$. For any two (not necessarily disjoint) subsets A and B of V^c, we define the measure of edges from A to B as*

$$E(A,B) = \mu^2\left(\{(x,y) \mid x \in A,\ y \in B,\ (x,y) \in E^c\}\right).$$

Definition 2.2 (Sphere caps) *Let $a \in [0,1]$, and $x \in S^{d-1}$. An a-cap centered at x is defined to be the set $\mathcal{C}_a = \{u \in S^{d-1} \mid \langle u, x\rangle \geq a\}$. Denote the measure of an a-cap by $\rho(a)$. A few remarks are in place. Notice that for every $x,\ x'$, an a-cap centered at x has the same measure as an a-cap centered at x' (sphere symmetry). Furthermore, notice that* large *caps have* small *corresponding values of a and vice versa.*

The value of $\rho(a) = \mu(\mathcal{C}_a)$ is approximately given by the following lemma proven for example in [8].

Lemma 2.3 *Let $\mathcal{C}_a$ be an a-cap centered at some $x \in S^{d-1}$. There exists a constant $c > 0$ (independent of a and d) s.t.*

$$\frac{c}{\sqrt{d}}\left(1-a^2\right)^{\frac{d-1}{2}} \leq \mu(\mathcal{C}_a) = \rho(a) \leq \frac{1}{2}\left(1-a^2\right)^{\frac{d-1}{2}}.$$

The main property of G^c_k of our interest is the measure of edges between any two given subsets of V^c of a specified size. We first prove that the sets in G^c_k which share the least amount of edges are caps with the same center.

Theorem 2.4 *Let $1 \geq a > 0$ and let A and B be two (not necessarily disjoint) measurable sets in V^c of measure $\rho(a)$. Let x be an arbitrary vertex of S^{d-1}. The minimum of $E(A,B)$ is obtained when $A = B = \mathcal{C}_a$ where $\mathcal{C}_a$ is an a-cap of measure $\rho(a)$ centered at x.*

The proof of Theorem 2.4 is based on symmetrization techniques similar to those presented in [8]. Due to space limitations, detailed proof is omitted. We now turn to analyze the measure of edges between caps of measure $\rho(a)$. Namely we study the value of $E(\mathcal{C}_a, \mathcal{C}_a)$.

Theorem 2.5 *Let* $\sqrt{\frac{k-2}{2(k-1)}} > a > 0$ *and* $k > 2$ *be constant. Let* $x \in S^{d-1}$*, let* $\mathcal{C}_a$ *be an* a*-cap centered at* x*. Let* $\varepsilon(a)$ *be the value of* $E(\mathcal{C}_a, \mathcal{C}_a)$*. Finally let*

$$\lambda(a) = \left(\left(1 - \frac{1}{(k-1)^2}\right) \left(1 - \frac{2(k-1)}{k-2} a^2\right) \right)^{\frac{d-1}{2}}.$$

Then $\varepsilon(a) \in \left[\left(1 - c\left(\sqrt{\frac{\log(d)}{d}}\right)\right)^{\frac{d-1}{2}} \lambda(a), \lambda(a) \right]$ *for some constant* $c > 0$*.*

Proof (sketch): Let $x \in S^{d-1}$. Let $\mathcal{C}_a$ be an a-cap centered at x. W.l.o.g. we will assume that $x = (1, 0, \ldots, 0)$. Consider a vertex $v \in \mathcal{C}_a$ on the boundary of $\mathcal{C}_a$. Let $N(v)$ be the set of vertices adjacent to v. We start by computing the measure of vertices that are neighbors of v and are in the cap $\mathcal{C}_a$, *i.e.* the measure of $N(v) \cap \mathcal{C}_a = \{u = (u_1 \ldots u_d) \in S^{d-1} \mid u_1 \geq a \text{ and } \langle v, u\rangle \leq -1/(k-1)\}$.

Claim 2.6 *Let* a, k *be as in Theorem 2.5. Let* $v = (a, \sqrt{1-a^2}, 0, \ldots, 0)$ *be a vertex on the boundary of* $\mathcal{C}_a$*. Let* $N(v)$ *be the set of neighbors of* v*. Let* $z = \sqrt{a^2 + \frac{(1/(k-1)+a^2)^2}{1-a^2}}$*. Finally let* $\delta = c\sqrt{\frac{\log(d)}{d}}$ *for a sufficiently large constant* c*. The measure of vertices in* $N(v) \cap \mathcal{C}_a$ *satisfies*

$$(1-\delta)^{\frac{d-1}{2}} \left(1 - z^2\right)^{\frac{d-1}{2}} \leq \mu(N(v) \cap \mathcal{C}_a) \leq \left(1 - z^2\right)^{\frac{d-1}{2}}$$

Claim 2.6 addresses the measure of $\mathcal{C}_a \cap N(v)$, and states that it is essentially the measure of a z-cap (where $z = \sqrt{a^2 + \frac{(1/(k-1)+a^2)^2}{1-a^2}}$). To prove Claim 2.6 we study the measure of certain restricted sets in S^{d-1}. These sets are studied in Claim 2.7. Claims 2.6 and 2.7 are depicted in Figure 1.

Claim 2.7 *Let* a, z, δ *be as in Claim 2.6. Let* $b^2 = z^2 - a^2$*. Let* $N = \{(u_1 \ldots u_d) \in S^{d-1} \mid u_1 = a,\ u_2 = b\}$*. Let* N_δ *be a* δ *neighborhood of* N *(*i.e. *all points which are of distance less that* δ *from the set* N*). The measure of the set* N_δ *is at least* $\left(1 - z^2\right)^{\frac{d-1}{2}}$*.*

To complete the proof of Theorem 2.5, let $a,\ z,\ \delta$ be as in Claim 2.6. For the upper bound, use the bound on the measure of $\mathcal{C}_a$ (Lemma 2.3) and the upper bound appearing in Claim 2.6. As for the lower bound, let $w = (w_1, w_2, \ldots, w_d) \in \mathcal{C}_a$ with first coordinate w_1 of value $a + \delta$. Consider a vertex $v = (v_1, v_2, \ldots, v_d) \in \mathcal{C}_a$

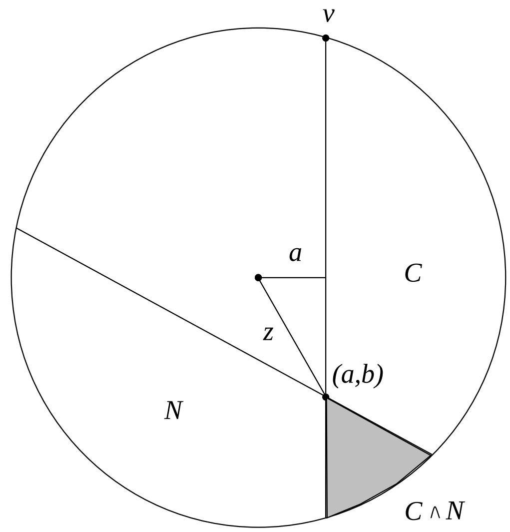

Figure 1. Projecting S^{d-1} onto the two dimensional subspace $\{(r_1, r_2, 0, \ldots, 0) \mid r_1, r_2 \in R\}$ of R^d we obtain the circle above. C is the projection of an a-cap $\mathcal{C}_a$ centered at $(1, 0, \ldots, 0)$ (where a is the distance of the cap from the origin). The vertex $v = (a, \sqrt{1-a^2}, 0, \ldots, 0) \in \mathcal{C}_a$ is on the boundary of $\mathcal{C}_a$. N is the projection of the set of points $N(v)$ on the sphere that are adjacent to v (*i.e.* form an angle of at least $\arccos(-1/(k-1))$ with v). The shaded section is $C \cap N$ (the projection of $\mathcal{C}_a \cap N(v)$). The point (a, b) is the closest point of the projection of $\mathcal{C}_a \cap N(v)$ to the origin. It is not hard to verify that $b^2 = \frac{(1/(k-1)+a^2)^2}{1-a^2}$. Finally, we denote the value of $\sqrt{a^2+b^2}$ by z. Claim 2.6 addresses the measure of $\mathcal{C}_a \cap N(v)$, and states that it is essentially the measure of a z-cap. This is done by studying the points in S^{d-1} who's projection falls *close* to (a, b). Roughly speaking, we first show that such points are in $\mathcal{C}_a \cap N(v)$, then using Claim 2.7 we show that the measure of these points is essentially the measure of a z-cap.

with first coordinate v_1 of value less than $a + \delta$. It is not hard to verify that the measure of $N(v) \cap \mathcal{C}_a$ is greater than the measure of $N(w) \cap \mathcal{C}_a$ which in turn is greater than $(1 - c\delta)^{\frac{d-1}{2}} \left(1 - z^2\right)^{\frac{d-1}{2}}$ for a sufficiently large constant c. Furthermore, for our choice of δ, the measure of vertices $v = (v_1, v_2, \ldots, v_d)$ in $\mathcal{C}_a$ with $v_1 \leq a + \delta$ is at least $\rho(a)/2$. Hence, we conclude that $E(\mathcal{C}_a, \mathcal{C}_a)$ is at least $\frac{\rho(a)}{2}(1-c\delta)^{\frac{d-1}{2}} \left(1 - z^2\right)^{\frac{d-1}{2}}$. Simplifying the above expression, we conclude our assertion. □

Theorem 2.5 addresses the case in which k is constant and the caps considered are both of measure $\rho(a)$ for a constant value of a. For the proof of Theorem 1.1 we also need to addresses nonconstant values of a and k which depend on d.

Theorem 2.8 *Let $a = \left(\frac{\log(d)}{d}\right)^{\frac{1}{4}}$. Let k satisfy $\frac{1}{k-1} = a^2$. Let $x \in S^{d-1}$, let $\mathcal{C}_a$ be an a-cap centered at x. Let $\varepsilon(a)$ be the value of $E(\mathcal{C}_a, \mathcal{C}_a)$. The value of $\varepsilon(a)$ is in the range*

$$\left[\frac{1}{poly(d)}\left(1 - \frac{2(k-1)}{k-2}a^2\right)^{\frac{d-1}{2}}, \left(1 - \frac{2(k-1)}{k-2}a^2\right)^{\frac{d-1}{2}}\right].$$

The outline of the proof of Theorem 2.8 is similar to that of Theorem 2.5. Combining Theorems 2.4, 2.5, 2.8 we obtain:

Corollary 2.9 *Let a, k, $\varepsilon(a)$ be defined as in Theorem 2.5 or 2.8. The value of $E(A, B)$ for any two subsets A and B both of measure $\rho(a)$ is at least $\varepsilon(a)$.*

2.2. The dense graph G^d

We now define the dense graph $G_k^d = (V, E)$ that corresponds to G_k^c. It is shown in [8] that S^{d-1} can be partitioned into $n = 2^{\theta(d^2)}$ cells each of equal size and of diameter at most 2^{-d}. Let $\mathcal{P} = \{C_1 \dots C_n\}$ denote the cells obtained in the above partition. The graph G_k^d will be of size n, in which each vertex $v_i \in V$ corresponds to a cell $C_i \in \mathcal{P}$. The edge set of G_k^d consists of the edges (u, v) iff there is a positive measure of edges in G_k^c between their corresponding cells C_u, C_v. We show that G_k^d and G_k^c share the same expansion properties.

Definition 2.10 *Let A and B be (not necessarily disjoint) subsets of G_k^d. For each vertex $v \in A$ let $d_v(B)$ be the number of neighbors v has in B. Let $E(A, B) = \sum_{v \in A} d_v(B)$.*

The proof of the following claims is omitted.

Lemma 2.11 *Let A and B be subsets of G_k^d, and let A_c and B_c be the corresponding subsets of G_k^c. (a) The size of A (B) is ρn iff A_c (B_c) has measure ρ. (b) $E(A, B) \geq E(A_c, B_c)n^2$.*

Theorem 2.12 *Let $a, k,$ $\varepsilon(a)$ be as defined in Theorem 2.5 or 2.8. For any two subsets A and B in G_k^d both of size $\rho(a)n$, the value of $E(A, B)$ is at least $\varepsilon(a)n^2$.*

Lemma 2.13 *The graph G_k^d is vector $k\left(1 + \frac{ck^2}{2^d}\right)$-colorable for some constant $c > 0$.*

By definition, the continuous graph G_k^c is vector k-colorable. In Lemma 2.13 we show that the finite approximation G_k^d to G_k^c is *almost* vector k-colorable. In general, this does not suffice for the proof of Theorem 1.1 as we are interested in graphs which are vector k-colorable (rather than "almost vector k-colorable"). This can be fixed by starting with a continuous graph with vector coloring number slightly less than k (e.g. $k/\left(1 + \frac{ck^2}{2^d}\right)$). In order to simplify our presentation, we neglect this point and consider the graph G_k^d to be exactly vector k-colorable. This is possible due to the fact that the properties of G_k^c are *continuous* in k. Namely, choosing d large enough, it can be seen that the multiplicative error of $(1 + \frac{ck^2}{2^d})$ in the value of k does not effect the analysis appearing throughout our work.

Lemma 2.14 *Let $\rho(\frac{1}{k-1})$ be the measure of a $\frac{1}{(k-1)}$-cap. Every vertex v in the graph G_k^d has degree $d_v \in [\frac{1}{poly(d)}\rho(\frac{1}{k-1})n, poly(d)\rho(\frac{1}{k-1})n]$.*

3. Constructing the sparse graph, and testing of $\alpha(G)$

Let G be a graph of size n which does not have an independent set of size ρn (*i.e.* $\alpha(G) < \rho n$). Let R be a random subgraph of G of size s (*i.e.* R is the subgraph induced by a random subset of vertices in G of size s). In this section we study the minimal value of s for which $\alpha(R) < \rho s$ with high probability.

In general, if our only assumption on G is that $\alpha(G) < \rho n$, we cannot hope to set s to be smaller than n. Hence, we strengthen our assumption on G, to graphs G which not only satisfy $\alpha(G) < \rho n$ but are also *far* from having an independent set of size ρn (we defer defining the exact notion of "far" until later in this discussion). That is, given a graph G which is *far* from having an independent set of size ρn, we ask for the minimal value of s for which (with high probability) a random subgraph of size s does not have an independent set of size ρs. This question (and many other closely related ones) have been studied in [9] under the title of *property testing*.

In [9], a graph G of size n is said to be ε-far from having an independent set of size ρn if any set of size ρn in G has at least εn^2 induced edges. It was shown in [9] that if G is ε-far from having an independent set of size ρn then with high probability a random subgraph of size $s = \frac{c \log(1/\varepsilon)\rho}{\varepsilon^4}$, for a sufficiently large constant c, does not have an independent set of size ρs.

The results of [9] do not suffice for the proof (as we present it) of Theorem 1.1. We thus turn to strengthening their results. To do so, we introduce a stronger notion of being "ε-far". Roughly speaking, we prove that under our new notion of distance, choosing s to be of size $\frac{\rho}{\varepsilon}$ suffices. Furthermore, by showing a connection between our notion of distance and that presented in [9], we improve the result of [9] stated above and obtain a sample size proportional to $1/\varepsilon^3$. The proof of the following theorem will appear at the end of this section.

Theorem 3.1 *Let G be a graph of size n. Let $\rho < 1$. Assume that every subset of size ρn in G induces a subgraph with at least εn^2 edges. Let c be a large constant. Let $s = \frac{c\rho^4 \log(\rho^4/\varepsilon^3)}{\varepsilon^3}$. If $s \leq \frac{\varepsilon}{3\rho} n$ then with probability at least $3/4$ a random set R of size s of G does not have an independent set of size ρs.*

The choice of the new notion of ε-far that we introduce is governed by two concerns. One is that we need a notion that will allow us to reduce the sample size s considerably compared to the known bounds (such as those of Theorem 3.1). The other is that we need a notion that is satisfied by the dense graph of Section 2.2. One aspect of our new notion is that rather than considering the number of edges in subsets of size ρn, we consider somewhat smaller sets, whose size is parameterized by an additional parameter $\delta \in (0, 1]$. Hence we would like sets B of small size $\delta\rho n$ to have many edges. The smaller we can make δ, the better our eventual bounds on the sample size. Another aspect of our new notion is that we shall consider not only edges within B, but also edges going out of B. We will want that the number of edges going out of B to any large enough set A grows at least linearly with the size of A. As a special case we can take $A = B$, meaning that there are many edges in the subgraph induced by B. In fact, we shall be interested in the case that $B \subseteq A$. In this case, we may think of B as the $\delta\rho n$ vertices of lowest degree in the subgraph induced by A, and we want the degree of at least one of these vertices to be high, and to grow linearly with the size of A.

We use $\mathcal{D}(\varepsilon, \delta, \rho)$ to denote the set of graphs of size n which are ε-far from having an independent set of size ρn, according to our new notion. The exact definition follows below.

Definition 3.2 *Let G be a graph of size n and let $\rho < 1$. Let $\delta \in (0, 1]$. If for any $\alpha \in [1, \frac{1}{\delta\rho}]$ and any subset A of G of size $\alpha\delta\rho n$, A induces a subgraph with at most $\delta\rho n$ vertices of degree less than $\alpha\frac{\varepsilon}{\delta\rho} n$ then $G \in \mathcal{D}(\varepsilon, \delta, \rho)$.*

One may rephrase the definition if $\mathcal{D}(\varepsilon, \delta, \rho)$ as follows. Let A be any subset in G of size $\alpha\delta\rho n$. Let L (for *low degree*) be the $\delta\rho n + 1$ vertices of minimal degree in the subgraph induced by A. $G \in \mathcal{D}(\varepsilon, \delta, \rho)$ iff for every such A, the corresponding subset L has a vertex of degree at least $\alpha\frac{\varepsilon}{\delta\rho} n$ (in the subgraph induced by A).

Consider a graph G which is ε-far from having an independent set of size ρn according to the original notion of [9]. We show that $G \in D(\varepsilon\rho, 1, \rho)$. Let $\alpha \in [1, 1/\rho]$ and let A be any subset of vertices in G of size $\alpha\rho n$. Let L be the ρn vertices in A with minimal degree. As L induces at least εn^2 edges there must be a vertex in L of degree at least $\frac{\varepsilon}{\rho} n$. Hence, as α is at most $1/\rho$, we conclude that $G \in D(\varepsilon\rho, 1, \rho)$. We now turn to prove the main theorem of this section.

Theorem 3.3 *Let G be a graph of size n. Let $\rho < 1$. Let $\delta \in (0, 1]$, let $\varepsilon > 0$, and let c be a large constant. Let $s = \frac{c\rho}{\varepsilon(1-\delta)^2} \log \frac{1}{\varepsilon(1-\delta)^2} \log \frac{1}{\delta\rho}$ If $s \leq \frac{\delta(1-\delta)}{1+\delta}\rho n$, and $G \in \mathcal{D}(\varepsilon, \delta, \rho)$ then with probability at least $3/4$ a random set R of size s of G does not have an independent set of size ρs.*

Two remarks are in place before we present the proof of Theorem 3.3. First we would like to address the restriction $s \leq \frac{\delta(1-\delta)}{1+\delta}\rho n$. In our applications (and also in standard ones) the values of $1 - \delta$, ρ and ε are assumed to be small but still large enough to satisfy $s = \frac{c\rho}{\varepsilon(1-\delta)^2} \log \frac{1}{\varepsilon(1-\delta)^2} \log \frac{1}{\delta\rho} \leq \frac{\delta(1-\delta)}{1+\delta}\rho n$. Hence the reader may ignore this restriction. Secondly, the result of Theorem 3.3 can be improved to $s = \frac{c\rho}{\varepsilon(1-\delta)^2} \log \frac{1}{\varepsilon(1-\delta)^2}$ by considering a slightly different definition of $\mathcal{D}(\varepsilon, \delta, \rho)$ (details are omitted).

Proof : Our proof uses ideas appearing in [2] (in which testing the chromatic number of G is considered). We would like to prove that a sufficiently large subset R of size s of G does not have an independent set of size ρs. Let $\{r_1 \ldots r_s\}$ be the vertices of R. Consider choosing the vertices of R one by one, such that at each step the random subset chosen so far is $R_i = \{r_1 \ldots r_i\}$. Consider an independent set $I \subseteq R_i$ of size less than ρs. We would like to show that (with high probability) I cannot be extended to an independent set of size ρs by choosing additional vertices from R. For each such I, the vertices in $V \setminus R_i$ that cannot be added to I are exactly those adjacent to some vertex in I. Let $F(I)$ (for *free*) be the set of vertices in $V \setminus R_i$ which are not adjacent to any vertices in I, and let $N(I)$ be the set of vertices that are adjacent to a vertex in I. Consider the next random vertex $r_{i+1} \in R$. If r_{i+1} is chosen from $N(I)$ then it cannot be added to I, and we view this round as a success regarding the set I. Otherwise, r_{i+1} happens to be in $F(I)$ and can be added to I. But if r_{i+1} also happens to have many neighbors in $F(I)$ then adding it to I will substantially reduce the size of $F(I \cup \{r_{i+1}\})$ which works in our favor. This later case is also viewed as a successful round regarding I.

Motivated by the discussion above, for each subset $I \subseteq R_i$ we define the following set $RES(I)$ of vertices that *restrict* upon I. A vertex v in $V \setminus R_i$ is in $RES(I)$ iff one of the following occur

1. $v \in N(I)$
2. $|F(I)| \geq \delta\rho n$, **and** $v \in F(I)$, **and** v has at least $\frac{\varepsilon|F(I)|}{(\delta\rho)^2}$ neighbors in $F(I)$.

Lemma 3.4 *Let R_i be of size at most $s = \frac{\delta(1-\delta)}{1+\delta}\rho n$. For any subset $I \subseteq R_i$, the subset $RES(I)$ is of size at least $n(1 - \frac{2\delta}{1+\delta}\rho)$.*

Proof : Let I be any subset of R_i. The size of $RES(I)$ is at least $|N(I)|$. Hence if $N(I)$ is of size greater than $n -$

$s - \delta\rho n = n(1 - \frac{2\delta}{1+\delta}\rho)$ then so is $RES(I)$. Otherwise, the size of $F(I)$ is at least $\delta\rho n$ and by the definition of $\mathcal{D}(\varepsilon, \delta, \rho)$ we have that all but $\delta\rho n$ vertices in $F(I)$ have at least $\alpha\frac{\varepsilon}{\delta\rho}n$ neighbors in $F(I)$ (where α satisfies $|F(I)| = \alpha\delta\rho n$). We conclude that $|RES(I)| \geq n - s - \delta\rho n = n(1 - \frac{2\delta}{1+\delta}\rho)$. □

Corollary 3.5 *For any subset R_i of size less than $s = \frac{\delta(1-\delta)}{1+\delta}\rho n$ and any subset $I \subseteq R_i$, a random vertex in $V \setminus R_i$ is in the subset $RES(I)$ with probability at least $1 - \frac{2\delta}{1+\delta}\rho$.*

Consider the following tree T defined by choosing the vertices r_i of R one at a time. Each vertex of T is indexed by a pair (X, Y) where both X and Y are subsets of R. If for a vertex (X, Y) it is the case that the set X is an independent set we call (X, Y) an open vertex. Otherwise, (X, Y) is referred to as closed. The root of T is the vertex (ϕ, ϕ) (which is open).

Assume that the vertices $R_{i-1} = \{r_1 \ldots r_{i-1}\}$ have been chosen, and that T is the tree obtained so far. Let (X, Y) be an open leaf of T. Let r_i be the next random vertex in G. If $r_i \notin RES(X)$ we do not do anything with respect to the vertex (X, Y). Otherwise, when $r_i \in RES(X)$ we consider two cases. If $|Y| < s - \rho s$ then we add two sons to (X, Y), one son labeled $(X \cup \{r_i\}, Y)$ and another labeled $(X, Y \cup \{r_i\})$. In the second case in which $|Y| \geq s - \rho s$ we add only one son to T labeled $(X \cup \{r_i\}, Y)$. If the leaf (X, Y) is closed, we do nothing with respect to r_i.

In this process it is not hard to verify that after choosing the vertices $r_1 \ldots r_i$ each node (X, Y) of T satisfies: *(a)* X and Y are disjoint subsets of $\{r_1 \ldots r_i\}$, and *(b)* $|Y| \leq s - \rho s$.

Lemma 3.6 *For each vertex (X, Y) in T, the size of X is bounded by $t = \frac{(\delta\rho)^2 \lceil \log(1/\delta\rho) \rceil}{\varepsilon} + 2$.*

Proof : Let $P = \langle v_0, v_1, \ldots, v_l \rangle$ be a path in T with $v_0 = (\phi, \phi)$ and $v_l = (X, Y)$. For $i = 0 \ldots l$ let $(X_i, Y_i) = v_i$ be the subsets corresponding to v_i and for $i = 1 \ldots l$ let $\hat{r}_i = (X_i \cup Y_i) \setminus (X_{i-1} \cup Y_{i-1})$. Assume that v_l is an open vertex of T. Let $e_i = (v_{i-1}, v_i)$ be the edges of P. An edge e_i is said to be an X-edge if it is obtained by adding the vertex $\hat{r}_i$ to X_i (*i.e.* $X_i = X_{i-1} \cup \{\hat{r}_i\}$), otherwise e_i is said to be a Y-edge. The size of X is thus bounded by the number of X-edges in P.

Consider an X-edge $e_i = (v_{i-1}, v_i)$ where $v_{i-1} = (X_{i-1}, Y_{i-1})$ and $v_i = (X_i, Y_i)$. Both vertices v_{i-1} and v_i are open, implying that $\hat{r}_i$ is in the subset $RES(X_{i-1}) \setminus N(X_{i-1})$. This implies that: *(1)* $|F(X_{i-1})| \geq \delta\rho n$, and *(2)* $|F(X_{i-1})| - |F(X_i)|$ is at least $\alpha\frac{\varepsilon}{\delta\rho}n$ where $|F(X_{i-1})| = \alpha\delta\rho n$.

We would like to bound the number of X-edges in P. The main idea used in our proof is that (as stated above) each X-edge reduces the size of the corresponding set F significantly. Namely, for each x between 0 and $\lceil \log(1/(\delta\rho)) \rceil - 1$, we bound the number of X-edges $e_i = (v_{i-1}, v_i)$ for which $|F(X_{i-1})| \in [\frac{n}{2^x}, \frac{n}{2^{x+1}})$. Each such X-edge $e_i = (v_{i-1}, v_i)$ satisfies $|F(X_{i-1})| - |F(X_i)| \geq \frac{\varepsilon}{(\delta\rho)^2 2^{x+1}}n$. Furthermore, any two X-edges $e_i = (v_{i-1}, v_i)$ and $e_j = (v_{j-1}, v_j)$ (here $i < j$) for which $|F(X_{i-1})|,\ |F(X_{j-1})| \in [\frac{n}{2^x}, \frac{n}{2^{x+1}})$ satisfy (*a*) $F(X_{j-1}) \subseteq F(X_i)$ and (*b*) $|F(X_{i-1})| - |F(X_{j-1})| < \frac{n}{2^{x+1}}$. We conclude that the number of X-edges $e_i = (v_{i-1}, v_i)$ for which $|F(X_{i-1})| \in [\frac{n}{2^x}, \frac{n}{2^{x+1}})$ is bounded by $\frac{(\delta\rho)^2}{\varepsilon}$. Each X-edge $e_i = (v_{i-1}, v_i)$ satisfies $|F(X_{i-1})| \geq \delta\rho n$, thus all in all, the number of X-edges in P is bounded by $\frac{(\delta\rho)^2 \lceil \log(1/\delta\rho) \rceil}{\varepsilon} + 1$.

We conclude that if (X, Y) is open then the size of X is at most $\frac{(\delta\rho)^2 \lceil \log(1/\delta\rho) \rceil}{\varepsilon} + 1$. It is left to consider the case in which $v_l = (X, Y)$ is a closed leaf. The asserted bound follows from the fact that any path in T has at most one closed vertex (which must appear as the last vertex in the path). □

Corollary 3.7 *Let t be as in Lemma 3.6. The depth of T is bounded by $s - \rho s + t$.*

Proof : Follows from the fact that in each vertex (X, Y) of T the size of X is bounded by t, and the size of Y is bounded by $s - \rho s$. Notice than if a path $P = \langle v_0 = (\phi, \phi), v_1, \ldots, v_l = (X, Y) \rangle$ is of length l then $|X| + |Y| = l$. □

Lemma 3.8 *Let T be the tree obtained after choosing $R = \{r_1 \ldots r_s\}$, let t be as in Lemma 3.6. Then T has at most $\sum_{i=1}^{t} \binom{s}{i}$ leaves.*

Proof : Let (X_1, Y_1) and (X_2, Y_2) be two leaves of T. Let (X, Y) be their least common ancestor. (X, Y) has two sons $(X \cup \{r_i\}, Y)$ and $(X, Y \cup \{r_i\})$. We conclude that r_i is in X_1 and not in X_2 (or vice versa) implying that $X_1 \neq X_2$. Furthermore, X_1 and X_2 are both subsets of R of size at most t (Lemma 3.6). Hence, the number of leaves in T is bounded by the number of different subsets of R of size at most t. □

Lemma 3.9 *If after choosing s random vertices $R = \{r_1 \ldots r_s\}$ all leaves of T are closed then R does not have any independent sets of size ρs.*

Proof : Let A be a subset of size ρs in R. To show that A cannot be an independent set in R, we traverse T starting with the root (ϕ, ϕ).

In this process let us only visit vertices (X, Y) for which $X \subseteq A$ and $A \cap Y = \phi$. Suppose we are at a vertex (X, Y) of T. If (X, Y) is a closed vertex then X is not an independent set implying that A isn't either. Otherwise, if $|Y| < s - \rho s$, there is a vertex r_i of R such that $(X \cup \{r_i\}, Y)$

and $(X, Y \cup \{r_i\})$ are sons of (X, Y). In this case we continue to traverse T by moving to $(X \cup \{r_i\}, Y)$ if $r_i \in A$ and to $(X, Y \cup \{r_i\})$ if $r_i \notin A$. Finally, if $|Y| = s - \rho s$ the vertex (X, Y) only has one son $(X \cup \{r_i\}, Y)$. As $|A| = \rho s$ and $A \cap Y = \phi$ we conclude that r_i must be in A and we continue to traverse T by moving to $(X \cup \{r_i\}, Y)$.

By our assumption, after choosing s random variables all leaves of T are closed. This implies that our traversal procedure eventually ends at a closed leaf, therefore A is not an independent set. $\square$

Lemma 3.10 *Let c be a sufficiently large constant. Let t be as in Lemma 3.6 Let $s = c\frac{t}{\rho(1-\delta)^2} \log \frac{t}{\rho(1-\delta)^2}$. If $s \leq \frac{\delta(1-\delta)}{1+\delta}\rho n$, then after choosing s random vertices $R = \{r_1 \dots r_s\}$, all leaves of T are closed with probability at least $3/4$.*

Proof (sketch): By Lemma 3.8 the number of leaves in T is bounded by $\sum_{i=1}^{t} \binom{s}{i} \leq t\binom{s}{t}$. Consider a (potential) path $P = \langle v_0, v_1, \dots v_l \rangle$ in T form the root $v_0 = (\phi, \phi)$ to a leaf $v_l = (X, Y)$. Recall the each vertex $v_i = (X_i, Y_i)$ along the path P defines a set of vertices $RES(X_i)$. The probability that v_l is open after choosing s random vertices is at most the probability that the total number of steps in which the vertex r_i chosen hits the set $RES(X_i)$ corresponding to that specific step is less than $s - \rho s + t$ (Corollary 3.7).

As $|R| = s \leq \frac{\delta(1-\delta)}{1+\delta}\rho n$, at each stage the probability of choosing a vertex in the set $RES(X_i)$ is at least $p = 1 - \frac{2\delta}{1+\delta}\rho$ (Corollary 3.5). Therefore, the probability that v_l is an open leaf is at most the probability that the Binomial random variable $B(s, p)$ is less than $s - \rho s + t$. This is equal to the probability that the Binomial random variable $B(s, 1-p)$ is more than $\rho s - t$. Setting s to be $c\frac{t}{\rho(1-\delta)^2} \log \frac{t}{\rho(1-\delta)^2}$ for a sufficiently large constant c, it can be seen that the above probability is less than $\frac{3}{4}\left(t\binom{s}{t}\right)^{-1}$ (standard Chernoff bound). Hence, (union bound) the probability that some leaf $v \in T$ is open after choosing s random vertices is at most $3/4$. $\square$

The proof of Theorem 3.3 follows from Lemmas 3.9 and 3.10. $\square$

Next we state the main Corollary that we use in the proof of Theorem 1.1. We start by recalling Definition 2.10 from Section 2.

Definition 2.10 *Let A and B be (not necessarily disjoint) subsets of G. For each vertex $v \in A$ let $d_v(B)$ be the number of neighbors v has in B. Let $E(A, B) = \sum_{v \in A} d_v(B)$.*

Corollary 3.11 *Let G be a graph of size n, and $\rho < 1$. Let $\delta \in (0, 1]$, let $\varepsilon > 0$, and let c be a large constant. Let $s = \frac{c\rho}{\varepsilon(1-\delta)^2} \log \frac{1}{\varepsilon(1-\delta)^2} \log \frac{1}{\delta\rho}$. Assume $s \leq \frac{\delta(1-\delta)}{1+\delta}\rho n$. If for any two subsets A and B both of size $\delta\rho n$ it is the case that $E(A, B) \geq \varepsilon n^2$ then with probability at least $3/4$ a random set R of size s of G does not have an independent set of size ρs.*

Proof : Let $\alpha \in [1, 1/(\delta\rho)]$. Let X be a subset of G of size $\alpha\delta\rho n$. Let $X = \{X_1, X_2, \dots, X_\ell\}$ be a partition of X into ℓ sets in which the size of X_i for all $i \neq \ell$ is $\delta\rho n$. Notice that $\ell = \lfloor \alpha \rfloor + 1$. For each $v \in X$ let $d_v(X)$ be the degree of v in the subgraph induced by X. Let L be the $\delta\rho n$ vertices in X with minimal degree. By our assumption, the value of $E(L, X_i)$ for each i (except $i = \ell$) is at least εn^2. Hence, $\sum_{v \in L} d_v(X)$ is at least $\lfloor \alpha \rfloor \varepsilon n^2 \geq \frac{\alpha}{2}\varepsilon n^2$. This implies that L must include a vertex v with degree $d_v(X) \geq \alpha\frac{\varepsilon}{2\delta\rho}n$. We conclude that $G \in \mathcal{D}(\frac{\varepsilon}{2}, \delta, \rho)$ (as defined in Definition 3.2). Now Theorem 3.3 implies our assertion. $\square$

We now present the proof of Theorem 3.1 stated in the beginning of this section.

Proof (Theorem 3.1): Let $\delta = 1 - \frac{\varepsilon}{2\rho^2}$. We start by showing that $G \in \mathcal{D}(\delta\rho\varepsilon, \delta, \rho)$. Let A be some subset of G of size $l = \alpha\delta\rho n$ (for $\alpha \in [1, 1/\delta\rho]$). For each vertex v in A let $d_v(A)$ be the number of neighbors v has in A. Let $\{v_1 \dots v_l\}$ be the vertices of A ordered by their degree $d_v(A)$. That is $d_{v_i}(A) \leq d_{v_j}(A)$ if $i < j$. Let L be the subset of $\delta\rho n$ *low-degree* vertices in A, that is $L = \{v_1 \dots v'_l\}$ for $l' = \delta\rho n$. Also let $d = d_{v_{l'}}(A)$. Notice, that the degree (in A) of every vertex in L is bounded by d.

Consider the subgraph induced by L. We will show that L has many induced edges. As the number of edges in L is bounded by $|L|d$ this will imply a lower bound on d. Details follow.

Let L^c be any set in $V \setminus L$ of size $(1-\delta)\rho n$. It is known that the number of edges induced by the set $L \cup L^c$ is at least εn^2 (notice that $|L \cup L^c| = \rho n$). The number of edges (in $L \cup L^c$) adjacent to vertices in L^c is bounded by $(1-\delta)\rho^2 n^2 = \frac{\varepsilon n^2}{2}$ (recall that we set $\delta = 1 - \frac{\varepsilon}{2\rho^2}$). Hence, the number of edges induced by vertices in L is at least $\frac{\varepsilon n^2}{2}$.

For every vertex $v \in L$ the degree of v in the subgraph induced by L is at most the degree of v in A (that is $d_v(A)$) which is at most d. Hence, $d|L|/2$ is an upper bound on the number of edges in the subgraph induced by L. We conclude that $d \geq \frac{\varepsilon}{\delta\rho}n \geq \alpha\varepsilon n$ (we use the fact that $\alpha \leq 1/(\delta\rho)$). This implies that $G \in \mathcal{D}(\delta\rho\varepsilon, \delta, \rho)$ for $\delta = 1 - \frac{\varepsilon}{2\rho^2}$. Now using Theorem 1.1 and noticing that in this case the value of t in Lemma 3.6 can be set to $\frac{O(\delta\rho)}{\varepsilon}$ we conclude our assertion. $\square$

4. Proof of Theorem 1.1

In the following section we address the proof of Theorem 1.1 presented in the Introduction. The outline of our proof for Theorem 1.1 is as follows. Recall that we are looking for a graph G which on one hand has small vector chromatic number but on the other $\alpha(G)$ is small. We start with the graph G_k^d presented in Section 2. It is not hard to verify that G_k^d satisfies the first property but is far from satisfying the latter. We thus consider a random subgraph R of G_k^d. Let s be as in Theorem 3.3. Choosing R to be of size s and using Theorem 3.3 of Section 3, we will obtain the graphs needed to prove our theorem.

Next we prove the first assertion of Theorem 1.1. Proof of the remaining assertions in Theorem 1.1 are omitted. The three assertions are all proven similarly, the main difference between their proof is the choice of parameters used. To avoid confusion, we restate the first assertion of Theorem 1.1 using a slightly different notation then that appearing in the original presentation.

Theorem 1.1 (a) *For every constant $\varepsilon > 0$ and constant $k > 2$, there are infinitely many graphs R that are vector k-colorable and satisfy $\alpha(R) \leq \frac{s}{\Delta_R^{1-\frac{2}{k}-\varepsilon}}$, where s is the number of vertices in R, Δ_R is the maximal degree in R, and $\Delta_R > s^\delta$ for some constant $\delta > 0$.*

Proof : Let $k > 2$ be constant. Let $a > 0$ be an arbitrarily small constant. Let $G = G_k^d = (V, E)$ be the graph from Section 2. Let n be the size of the vertex set V of G, and let Δ be the maximal degree of G (recall, by Lemma 2.14, that all vertices in G are of degree approximately Δ). Recall that $n = 2^{\theta(d^2)}$ where d is the dimension in which the corresponding graph G_k^c was defined. We will assume the dimension d is a very large constant determined after fixing a. We would like to apply the Theorems of Section 3 on the graph G_k^d. Fix δ to be of size $1/2$. Let $\rho = \rho(a)/\delta$ (*i.e.* $\delta\rho = \rho(a)$), where the function $\rho(a)$ is as defined in Section 2. We start by proving the following Corollary of Theorem 2.12.

Corollary 4.1 *Any two subsets A and B of vertices in G of size $\delta\rho n$ satisfy $E(A, B) \geq \Delta\rho^{\frac{2(k-1)}{k-2}+a} n$.*

Proof : We will prove the assertion for $k = 3$. The case in which k is an arbitrary constant > 2 is analogous. Recall that $\delta\rho = \rho(a)$ for a small constant $a > 0$.

Let A and B be any subsets of vertices in G of size $\delta\rho n = \rho(a)n$. By Theorem 2.12, we have that $E(A, B)$ is at least $\left[\left(1 - c\left(\sqrt{\frac{\log(d)}{d}}\right)\right)\left(\frac{3}{4}(1-4a^2)\right)\right]^{\frac{d-1}{2}} n^2$ for some constant c. It is not hard to verify that $\left(1 - c\left(\sqrt{\frac{\log(d)}{d}}\right)\right)^{\frac{d-1}{2}} > \rho^{a/4}$. Furthermore, by Lemma 2.14 we have that $\left(\frac{3}{4}\right)^{\frac{d-1}{2}} > \frac{\Delta}{n}\rho^{a/4}$. Hence, $E(A, B) \geq \rho^{a/2}\left(1-4a^2\right)^{\frac{d-1}{2}} \Delta n$. It is not hard to verify that $\left(1-4a^2\right)^{\frac{d-1}{2}} \geq \left(1-a^2\right)^{\frac{(4+a/3)(d-1)}{2}} \geq \rho^{4+a/2}$. We conclude that any subsets A and B both of size $\delta\rho n$ satisfy $E(A, B) \geq \Delta\rho^{4+a} n$. ☐

Applying Corollary 3.11 on G with $\varepsilon = \frac{\Delta}{n}\rho^{\frac{2(k-1)}{k-2}+a}$ and $\delta = 1/2$ we obtain

Corollary 4.2 *Let c be a large constant. With probability $\geq 3/4$ a random set R of G of size $s = \frac{c\log(1/\rho)\log(n/(\Delta\rho))}{\rho^{\frac{k}{k-2}+a}}\frac{n}{\Delta}$ does not have an independent set of size ρs.*

Let R be a random subset of G of size s as above. The vector chromatic number of G is k (recall the discussion following Lemma 2.13). This implies that any subgraph H of G (including that induced by R) is vector k-colorable. Consider the maximal degree Δ_R of R, recall that all vertices of G have degree approximately Δ.

Claim 4.3 *With probability greater than $3/4$, R will have maximal degree $\Delta_R \in [\frac{1}{poly(d)}\Delta\frac{s}{n}, 2\Delta\frac{s}{n}]$.*

Proof : Consider any vertex v in G. Denote the size of R by s. Let $d_v(R)$ be the number of neighbors v has in the subset R. For each vertex v, the expected value of $d_v(R)$ is in the range $[\frac{1}{poly(d)}\Delta\frac{s}{n}, \Delta\frac{s}{n}]$ (Lemma 2.14). Thus (using standard bounds) the probability that $d_v(R)$ deviates from its expectation by more than a constant fraction of its expectation is at most $2^{-\Omega(\Delta\frac{s}{poly(d)n})}$. The probability that some vertex in G has degree (in R) $\notin [\frac{1}{poly(d)}\Delta\frac{s}{n}, 2\Delta\frac{s}{n}]$ is thus at most $2^{\log(n)-\Omega(\Delta\frac{s}{poly(d)n})} \leq 3/4$ for our choice of s. ☐

We conclude that with probability at least $1/2$ the subgraph induced by R has maximal degree $\Delta_R \in [\frac{1}{poly(d)}\Delta\frac{s}{n}, 2\Delta\frac{s}{n}]$ (Claim 4.3), and does not have an independent set of size ρs (Corollary 4.2). As $\Delta_R \leq O(\log(1/\rho)\log(n/(\Delta\rho))(1/\rho)^{\frac{k}{k-2}+a})$, and $\Delta_R \geq \frac{1}{poly(d)}\Delta\frac{s}{n} \geq s^\delta$ for some constant $\delta > 0$ we conclude our assertion. ☐

References

[1] N. Alon and N. Kahale. Approximating the independence number via the θ function. *Mathematical Programming*, 80(2):253–264, 1998.

[2] N. Alon and M. Krivelevich. Testing k-colorability. *SIAM Journal on Discrete Mathematics*, 15(2):211–227, 2002.

[3] A. Blum and D. Karger. An $o(n^{3/14})$-coloring for 3-colorable graphs. *Information Processing Letters*, 61(6):49–53, 1997.

[4] M. Charikar. On semidefinite programming relaxations for graph coloring and vertex cover. *Proceedings of SODA*, pages 616–620, 2002.

[5] U. Feige. Randomized graph products, chromatic numbers, and the Lovász theta function. *Combinatorica*, 17(1):79–90, 1997.

[6] U. Feige. Approximating maximum clique by removing subgraphs. *Manuscript*, 2002.

[7] U. Feige and J. Kilian. Zero knowledge and the chromatic number. *Journal of Computer and System Sciences*, 57:187–199, 1998.

[8] U. Feige and G. Schechtman. On the optimality of the random hyperplane rounding technique for MAX CUT. *Random Structures and Algorithms*, 20(3):403–440, 2002.

[9] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of ACM*, 45(4):653–750, 1998.

[10] M. Halldorsson. A still better performance guarantee for approximate graph coloring. *Information Processing Letters*, 45:19–23, 1993.

[11] E. Halperin, R. Nathaniel, and U. Zwick. Coloring k-colorable graphs using smaller palettes. *Proceedings of SODA*, pages 319–326, 2001.

[12] J. Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *In Proceedings of 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 627–636, 1996.

[13] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. *Journal of ACM*, 45(2):246–265, 1998.

[14] S. Khot. Improved inapproximability results for max clique, chromatic number and approximate graph coloring. *In Proceedings of 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[15] L. Lovász. On the Shannon capacity of a graph. *IEEE Transactions on Information Theory*, 25:2–13, 1979.

[16] M. Szegedy. A note on the θ number of Lovász and the generalized Delsarte bound. *In proceedings of 35th Annual Symposium on Foundations of Computer Science*, pages 36–39, 1994.

Appendix

A. Semidefinite relaxations for the vector chromatic number and the Lovász θ function

There are many equivalent ways to define the Lovász θ function and its variants. We will follow the definitions suggested in [13, 4]. Let $G = (V, E)$ be a graph of size n. For convenience we will assume that $V = [1 \dots n]$. The semidefinite relaxations below assign unit vectors to every vertex $i \in V$. These unit vectors are to satisfy certain constraints which will in turn determine the value of the relaxations.

$COL_1(G)$ Minimize k
subject to:

$$\langle v_i, v_j \rangle \le -\tfrac{1}{k-1} \quad \forall (i,j) \in E$$
$$\langle v_i, v_i \rangle = 1 \quad \forall i \in V$$

$COL_2(G)$ Minimize k
subject to:

$$\langle v_i, v_j \rangle = -\tfrac{1}{k-1} \quad \forall (i,j) \in E$$
$$\langle v_i, v_i \rangle = 1 \quad \forall i \in V$$

$COL_3(G)$ Minimize k
subject to:

$$\langle v_i, v_j \rangle = -\tfrac{1}{k-1} \quad \forall (i,j) \in E$$
$$\langle v_i, v_j \rangle \ge -\tfrac{1}{k-1} \quad \forall i,j \in V$$
$$\langle v_i, v_i \rangle = 1 \quad \forall i \in V$$

The function $COL_1(G)$ is the vector chromatic number of G as defined in [13]. The function $COL_2(G)$ is the strict vector chromatic number of G and is equal to the Lovász θ function on $\bar{G}$, where $\bar{G}$ is the *complement* graph of G. Finally, the function $COL_3(G)$ will be referred to as the *strong* vector chromatic number as defined in [16, 4]. Let $\omega(G)$ denote the size of the maximal clique in G, following we will show that

$$\omega(G) \le COL_1(G) \le COL_2(G) \le COL_3(G) \le \chi(G)$$

It is not hard to verify that $COL_1(G) \le COL_2(G) \le COL_3(G)$. To show the other inequalities we need the following fact. For every integer k, the k unit vectors $\{v_1 \dots v_k\}$ that minimize the value of $\max_{i \ne j \in [1 \dots k]} \langle v_i, v_j \rangle$ are the vertices of a simplex in R^{k-1} centered at the origin. For each $i, j \in [1 \dots k]$, these vectors satisfy $\langle v_i, v_j \rangle = -\frac{1}{k-1}$.

Now to prove the inequality $COL_3(G) \le \chi(G)$, consider a k coloring σ of G. The coloring σ partitions the vertex set V into k color classes $\{V_1 \dots V_k\}$. Assigning each color class V_i with the corresponding vector v_i above, we obtain a valid assignment for $COL_3(G)$. To show that $\omega(G) \le COL_1(G)$, consider a graph G with maximal clique size $\omega(G)$. To obtain a valid assignment of vectors to $COL_1(G)$ of value k we require that all pairs of vectors corresponding to the vertices of the maximal clique will have inner product of value at most $-\frac{1}{k-1}$. As mentioned above, this can happen only if $\omega(G) \le k$.

Remark A.1 *The results of this paper show a large gap between $COL_1(G)$ and $\chi(G)$ (Theorem 1.1). Using the line of proof appearing in [16], it can be seen that similar gaps also apply to the functions $\omega(G)$ and $COL_3(G)$. Details omitted.*