Approximating the independence number via the ϑ-function

Noga Alon * Nabil Kahale †

Abstract

We describe an approximation algorithm for the independence number of a graph. If a graph on n vertices has an independence number $n/k + m$ for some fixed integer $k \geq 3$ and some $m > 0$, the algorithm finds, in random polynomial time, an independent set of size $\tilde{\Omega}(m^{3/(k+1)})$, improving the best known previous algorithm of Boppana and Halldorsson that finds an independent set of size $\Omega(m^{1/(k-1)})$ in such a graph. The algorithm is based on semi-definite programming, some properties of the Lovász ϑ-function of a graph and the recent algorithm of Karger, Motwani and Sudan for approximating the chromatic number of a graph. If the ϑ-function of an n vertex graph is at least $Mn^{1-2/k}$ for some absolute constant M, we describe another, related, efficient algorithm that finds an independent set of size k. Several examples show the limitations of the approach and the analysis together with some related arguments supply new results on the problem of estimating the largest possible ratio between the ϑ-function and the independence number of a graph on n vertices.

*Department of Mathematics, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, Israel and AT & T Bell Laboratories, Murray Hill, NJ 07974, USA. Email: noga@math.tau.ac.il. Research supported in part by a USA-Israel BSF grant and by the Fund for Basic Research administered by the Israel Academy of Sciences.

†AT & T Labs-Research, Murray Hill, NJ 07974. Email: kahale@research.att.com. This work was partly done while the author was at XEROX PARC, and partly at DIMACS.

[0]Keywords: Independence number of a graph, semi-definite programming, approximation algorithms.

1 Introduction

An *independent set* of a graph is a subset of vertices that contains no pair of neighbors. The *independence number* $\alpha(G)$ of a graph G is the size of a largest independent set in G. Determining or estimating $\alpha(G)$ is a fundamental problem in Theoretical Computer Science. The problem of computing $\alpha(G)$ is known to be NP-hard [19]. The best known approximation algorithm for the independence number, designed by Boppana and Halldorsson [7], has a performance guarantee of $O(n/(\log n)^2)$, where n is the number of vertices in the graph. Boppana and Halldorsson's algorithm performs better when the graph contains a large independent set. Indeed, they showed that if the independence number exceeds $n/k + m$, where k is a fixed integer and $m > 0$, then an independent set of size $\Omega(m^{1/(k-1)})$ can be found in polynomial time. On the negative side, it has recently been shown in [3], improving previous results in [11], [4], that for some $\epsilon > 0$ it is impossible to approximate in polynomial time the independence number of a graph within a factor of n^ϵ, assuming P $\neq$ NP. The exponent ϵ has since been improved under similar hardness assumptions, and very recently it has been shown by Håstad [16] that it is in fact larger than $(1-\delta)$ for every positive δ, assuming NP does not have polynomial time randomized algorithms.

Another fundamental quantity associated with a graph G is its *chromatic number* $\chi(G)$. A *proper coloring* of a graph is an assignment of colors to each vertex of the graph so that adjacent vertices have different colors. The chromatic number is the minimum number of colors used in a proper coloring. The best known approximation algorithm [15] for the chromatic number of a graph on n vertices has a performance guarantee of $O(n(\log\log n)^2/(\log n)^3)$.

In this paper we obtain an improved approximation algorithm for the independence number by considering the ϑ-function of the graph. This function, introduced by Lovász [23], can be defined as follows. Given a graph $G = (V, E)$, an *orthonormal labeling* (or *orthonormal representation*) of G is an assignment of a unit vector a_v in an Euclidean space to each vertex v of G, such that $a_u \cdot a_v = 0$ if $u \neq v$ and $(u, v) \notin E$. The ϑ-function $\vartheta(G)$ is equal to the minimum over all unit vectors d and all orthonormal labelings (a_v) of G of

$$\max_{v\in V} \frac{1}{(d \cdot a_v)^2}.$$

The ϑ-function satisfies the inequality

$$\alpha(G) \leq \vartheta(G) \leq \chi(\overline{G}),$$

where $\overline{G}$ is the complement of G. Moreover, the ϑ-function can be computed in polynomial time at an arbitrary precision [17]. The number $\chi(\overline{G})$ is also refered to as the *clique cover number* of the graph.

Here we study the gap between the ϑ-function and the independence number. We show in Section 3 that for any fixed integer $k \geq 3$, if $\vartheta(G) \geq n/k + m$ then $\alpha(G) \geq \tilde{\Omega}(m^{3/(k+1)})$. Here, and in what follows, the notation $g(n) = \tilde{\Omega}(f(n))$ means, as usual, that $g(n) \geq \Omega(f(n)/(\log n)^c)$ for some constant c independent of n. The notation $g(n) = \tilde{O}(f(n))$ is defined similarly. Our proof is algorithmic, that is, if $\vartheta(G) \geq n/k + m$ then an independent set of size $\tilde{\Omega}(m^{3/(k+1)})$ can be found in randomized polynomial time, thus improving Boppana and Halldorsson's result. Our proof and algorithm uses semi-definite programming, along the ideas in [17, 13], together with the recent work by Karger, Motwani and Sudan [18]. It is worth noting that the authors of [7] showed that no approximation algorithm (with an arbitrary running time) which is based on a subgraph exclusion procedure like most of the previous algorithms for the independent set problem (including the one in [7]), can approximate the maximum independent set as well as our algorithm here, showing that the application of some other tools is indeed crucial.

In Section 4 we show that if $g(n)$ is a function of n such that, for any graph G on n vertices, $\chi(\overline{G}) \leq g(n)\vartheta(G)$, then for any graph G on n vertices, $\vartheta(G) \leq H_n g(n)\alpha(G)$, where $H_n = 1 + 1/2 + \ldots + 1/n$ ($= O(\log n)$). This improves a recent result of Szegedy [24] by a $\log n$ factor.

In Section 5 we bound the ϑ-function of graphs with small independence number. We show that if $\alpha(G) < k$, then $\vartheta(G) \leq Mn^{1-2/k}$, where M is an absolute constant. This generalizes a result in [20] where the case $k = 3$ was treated (in a disguised form.) For $k = 3$ the above estimate is shown to be tight in [1]. We also show that if $\vartheta(G) > Mn^{1-2/k}$, then an independent set of size k can be found in polynomial time in n (independent of k). By a very recent result of Feige [10] that applies the randomized graph products technique of Berman and Schnitger [6], there are graphs G on n vertices with an independence number $\alpha(G) < k$ whose ϑ-function satisfies $\vartheta(G) \geq \Omega(n^{1-O(1/\log k)})$, showing that our $O(n^{1-2/k})$ upper bound is not very far from being best possible. We also generalize Kashin and Konyagin's result in a different direction by showing that if the complement of a graph G has no odd cycle of length at most $2s+1$, then $\vartheta(G) \leq 1 + (n-1)^{1/(2s+1)}$. This bound can be shown to be nearly tight by modifying the construction in [1].

In Section 6 we show that the result in Section 3 cannot be significantly improved by giving, for every $\epsilon > 0$, an explicit family of graphs on n vertices whose ϑ-function is at least $(\frac{1}{2} - \epsilon)n$ and whose independence number is $O(n^{\delta})$, where $\delta = \delta(\epsilon) < 1$. We note that this construction is tight in the sense that if the ϑ-function exceeds $(\frac{1}{2} + \epsilon)n$, then the independence number is $\Omega(n)$. Our construction is based on a combinatorial result of Frankl and Rödl [12] and extends a construction in [2]. The final Section 7 contains some concluding remarks and open problems.

2 The ϑ-function and Ramsey theory

For integers $k, s, n \geq 2$, let $r(k,s) = \binom{k+s-2}{k-1}$, and $t_k(n) = \max\{s | r(k,s) \leq n\}$. It is well known in Ramsey theory [9] that any graph G with at least $r(k,s)$ vertices contains either a clique of size k or an independent set of size s. Moreover, a clique of size k or an independent set of size s can be found in G in polynomial time (as a function of the input size.) Boppana and Halldorsson [7] show that if a graph G on n vertices contains an independent set of size $n/k + m$, then an independent set of size $t_k(m)$ can be found in polynomial time. Their strategy is to repeatedly delete from G a clique of size k until the remaining graph contains no such clique. Since the number of cliques removed is obviously at most n/k and since each clique contains at most one vertex from an independent set, the remaining graph has at least m vertices. Moreover, it contains no clique of size k. Thus an independent set of size $t_k(m)$ can be found in polynomial time in the remaining graph. Note that, for fixed k, $t_k(m) = \Omega(m^{1/(k-1)})$.

A careful look at Boppana and Halldorsson's algorithm yields the following.

Proposition 2.1 *If $\chi(\overline{G}) \geq n/k + m$, then an independent set of size $t_k(m)$ can be found in G polynomial time.*

Proof Each time a clique is removed from the graph, the clique cover number diminishes by at most 1. Since at most n/k cliques have been removed, the clique cover number of the remaining graph is at least m. Thus the remaining graph has at least m vertices. We conclude as before that an independent set of size $t_k(m)$ can be found in polynomial time in the remaining graph. □

Corollary 2.2 *If $\vartheta(G) \geq n/k + m$, then an independent set of size $t_k(m)$ can be found in G in polynomial time.*

3 Improved approximation for the independence number

When k is a fixed integer, we have the following.

Theorem 3.1 *For any fixed integer $k \geq 3$, if $\vartheta(G) \geq n/k + m$, then an independent set of size $\tilde{\Omega}(m^{3/(k+1)})$ can be found in randomized polynomial time.*

Note that as shown in [7] such an approximation algorithm cannot be based on a subgraph exclusion procedure as in [7]. A similar result can be proved for non integer values of k, but since its precise statement is somewhat cumbersome we omit it here. The need for the integrality of k is in the proof of the main result of [18], which can be modified to yield certain estimates for non-integral k as well.

The proof of Theorem 3.1 uses a recent result by Karger, Motwani and Sudan [18]. As defined in [18], the *vector chromatic number* of a graph is the minimum real number h such that there exists an assignment of a unit vector a_v to each vertex v satisfying the inequality $a_v \cdot a_w \leq -1/(h-1)$ whenever (v, w) is an edge. It is shown in [18] that if the vector chromatic number of a graph G on n vertices is at most h for some fixed integer $h \geq 3$, then G can be properly colored with at most $\tilde{O}(n^{1-3/(h+1)})$ colors in randomized polynomial time. Karger, Motwani and Sudan [18] also define the *strict vector chromatic number* as the minimum real number h such that there exists an assignment of unit vectors a_v to each vertex v satisfying the equality $a_v \cdot a_w = -1/(h-1)$ whenever (v, w) is an edge. They show that the strict vector chromatic number of a graph G is equal to $\vartheta(\overline{G})$. By definition, the vector chromatic number is always upper bounded by the strict vector chromatic number.

We now turn to the proof of Theorem 3.1. We first show that if $\vartheta(G) \geq n/k + m$, then G contains an independent set of size $\tilde{\Omega}(m^{3/(k+1)})$. It is known [23] that $\vartheta(G)$ is the maximum over all unit vectors d and all orthonormal labelings (b_v) of the complement $\overline{G}$ of G of $\sum_{v \in V}(d \cdot b_v)^2$, and that the maximum is attained. This characterization of the ϑ-function will be called the *dual characterization*. It implies immediately that $\alpha(G) \leq \vartheta(G)$. Indeed, if I is an independent set, then by setting $b_v = e$ for $v \in I$, where e is any unit vector, and by assigning an orthonormal family orthogonal to e to the remaining vertices, we get an orthonormal representation of $\overline{G}$. For this representation, it is clear that $\sum_{v \in V}(b_v \cdot e)^2 = |I|$.

Let d be a unit vector and (b_v) an orthonormal labeling of $\overline{G}$ such that $\vartheta(G) = \sum_{v \in V}(d \cdot b_v)^2$. We will use the family (b_v) to find a large independent set in G. Without loss of generality, label the vertices from 1 to n and assume that $(d \cdot b_1)^2 \geq (d \cdot b_2)^2 \geq \cdots \geq (d \cdot b_n)^2$. The inequalities $(d \cdot b_1)^2 + (d \cdot b_2)^2 + \cdots + (d \cdot b_n)^2 \geq n/k + m$ and $(d \cdot b_i)^2 \leq 1$ for $1 \leq i \leq m$ imply that $(d \cdot b_m)^2 \geq 1/k$. Let K be the subgraph of G induced on $\{1, 2, \ldots, m\}$. The family $(b_1, b_2, \ldots, b_m)$ is clearly an orthonormal labeling of $\overline{K}$. It follows from the definition of the ϑ-function in Section 1 that $\vartheta(\overline{K}) \leq k$. From the discussion in the beginning of this section, we conclude that the vector chromatic number of K is at most k, and thus K can be properly colored with $\tilde{O}(m^{1-3/(k+1)})$ colors in randomized polynomial time. The largest color class forms an independent set of K (and thus an independent set of G) of size $\tilde{\Omega}(m^{3/(k+1)})$.

To conclude the proof of the theorem, we show how to find in polynomial time a unit vector d and an orthonormal labeling $(b_v), v \in V$ of G such that $\sum_{v\in V}(d \cdot b_v)^2 \geq \vartheta(G) - 1$. (The preceding argument shows that this inequality suffices for our needs, since the same argument would still be valid by replacing m with $m-1$.) One way to achieve this goal is to use another characterization of the ϑ-function. Let B range over all positive semi-definite symmetric matrices indexed by V such that $\mathrm{tr}(B) = 1$ and $b_{uv} = 0$ whenever $(u, v) \in E, u \neq v$. Then $\vartheta(G)$ is the maximum over all such matrices [23] of $\sum_{u,v\in V} b_{uv}$. A matrix B satisfying the above conditions and such that $\sum_{u,v\in V} b_{uv} \geq \vartheta(G) - 1$ can be found in polynomial time in n using the ellipsoid method [17]. Since B is positive semi-definite, there exist vectors c_u such that $c_u \cdot c_v = b_{uv}$, for all $u, v \in V$. Let $b_v = c_v/||c_v||$, and

$$d = \frac{\sum_{v\in V} c_v}{||\sum_{v\in V} c_v||}.$$

Clearly, the family (b_v) is an orthonormal representation of $\overline{G}$. Moreover, it is shown implicitly in [23, Th. 5] that $\sum_{v\in V}(d \cdot b_v)^2 \geq \sum_{u,v\in V} b_{uv}$. Thus $\sum_{v\in V}(d \cdot b_v)^2 \geq \vartheta(G) - 1$. Note that, given B, the vectors c_v can be computed at an arbitrary precision in polynomial time using a Cholesky factorization [14, Sec. 4.2] of B. □

4 Comparing the worst-case ratios

It is shown in [24] that if $f(n)$ is a monotone increasing function such that, for every n and every graph G on n vertices $\vartheta(G) \leq \alpha(G)f(n)$ holds, then for every n and for every graph G on n vertices $\vartheta(G) \geq \chi(\overline{G})/(f(n)\log n)$ holds. It is also shown that if $g(n)$ is a monotone increasing function such that, for every n and every graph G on n vertices $\vartheta(G) \geq \chi(\overline{G})/g(n)$ holds, then for every n and for every graph G on n vertices $\vartheta(G) \leq 8\log^2 n g(n)\alpha(G)$ holds. We improve the latter result by a logarithmic factor and observe that it is not needed to require that g be monotone.

Theorem 4.1 *Let $g(n)$ be a function of n such that, for any graph G on n vertices, $\chi(\overline{G}) \leq g(n)\vartheta(G)$. Then for any graph G on n vertices, $\vartheta(G) \leq H_n g(n)\alpha(G)$.*

Proof Without loss of generality, assume that $g(n)$ is the maximum over all graphs G on n vertices of the ratio $\chi(\overline{G})/\vartheta(G)$. Consider the operation of adjoining an extra vertex to a graph by connecting it to every vertex. It is known (see e.g. [21, p. 20]) that the ϑ-function remains the same under this operation. It is also easy to see that the clique cover number remains the same. It follows that $g(n)$ is an increasing function of n.

Let G be a graph on n vertices. Following the notation of Section 3, let d be a unit vector, let (b_v) be an orthonormal labeling of $\overline{G}$ such that $\vartheta(G) = \sum_{v\in V}(d \cdot b_v)^2$, and assume that $(d \cdot b_1)^2 \geq (d \cdot b_2)^2 \geq \cdots \geq (d \cdot b_n)^2$. Note that $(d \cdot b_i)^2 \geq \vartheta(G)/(H_n i)$, for some $i \in [1, n]$. This is because otherwise $(d \cdot b_i)^2 < \vartheta(G)/(H_n i)$ for all i, and by summing these inequalities for $1 \leq i \leq n$ we get a contradiction. Let K be the subgraph of G induced on $\{1, 2, \ldots, i\}$. The definition of the ϑ-function in Section 1 shows that $\vartheta(\overline{K}) \leq 1/(d \cdot b_i)^2 \leq H_n i/\vartheta(G)$. Since $\chi(K) \leq g(i)\vartheta(\overline{K}) \leq g(n)\vartheta(\overline{K})$, we conclude that $\chi(K) \leq g(n)H_n i/\vartheta(G)$. Thus K contains an independent set of size $i/\chi(K) \geq \vartheta(G)/(H_n g(n))$. Hence $\alpha(G) \geq \vartheta(G)/(H_n g(n))$, as desired. □

5 Graphs with a small independence number

Kashin and Konyagin [20] show (in a disguised form) that for any graph G on n vertices, if $\alpha(G) < 3$ then $\vartheta(G) \leq 2^{2/3} n^{1/3}$. We generalize their result for larger bounds on $\alpha(G)$, and also for graphs whose complement contains no short odd cycles.

Theorem 5.1 *There exists an absolute constant M such that for any graph $G = (V, E)$ on n vertices and any integer $k \geq 2$, if $\alpha(G) < k$ then $\vartheta(G) \leq M n^{1-2/k}$.*

The proof of Theorem 5.1 is based on the following lemma.

Lemma 5.2 *Let $f_k(n) = \max \vartheta(H)$, where H ranges over all graphs on n vertices satisfying the condition $\alpha(H) < k$. If $G = (V, E)$ is a graph such that $\alpha(G) < k$ and Δ is the maximum degree of $\overline{G}$, then $\vartheta(G) \leq 1 + \sqrt{\Delta f_{k-1}(\Delta)}$.*

Proof The ϑ-function can be shown (see e.g. [21]) to be equal to the maximum over all orthonormal labelings (b_v) of $\overline{G}$ of the largest eigenvalue of the matrix $(b_u \cdot b_v)$ indexed by the vertices of G.

For $u \in V$, let H_u be the subgraph of G induced on the set of neighbors of u in $\overline{G}$. Since H_u contains no neighbor of u (in G), $\alpha(H_u) < k - 1$. Moreover, H_u has at most Δ vertices. The argument in Section 4 shows that $f_k(n)$ is an increasing function of n. Thus $\vartheta(H_u) \leq f_{k-1}(\Delta)$.

Since $(b_v), v \in H_u$ is an orthonormal labeling of $\overline{H_u}$, it follows from the dual characterization of the ϑ-function (taking $d = b_u$) that $\sum_{v \in H_u} (b_u \cdot b_v)^2 \leq \vartheta(H_u) \leq f_{k-1}(\Delta)$. By the Cauchy-Schwartz inequality, it follows that $\sum_{v \in H_u} |b_u \cdot b_v| \leq \sqrt{\Delta f_{k-1}(\Delta)}$. On the other hand, since (b_v) is an orthonormal labeling of $\overline{G}$, $b_u \cdot b_v = 0$ if $v \neq u$ and v is not in H_u. We conclude that $\sum_{v \in V} |b_u \cdot b_v| \leq 1 + \sqrt{\Delta f_{k-1}(\Delta)}$, for any $u \in V$. Consequently, the largest eigenvalue of the matrix $(b_u \cdot b_v)$ is at most $1 + \sqrt{\Delta f_{k-1}(\Delta)}$. Since this inequality holds for all orthonormal labelings of G, we conclude that $\vartheta(G) \leq 1 + \sqrt{\Delta f_{k-1}(\Delta)}$. $\square$

Fact 5.3 *If the vertex set V of a graph G is split into l pairwise disjoint subsets $V_1, V_2, \ldots, V_l$, for an integer $l \geq 1$, the ϑ-function of G is upper bounded by the sum of the ϑ-functions of the subgraphs induced on the V_i, $1 \leq i \leq l$.*

Proof This follows immediately from the dual characterization of the ϑ-function. $\square$

We are now ready to prove Theorem 5.1. The proof is by induction on k. The base case $k = 2$ is easy since the ϑ-function of the complete graph is 1. Assume now that the induction hypothesis holds for $k - 1$, that is $f_{k-1}(n) \leq M n^{1-2/(k-1)}$, where M is a constant to be determined later. We prove by induction on n that $f_k(n) \leq M n^{1-2/k}$. Since $\vartheta(G) \leq n$, the inequality $f_k(n) \leq M n^{1-2/k}$ is trivial when $n \leq M^{k/2}$. Assume now that $f_k(m) \leq M m^{1-2/k}$ for $m < n$. Let G be a graph on n vertices such that $\alpha(G) < k$, and define $\Delta = 9 n^{1-1/k}$. Assume for simplicity that Δ is an integer. We distinguish two possible cases:

1. The maximum degree of $\overline{G}$ is at most Δ. By Lemma 5.2 and the induction hypothesis,

$$\begin{aligned} \vartheta(G) &\leq 1 + \sqrt{\Delta f_{k-1}(\Delta)} \\ &\leq 1 + \sqrt{\Delta M \Delta^{1-2/(k-1)}} \\ &\leq 1 + 9\sqrt{M} n^{1-2/k} \\ &\leq M n^{1-2/k}, \end{aligned}$$

where the last inequality holds if M is a sufficiently large constant.

2. There exists a vertex u of G that has more than Δ neighbors in $\overline{G}$. Let $U \subset V$ be a subset of Δ neighbors of u in $\overline{G}$, H the subgraph of G induced on U, and K the subgraph of G induced on $V - \{u\} - U$. It follows from Fact 5.3 that $\vartheta(G) \leq 1 + \vartheta(H) + \vartheta(K)$. But $\vartheta(H) \leq f_{k-1}(\Delta)$ since $\alpha(H) < k-1$, and $\vartheta(K) \leq f_k(n - \Delta - 1)$. Thus

$$\begin{aligned}\vartheta(G) &\leq 1 + M\Delta^{1-2/(k-1)} + f_k(n-\Delta) \\ &\leq M\frac{k}{k-2}\Delta^{1-2/(k-1)} + M(n-\Delta)^{1-2/k}.\end{aligned}$$

(The second inequality holds since we are assuming $n \geq M^{k/2}$, and thus $\Delta^{1-2/(k-1)} \geq M^{(k-3)/2}$. So it suffices that $2M^{(k-1)/2} \geq k-2$, for all $k \geq 3$.) Since $n^{1-2/k}$ is a concave function of n, $(n-\Delta)^{1-2/k} \leq n^{1-2/k} - \Delta(1-2/k)n^{-2/k}$. Hence

$$\vartheta(G) \leq Mn^{1-2/k} - M\Delta(1-2/k)n^{-2/k} + M\frac{k}{k-2}\Delta^{1-2/(k-1)}.$$

But $\Delta^{-2/(k-1)} = 9^{-2/(k-1)}n^{-2/k} \leq (1-2/k)^2 n^{-2/k}$. This is because $\left(\frac{k}{k-2}\right)^{k-1} \leq 9$, for $k \geq 3$. We conclude that $\vartheta(G) \leq Mn^{1-2/k}$, and thus $f_k(n) \leq Mn^{1-2/k}$, as desired. □

Theorem 5.4 *If G is a graph on n vertices such that $\vartheta(G) > M'n^{1-2/k}$ for an appropriate absolute constant M', an independent set in G of size k can be found in polynomial time.*

Proof This follows from the proof of Theorem 5.1. □

Corollary 5.5 *If $u_1, u_2, \ldots, u_n$ are n unit vectors, and among any k of them some 2 are orthogonal, then $||\sum_{i=1}^n u_i|| \leq \sqrt{M}n^{1-1/k}$.*

Proof Consider the graph G on $\{1, 2, \ldots, n\}$, where (i,j) is an edge if and only if $u_i \cdot u_j = 0$. It is clear that (u_i) is an orthonormal representation of $\overline{G}$. Thus the largest eigenvalue of the matrix $P = (u_i \cdot u_j)$ is at most $\vartheta(G)$. In particular, $\sum_{ij} u_i \cdot u_j = \mathbf{1} \cdot P\mathbf{1} \leq n\vartheta(G)$. Equivalently, $||\sum u_i||^2 \leq n\vartheta(G)$. But $\alpha(G) < k$ by hypothesis, and so $\vartheta(G) \leq Mn^{1-2/k}$. Combining this with the preceding inequality we get the desired result. □

Corollary 5.5 has already been established [22, 20] for the special case $k = 3$. For this case it is tight up to a constant factor, as shown in [1].

It follows from Theorem 5.4 that if the independence number exceeds $M'n^{1-2/k}$, an independent set in G of size k ($\leq \log n$) can be found in polynomial time. A simpler algorithm can be used to achieve a slightly stronger result, however, following the ideas in [5]. Partition the vertices of the graph into $M'n^{1-2/k}/Ck$ subsets, each of size $Ckn^{2/k}/M'$, where $C > 0$ is any constant. The hypothesis implies that at least one of these subsets contains an independent set of size Ck. We can search for such an independent set in each of these subsets by brute-force search. The running time of the algorithm is polynomial since

$$\binom{\frac{Ckn^{2/k}}{M'}}{k} \leq n^{O(C)}.$$

5.1 Graphs with no short odd cycles

We give in this subsection another generalization of Kashin and Konyagin's aforementioned result.

Proposition 5.6 *Let G be a graph on a set of n vertices. If the complement of G has no odd cycle of length at most $2s+1$, the ϑ-function of G does not exceed $1+(n-1)^{1/(2s+1)}$.*

Proof Again, we use the fact that the ϑ-function is equal to the maximum over all orthonormal labelings (b_v) of $\overline{G}$ of the largest eigenvalue of the matrix $B = (b_u \cdot b_v)$ indexed by the vertices of G. Let (b_v) be an orthonormal labeling of $\overline{G}$ that achieves this maximum. Since b_{uv}, $u \neq v$, is non-zero only if $(u,v) \in \overline{G}$, the absence of odd cycles of length at most $2s+1$ in $\overline{G}$ implies that every diagonal entry of the matrix $(B-I)^{2s+1}$ is zero. In particular, $\mathrm{tr}((B-I)^{2s+1}) = 0$. Let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ be the eigenvalues of B. It follows that $\sum_{i=1}^{n}(\lambda_i - 1)^{2s+1} = 0$. Since B is positive semi-definite, the λ_i's are non-negative. Thus $(\lambda_1 - 1)^{2s+1} \leq n-1$, and so $\lambda_1 \leq 1 + (n-1)^{1/(2s+1)}$, as desired. □

Remark. Up to a multiplicative constant factor c_s depending on s, the bound in Proposition 5.6 can be shown to be tight by modifying the construction in [1].

6 Linear ϑ-function and sublinear independence number

By Corollary 2.2, if the ϑ-function of a graph on n vertices is at least $(\frac{1}{2}+\epsilon)n$, then the independence number is $\Omega(n)$. In this section we show that if ϑ is slightly smaller, then the independence number may be n^{δ} for some $\delta < 1$.

Theorem 6.1 *For every $\epsilon > 0$ there is an explicit family of graphs on n vertices whose ϑ-function is at least $(\frac{1}{2}-\epsilon)n$ and whose independence number is $O(n^{\delta})$, where $\delta = \delta(\epsilon) < 1$.*

The construction is based on a combinatorial result of Frankl and Rödl [12] and extends a construction in [2]. In fact, by interpreting the result in [2] appropriately one may note that it supplies (in a disguised form) graphs with n vertices, $\vartheta \geq n/16$ and independence number at most $O(n^{0.85002})$).

Proof of Theorem 6.1 For a pair of integers $q > s > 0$ let $G(q,s)$ denote the graph on $n = \binom{2q}{q}$ vertices corresponding to all q-subsets of the $2q$-element set $Q = \{1, 2, \ldots, 2q\}$, where two vertices are adjacent iff the intersection of their corresponding subsets is of cardinality precisely s. By the main result of Frankl and Rödl in [12], for every $\gamma > 0$ there is a $\mu = \mu(\gamma) > 0$ so that if $(1-\gamma)q > s > \gamma q$ then every family of more than $2^{2q(1-\mu)}$ subsets of cardinality q of Q contains some pair of subsets whose intersection is of cardinality s. This means that the independence number of the graph $G(q,s)$ for q and s that satisfy $(1-\gamma)q > s > \gamma q$ satisfies

$$\alpha(G(q,s)) \leq n^{\delta} \tag{1}$$

for some $\delta = \delta(\gamma) < 1$.

We next estimate the ϑ-function of $G(q,s)$. Let

$$x = \frac{(q-s) + \sqrt{(q-s)^2 - s^2}}{s}$$

be the bigger root of the quadratic polynomial $sx^2 - 2(q-s)x + s$. Associate with every vertex u of $G(q,s)$ that corresponds to a subset U of cardinality q of Q the vector $d_u = (x+1) \cdot 1_U - 1_Q$, where

1_U is the characteristic vector of U and 1_Q is the all 1 vector of length $2q$. Define $b_u = d_u/||d_u||$. A simple calculation shows that if u corresponds to subset U and v corresponds to subset V then $d_u \cdot d_v = |U \cap V|(x+1)^2 - 2qx$. In particular, $||d_u||^2 = qx^2 + q$. It also follows that the vectors b_u form an orthonormal labeling of $\overline{G(q,s)}$. Therefore, by the dual characterization of the ϑ-function and by letting d be the unit vector $\frac{1}{\sqrt{2q}}(1,1,\ldots,1)$ we conclude that

$$\vartheta(G(q,s)) \geq \sum_u (d \cdot b_u)^2 = n\frac{(qx-q)^2}{2q(qx^2+q)} = n\frac{q-2s}{2(q-s)},$$

since

$$(d \cdot b_u)^2 = \frac{(qx-q)^2}{2q(qx^2+q)}$$

for every vertex u of $G(q,s)$.

Given $\epsilon > 0$ we can now choose s to be the largest integer smaller than $q/2$ for which

$$\frac{q-2s}{2(q-s)} > (\frac{1}{2} - \epsilon).$$

Is is easy to check that for this s, $s/q > \gamma$ for an appropriate positive $\gamma = \gamma(\epsilon)$ and the desired result now follows from (1). □

7 Concluding Remarks

1. Any polynomial approximation algorithm that finds, in any n-vertex graph with independence number at least n/k, an independent set of size at least $\tilde{\Omega}(n^{\alpha_k})$ easily supplies a polynomial algorithm for coloring any k-colorable graph on n vertices by $\tilde{O}(n^{1-\alpha_k})$ colors. Indeed, this is done by simply applying the independence algorithm repeatedly. It follows that any improvement in the exponent of m in Theorem 3.1 will improve the exponent in the coloring algorithm of [18] (and vice versa, of course, as follows from the proof of Theorem 3.1). Note that the algorithm in Theorem 3.1 works for any graph with a large ϑ-function and, similarly, the algorithm of [18] works for any graph with a sufficiently small value of the ϑ-function of its complement. Therefore, the performance of both algorithms may be improved with a better understanding of the largest possible value of the ϑ-function of a graph on n vertices with a given independence number. It would be interesting to decide if this largest possible value is closer to the upper bounds provided for it by our results in Sections 3 and 5, or is closer to the lower bound given for it in [10]. A proof that the latter possibility holds would supply improved approximation algorithms for the independence number and chromatic number of a graph.

2. Estimating the largest possible ratio between the ϑ-function and the independence number of graphs on n vertices remains open, despite some recent progress. It is known [8] that $\chi(\overline{G}) \leq \alpha(G)n/\log^2 n$ for any graph on n vertices. As a consequence, $\vartheta(G) \leq \alpha(G)n/\log^2 n$. While the first inequality is tight up to a constant (e.g. for random graphs), it is an open question whether the same holds for the second inequality. The result of Feige [10] shows that there are graphs on n vertices for which this ratio is at least $\Omega(n/2^{O(\sqrt{\log n})})$, and it would be interesting to decide how tight this estimate is.

3. Improving a result in [20], it is shown in [1] that the bound in Corollary 5.5 is tight (up to a constant factor) when $k = 3$. Whether this is the case for higher values of k is another open question.

Acknowledgement We would like to thank Mario Szegedy for helpful discussions, suggestions and comments.

References

[1] N. Alon. Explicit Ramsey graphs and orthonormal labelings. *The Electronic Journal of Combinatorics*, 1 (1994), R12, 8pp.

[2] N. Alon and Y. Peres. Euclidean Ramsey Theory and a construction of Bourgain. *Acta Math. Hungar.*, 57:61–64, 1991.

[3] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and intractability of approximation problems. In *Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science*, pages 14–23. IEEE Computer Society Press, 1992.

[4] S. Arora and S. Safra. Probabilistic checking of proofs; a new characterization of NP. In *Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science*, pages 2–13. IEEE Computer Society Press, 1992.

[5] B. Berger and J. Rompel. A better performance guarantee for approximate graph coloring. *Algorithmica*, 5(4):459–466, 1990.

[6] P. Berman and G. Schnitger. On the complexity of approximating the independent set problem. *Information and Computation*, 96:77–94, 1992.

[7] R. Boppana and M. M. Halldorsson. Approximating maximum independent sets by excluding subgraphs. *BIT*, 32:180–196, 1992.

[8] P. Erdös. Some remarks on chromatic graphs. *Colloquium Mathematicum*, 16:253–256, 1967.

[9] P. Erdös and G. Szekeres. A combinatorial problem in geometry. *Compositio Mathematica*, 2:463–470, 1935.

[10] U. Feige. Randomized graph products, chromatic numbers, and the Lovász ϑ-function. In *27th Annual ACM Symposium on Theory of Computing*, pages 635–640. ACM Press, 1995.

[11] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating Clique is almost NP-complete. In *Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science*, pages 2–12. IEEE Computer Society Press, 1991.

[12] P. Frankl and V. Rödl. Forbidden intersections. *Trans. AMS*, 300:259–286, 1987.

[13] M. Goemans and D. Williamson. Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming. *J. ACM*, 42:1115–1145, 1995.

[14] G. H. Golub and C. F. Van Loan, **Matrix Computations**, The Johns Hopkins University Press, 1989.

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

[16] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. Proc. 37^{th} IEEE FOCS, IEEE (1996), 627 – 636.

[17] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1:169–197, 1981.

[18] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semi-definite programming. In *35th Symposium on Foundations of Computer Science*, pages 2–13. IEEE Computer Society Press, 1994.

[19] R. Karp. *Reducibility among combinatorial problems.* Plenum Press, New York, 1972. Miller and Thatcher (eds).

[20] B. S. Kashin and S. V. Konyagin. On systems of vectors in a Hilbert space. *Trudy Mat. Inst. imeni V. A. Steklova*, 157:64–67, 1981. English translation in: *Proceedings of the Steklov Institute of Mathematics* (AMS 1983), 67–70.

[21] D. E. Knuth. The sandwich theorem. *The Electronic Journal of Combinatorics*, 1(A1):48pp., 1994.

[22] S. V. Konyagin. Systems of vectors in Euclidean space and an extremal problem for polynomials. *Mat. Zametki*, 29:63–74, 1981. English translation in: *Mathematical Notes of the Academy of the USSR,* 29:33–39, 1981.

[23] L. Lovász. On the Shannon capacity of a graph. *IEEE Transactions on Information Theory IT-25*, pages 1–7, 1979.

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