COMBINATORICA
Bolyai Society – Springer-Verlag

COMBINATORICA **20** (3) (2000) 393–415

ON THE HARDNESS OF APPROXIMATING THE CHROMATIC NUMBER*

SANJEEV KHANNA†, NATHAN LINIAL, SHMUEL SAFRA‡

Received April 11, 1997
Revised June 10, 1999

We study the hardness of approximating the chromatic number when the input graph is k-colorable for some fixed $k \geq 3$. Our main result is that it is NP-hard to find a 4-coloring of a 3-chromatic graph. As an immediate corollary we obtain that it is NP-hard to color a k-chromatic graph with at most $k+2\lfloor k/3\rfloor - 1$ colors. We also give simple proofs of two results of Lund and Yannakakis [20]. The first result shows that it is NP-hard to approximate the chromatic number to within n^ϵ for some fixed $\epsilon > 0$. We point here that this hardness result applies only to graphs with large chromatic numbers. The second result shows that for any positive constant h, there exists an integer k_h, such that it is NP-hard to decide whether a given graph G is k_h-chromatic or any coloring of G requires $h \cdot k_h$ colors.

1. Introduction

A *legal coloring* of a graph G is an assignment of colors to its vertices such that every adjacent pair of vertices gets a different color. A variety of problems arising in practice can be modeled as the problem of finding a legal coloring of a certain graph using the smallest possible number of colors. As an example, consider the following computational problem: given a set of tasks to perform where some of the tasks are pair-wise conflicting (say they cannot be carried out at the same time or at the same place), find a partition

Mathematics Subject Classification (1991): 68Q17, 68Q25, 68R10

* A preliminary version of this paper appeared in the *Proceedings of the Second Israel Symposium on Theory and Computing Systems*, 1993, pp. 250–260.

† This work was done when the author was a graduate student at Stanford University.

‡ Part of this work was carried out while this author was visiting the Hebrew University.

of the set of tasks into a minimum number of subsets such that no subset in the partition contains two conflicting tasks. This problem is equivalent to finding a legal coloring with a minimum number of colors: consider the graph whose vertices are all tasks, where two vertices are connected if and only if the corresponding tasks are conflicting. A coloring of this graph is a non-conflicting partition of the set of tasks.

A k-coloring of a graph G is a legal coloring of G that uses only k different colors and the chromatic number of G, denoted $\chi(G)$, is the least integer k for which a k-coloring of G exists. Karp [17] showed for any $k \geq 3$, it is NP-hard to determine if a graph is k-colorable. In light of this result, it is natural to ask how well can one approximate the chromatic number in polynomial time. The best known approximation guarantee for general graphs is $O(n(\log\log n)^2/\log^3 n)$ [15]. Garey and Johnson [12] proved that it is NP-hard to approximate the chromatic number within a factor of $(2-\epsilon)$ for any $\epsilon > 0$. Later, Linial and Vazirani [19] used graph products to give evidence that for n-vertex graphs, the approximation ratio is either below $\log^{1+o(1)} n$ or above $n^{\Omega(1)}$. In the early 90's, the discovery of a surprising connection between Probabilistically Checkable Proofs (PCPs) and hardness of approximations [9,2,1], led to dramatically improved hardness result for many classical optimization problems, including graph coloring. Lund and Yannakakis [20] showed that chromatic number is hard to approximate within n^ϵ for some constant $\epsilon > 0$. More recently, Feige and Kilian [10], building on the strong PCP constructions due to Håstad [16], have shown that the chromatic number is hard to approximate within a factor of $n^{1-\epsilon}$ for any constant $\epsilon > 0$, unless NP $\subseteq$ ZPP. Thus at this point the approximability of chromatic number on general graphs is essentially well-understood.

However, the above hardness results do not tell us anything about the hardness of approximating chromatic number when the input graphs are k-colorable for some fixed constant $k \geq 3$. In fact much better algorithmic results are known for such graphs. Wigderson [21] gave a simple algorithm to color k-colorable graphs with $O(n^{1/(k-1)})$ colors. Blum [3,4], extending the work of Wigderson [21] provided a polynomial-time algorithm which finds a legal coloring of a 3-colorable graph using at most $O(n^{3/8}\log^{\frac{5}{8}} n)$ colors and more generally, a coloring of k-colorable graphs using $O(n^{1-1/(k-4/3)}\log^{8/5} n)$ colors. More recently, using semidefinite programming techniques, Karger, Motwani and Sudan [18], gave a polynomial-time algorithm which can color a k-colorable graph using $\min\{O(\Delta^{1-2/k}), O(n^{1-3/(k+1)})\}$ colors where Δ is the maximum vertex degree in the graph. Blum and Karger recently improved this guarantee to $\tilde{O}(n^{3/14})$ for 3-colorable graphs [7]. Yet, until now, no hardness of approximation result was known for 3-colorable graphs. In fact, the only hardness result known for k-colorable graphs when k is a constant is the following result due to Lund and Yannakakis [20]. For every constant h there exists a constant k_h such that it is NP-hard to color a k_h-colorable graph with $k_h \cdot h$ colors. However, k_h depends on h and is relatively large in comparison. Hence this result is not applicable to small fixed values of the chromatic number.

Our results and techniques. The main focus of this paper is to study the hardness of approximating chromatic number on graphs that are k-chromatic for some constant $k \geq 3$. Our main result is that it is NP-hard to decide whether $\chi(G) \leq 3$ or $\chi(G) \geq 5$ for a graph G. In other words, we show that it is NP-hard to find a 4-coloring of a 3-colorable graph. Our proof uses the result that clique is hard to approximate to within any constant factor, a consequence of the PCP theorem [2,1]. In a very recent development, Guruswami and Khanna [14], have discovered a new proof of this result that does not rely on the PCP theorem. We also note that a straightforward corollary of our main result is that for *any* fixed $k \geq 3$, it is NP-hard to color a k-colorable graph with $k + 2\lfloor \frac{k}{3} \rfloor - 1$ colors.

In addition to this result, we also present a simpler proof of the n^{ϵ}-hardness result of Lund and Yannakakis [20] (for the case when the chromatic number can be arbitrarily large). Building on the construction that we present here, Bellare and Sudan [8] designed a more efficient reduction, obtaining a stronger hardness result (better ϵ) than implied by either our construction or that of [20]. Also subsequent to our work, Fürer [11] used a geometric construction to obtain another simpler proof of the n^{ϵ}-hardness for approximating the chromatic number.

Finaly, we show that our construction above can be modified to also obtain a simpler proof of another hardness result of Lund and Yannakakis [20], namely, for every constant h there exists a constant k_h such that it is NP-hard to color a k_h-colorable graph with $k_h \cdot h$ colors.

The paradigm which underlies all our reductions can be informally described as follows. We start with a family of graphs such that given a member graph G, it is NP-hard to decide whether G has a “large” clique or if every clique in G is “small”. We construct a transformation τ that takes any such graph G to a graph H such that a large clique in G maps to a collection of large cliques in H which together cover the vertices of H. On the other hand, the transformation satisfies the property that if G does not have a large clique, then every clique in H is small as well (and hence H has a large clique cover number). Thus the problem of deciding whether or not G has large clique translates into the problem of deciding whether or not H has a small clique cover (chromatic number). This gives an appropriate hardness result for approximating the chromatic number of a graph. The range of chromatic number values to which the hardness result applies and the hardness of approximation factor obtained, depends on the properties of the transformation τ. As will be evident in the remainder of this paper, much stronger hardness results can be shown when the chromatic number itself is allowed to take large values.

Notation. We use $\omega(G)$ to denote the size of the largest clique in a graph G. The clique-cover number of a graph G is denoted by $\bar{\chi}(G)$. We use $\bar{G}$ to denote the complement graph of G. Thus $\bar{\chi}(G)=\chi(\bar{G})$. Finally, we denote by $[k]$ the set $\{0,\dots,k-1\}$.

Organization. In Sections 2 and 3, we present simpler proofs of the two theorems of Lund and Yannakakis [20]. In Section 4, we prove our main result, namely, it is NP-hard to color a 3-colorable graph using only 4 colors.

2. Hardness result for unrestricted values

In this section we give a simple proof of the result that chromatic number is hard to approximate to within a factor of n^{ϵ} when the optimal value itself may be $n^{\epsilon'}$ for some constant $0<\epsilon'<1$.

We start with r-partite graph instances G, that is, the vertices of G can be partitioned into r rows where each row is an independent set. Let q be the maximum number of vertices in any row of G. For such instances, it is NP-hard to determine whether $\omega(G)=r$ or $\omega(G)<\frac{r}{q^{\epsilon}}$ for any $q=o(r)$ and some constant $\epsilon>0$ [9, 2, 1, 16]. Given such a graph G with $q=r^{\delta}$ for some constant δ, $0<\delta<1$, we construct an r-partite graph H with rq' vertices where $q'=(rq)^{O(1)}$ such that $\bar{\chi}(H)=q'$ when $\omega(G)=r$ and $\bar{\chi}(H)>q'\cdot q^{\epsilon}$ otherwise. Since $\bar{\chi}(H)=\chi(\bar{H})$, our result follows from the construction of such a graph H.

2.1. Structure and properties of H

The graph H can be intuitively described as a transformation of G which makes it symmetric under "rotations". In particular, any r-clique in G would map to several symmetric images in H which together cover the vertices of H. On the other hand, the transformation also ensures that $\omega(H)=\omega(G)$.

Specifically, the graph H is an r-partite graph such that each row of H corresponds to a distinct row of G. The graph H satisfies the following two properties:

1. if $\omega(G)=r$ then $\bar{\chi}(H)=q'$, and
2. $\omega(H)=\omega(G)$.

By the first property above, if $\omega(G) = r$ then $\bar{\chi}(H) = q'$. Otherwise, $\omega(G) < \frac{r}{q^\epsilon}$ and by the second property above, $\bar{\chi}(H) \geq (rq')/\omega(G) > q' \cdot q^\epsilon$. Thus construction of a H that satisfies the above two properties will give us the desired hardness result.

2.2. Construction of H

The vertex set. The vertex set of H is defined via an image function that maps vertices in any row of G to vertices in the corrsponding row of H. If a vertex v in G maps to a vertex v' in H, we say that v' is the *image vertex* of v.

The edge set. The edge set of H is constructed as follows. For every edge (u,v) in G, there is an edge (u',v') in H where u' and v' are the images of u and v respectively. We refer to such an edge as a *direct* edge. In addition, the edge set of H is extended to include all the *rotations* of the above edges as follows. If H contains an edge connecting the $i^{\underline{\text{th}}}$ vertex in row k to the $j^{\underline{\text{th}}}$ vertex in row l, then we also add to H edges connecting the $(i+a \bmod q')^{\underline{\text{th}}}$ vertex in row k to the $(j+a \bmod q')^{\underline{\text{th}}}$ vertex in row l for every $m \in \{1,\dots,q'-1\}$.

The image function. In order to complete our description of H, it remains to describe the image function, which maps each vertex v of G to a unique vertex in H. We start by showing the existence of an injection T, of some special structure, from any domain $[n]$ to a range $[m]$, when $m = n^{O(1)}$ is sufficiently large.

Lemma 2.1. *For every positive integer n, there exists a function $T\colon [n] \to [m]$ where $m = \Omega(n^5)$, such that for every distinct multiset $\{i_1,i_2,i_3\}$, $i_1,i_2,i_3 \in [n]$, the sum $T(i_1)+T(i_2)+T(i_3) \bmod m$ is distinct. Moreover, this implies that the property holds for distinct multisets of size two as well.*

Proof. We use an inductive argument to show the existence of a mapping which satisfies the property that for every distinct multiset $\{i_1,i_2,i_3\}$, $i_1,i_2,i_3 \in [n]$, the sum $T(i_1)+T(i_2)+T(i_3) \bmod m$ is distinct. Suppose that $T(0),\dots,T(l-1)$ have already been selected from $[m]$ such that for $0 \leq i_1,i_2,i_3 \leq l-1$, all the sums $T(i_1)+T(i_2)+T(i_3) \bmod m$ are distinct. An element $y \in [m]$ cannot be chosen for $T(l)$ if and only if there are $0 \leq i_1,i_2,i_3,i_4,i_5 \leq l-1$ such that

$$T(i_1)+T(i_2)+T(i_3) \equiv T(i_4)+T(i_5)+y \pmod{m}.$$

Therefore, at most l^5 elements y are ineligible at any step and if $m = \Omega(n^5)$, then the process can be carried through to yield the desired mapping.

Now observe that the mapping T shown to exist above must also satisfy the property that for every distinct multiset of size two, namely $\{i_1, i_2\}$, the sum $T(i_1)+T(i_2)$ mod m is distinct. This follows because if there exists two distinct mulisets of size two, namely $\{i_1, i_2\}$ and $\{j_1, j_2\}$ such that

$$T(i_1) + T(i_2) \equiv T(j_1) + T(j_2) \pmod{m}$$

then we can simply add the same element y to both multisets and get a contradiction to the property proven above that the sums of all triplets are distinct. Finally, it is clearly the case that the mapping T is an injection. ∎

Applying the image function. Let T be the injection shown to exist above. We apply T to the domain of all vertices of G, i.e., $n = rq$, and let $q' = m$, the size of the range of T. That is,

$$T: \{0, \ldots, rq - 1\} \to \{0, \ldots, q' - 1\}.$$

The *image* of a vertex v in the i^{th} row of G is the vertex of H which is in the i^{th} row and column $T(v)$ of H.

2.3. Establishing the properties of H

We now show that H satisfies the two properties mentioned in Section 2.1.

Lemma 2.2. *If $\omega(G) = r$ then $\bar{\chi}(H) = q'$.*

Proof: The graph H is *symmetric under rotation*; that is, if we let $s^j(H)$ be the graph that results by rotating all the rows of H by j columns to the right (in a wraparound manner), then for every $j \in \{0, .., q'\}$, $s^j(H) = H$. Now given a clique C of size r in G, consider the set C' which consists of the images in H of the vertices in C. Clearly, C' is a clique of H, and due to the symmetry of H under rotation, C' has $q' - 1$ *rotational images* that together cover all the vertices of H. ∎

We next show that our transformation preserves the maximum clique size.

Lemma 2.3. $\omega(G) = \omega(H)$.

Proof. Since the image of a clique in G forms a clique in H, it is easily seen that $\omega(H)\geq\omega(G)$. It remains to show that $\omega(H)\leq\omega(G)$.

An edge (u,v) in G is said to be the *origin* of an edge $e=(u',v')$ in H if either (i) $u'=T(u)$ and $v'=T(v)$, or (ii) e is obtained via a rotation of the edge $(T(u),T(v))$. Using the properties of the image function T, we first show that every edge in H has a unique origin in G. Suppose, to the contrary, that an edge in H has two origins in G, namely the edges (u_1,v_1) and (u_2,v_2). Then we must have

$$T(u_1)-T(v_1)\equiv T(u_2)-T(v_2) \pmod{q'}.$$

Then it follows that

$$T(u_1)+T(v_2)\equiv T(u_2)+T(v_1) \pmod{q'},$$

which contradicts the properties of T.

Given a clique C' in H, consider the origins of the edges connecting the vertices of C' between themselves. We claim that these origin edges are related to *consistent* vertex origins in the graph G in the following sense: there exists a mapping f, from every vertex $v'\in C'$, to a vertex $v\in G$, such that the origin of the edge (v_1',v_2'), for $v_1',v_2'\in C'$, is the edge $(f(v_1'),f(v_2'))$ in G. Once this is proven, it follows immediately that the consistent vertex origins of C' form a clique in G, giving us the desired result.

Assume, by way of contradiction, that the origins of the edges connecting vertices in C' are not consistent. Then there exists a triangle $\{u',v',w'\}$ such that the origin of the edge (u',v') is (u_1,v) and the origin of the edge (u',w') is (u_2,w), where $u_1\neq u_2$. We show that such a triangle cannot exist in H.

Let (v_1,w_1) be the origin in G of the edge (v',w'). Observe that:

$$\begin{aligned} T(u_1)-T(v)&\equiv \mathrm{col}(u')-\mathrm{col}(v') \pmod{q'}\\ T(v_1)-T(w_1)&\equiv \mathrm{col}(v')-\mathrm{col}(w') \pmod{q'}\\ T(w)-T(u_2)&\equiv \mathrm{col}(w')-\mathrm{col}(u') \pmod{q}' \end{aligned}$$

where $\mathrm{col}(x)$, for a vertex x of H, denotes the column of x in H. Combining the above equivalence relationships, we get

$$T(u_1)-T(v)+T(v_1)-T(w_1)+T(w)-T(u_2)\equiv 0 \pmod{q'},$$

and therefore,

$$T(u_1)+T(v_1)+T(w)\equiv T(v)+T(w_1)+T(u_2) \pmod{q'}.$$

By the properties of T, and the fact that the origins of the edges connecting vertices u', v' and w' are in three distinct parts of G, we can conclude that $u_1=u_2$ (as well as that $v_1=v$ and $w_1=w$).

It follows that $\omega(H)\leq\omega(G)$. ∎

3. Hardness result for large constant values

The preceding construction results in graphs H whose chromatic number is $q' = n^{O(1)}$ where n is the number of vertices in the starting graph G. Our goal now is to adapt the argument in the previous section so that q' depends only on q, the maximum number of vertices in any row of G. Starting with an r-partite graph G such that G has $q = O(1)$ vertices in each row, we will show that for any constant h there exists a constant k_h such that it is NP-hard to determine whether $\chi(G) \le k_h$ or $\chi(G) \ge h \cdot k_h$.

3.1. Construction of H

The vertex set of H is once again constructed via an image function. Unlike the earlier construction where the image function T was defined over all the vertices in G, the domain of T is now only the vertices within a partite set of G. In other words, the domain of T is *restricted to* $[q]$. As a result, several vertices may now map to the same column in H and thus an edge in H may not have a unique origin. The analysis of the preceding section can not thus be carried over directly. In fact, to deal with the situation of multiple origins, we need to build some additional structure in our graph H.

We will use the mapping T to transform G into a $(k \cdot r)$-partite graph H for some integer $1 < k < q$, such that if $\omega(G) = r$ then $\bar{\chi}(H) = q'$, and otherwise $\bar{\chi}(H) > \frac{q' \cdot q^{\epsilon}}{2}$. This transformation is better described through an intermediate graph G' where the graph G' is also a $(k \cdot r)$-partite graph.

The intermediate graph G'. For every row of G, we include a block of k rows in G' such that the j^{th} row in the block corresponding to the row i of G is simply the i^{th} row shifted by j columns to the right in a wraparound manner. Thus each vertex of G has k copies in the graph G'. For every edge (u,v) in G, we insert an edge between every copy of vertex u and every copy of vertex v in G'. While doing so, we assume that each vertex of G is connected to itself. Thus all the k copies of any vertex form a clique in G'. It is easy to see that $\omega(G') = k \cdot \omega(G)$.

Transforming G' to H. We now transform G' to a graph H in a manner somewhat similar to the one described in Section 2. We apply the mapping T to the vertices in each row of G', however, the domain of the mapping is now restricted to $[q]$. The edge set of H is constructed via direct and rotated edges as before. For every edge (u,v) in G', we have an edge $(T(u),T(v))$ in H. We extend this edge set by including all the rotations of these edges. Let v be the k^{th} vertex in the j^{th} row of G, and let x be the vertex which is the copy of v in the $i^{\underline{\text{th}}}$ row of the $j^{\underline{\text{th}}}$ block of G', then the *image* of x is the vertex of H which is in the $i^{\underline{\text{th}}}$ row of the $j^{\underline{\text{th}}}$ block and column $T(k+i \bmod q)$ of H.

When the edge origins are not unique. As noted earlier, a consequence of the fact that the domain of T is now restricted to $[q]$ is that an edge (u',v') in H may now have multiple origins in G'. However, the following proposition shows that this can happen only when u' and v' are in the same column in H.

Proposition 1. *Every edge (u',v') in H such that $col(u') \neq col(v')$ has a unique origin edge in G'.*

Proof. Consider an edge (u',v') in H which has at least two distinct origin edges, say (x_1,y_1) and (x_2,y_2). Then we must have

$$T(x_1) - T(y_1) \equiv T(x_2) - T(y_2) \pmod{q'}.$$

By the special property of T, we can conclude $\{x_1,y_2\} = \{x_2,y_1\}$. But since (x_1,y_1) and (x_2,y_2) are distinct edges, either $x_1 \neq x_2$ or $y_1 \neq y_2$. So it must be the case that $x_1 = y_1$ and $x_2 = y_2$. This immediately implies $col(u') = T(x_1) = T(y_1) = col(v')$. ∎

3.2. Relating $\omega(G)$ and $\omega(H)$

We will now relate the clique numbers of the graphs G and H. Unlike the previous section where we obtained matching upper and lower bounds on $\omega(H)$ in terms of $\omega(G)$, we now characterize a range of values where $\omega(H)$ may lie.

Lemma 3.1. $k \cdot \omega(G) \leq \omega(H) \leq k \cdot \omega(G) + r$.

Proof. Since the image of any clique C in G' forms a clique in H and $\omega(G') = k\omega(G)$, it is easy to see that $\omega(H) \geq k \cdot \omega(G)$. We now show that $\omega(H) \leq k \cdot \omega(G) + r$.

To see this, consider any clique C' in H. Ignore any blocks in H where the clique C' contains at most one representative. Thus we restrict ourselves to a subset $C'' \subseteq C'$ such that in every block of H, C'' either has zero or at least two representatives. Clearly, $|C'| \leq |C''| + r$ because C' can have precisely one representative in at most r blocks. Let b denote the number of blocks in H such that C'' contains at least two representatives in them. We will show that $\omega(G) \geq b$.

Assigning vertex origins. Consider now an edge connecting two vertices $u', v' \in C''$ such that u' and v' are in the same block of H. We claim that it must be the case that these two vertices are in different columns and thus the edge (u', v') has a unique origin in G'. This follows rather easily from the observation that for every edge (u, v) in G' such that u and v are in the same block (and thus they are the copies of the same vertex x in G), $T(u) \neq T(v)$. We further use this observation to define a labeling L for each edge $e = (u', v')$ where u', v' are in the same block in H. We define $L(e) = x$ if the edge (u, v) is the origin of the edge (u', v') and u', v' correspond to two different copies of the vertex x in G.

Vertex origins form a clique. We now show that for any edge $e_1 = (u'_1, w'_1)$ in a block B_1 of H, and an edge $e_2 = (u'_2, w'_2)$ in a block B_2 of H (B_1 and B_2 may be the same blocks), such that u'_1, u'_2, w'_1, w'_2 form a 4-clique in H, it must be the case that $L(e_1)$ is connected to $L(e_2)$ in G. Therefore, the labels of the edges connecting vertices (within the same block) in C'' form a clique in G, which implies our claim that $\omega(G) \geq b$.

Since u'_1 and w'_1 are in different columns, as well as u'_2 and w'_2, we can assume, without loss of generality, that u'_1 and u'_2 are in different columns (as well as w'_1 and w'_2). Hence, the origin in G' of the edge (u'_1, u'_2) in H is uniquely defined, say (z_1, y_2); we show that z_1 is a copy of $L(e_1)$ and y_2 is a copy of $L(e_2)$ in G', which are connected in G', hence $L(e_1)$ is connected to $L(e_2)$ in G.

Let (x_1, y_1) be the origin in G' of the edge (u'_1, w'_1), we show that it must be the case that $x_1 = z_1$. A similar argument shows that the origin of e_2 is consistent with the origin of (u'_1, u'_2).

We break the argument into two cases:

- w'_1 and u'_2 are in different columns. The same argument as the triangle argument from the last section applies.
- w'_1 and u'_2 are in the same column. It must be the case that

$$T(x_1) - T(y_1) \equiv T(z_1) - T(y_2) \pmod{q'}.$$

  Using the special property of T, we know that $\{x_1, y_2\} = \{z_1, y_1\}$. Since $T(x_1) \neq T(y_2)$, it must be the case that $x_1 = z_1$.

∎

3.3. Putting together

This characterization immediately yields our goal concerning the chromatic number of $\bar{H}$ since, if $\omega(G) = r$, then H contains a clique C' of size $k \cdot r$ which along with its $q'-1$ rotational images covers all the vertices in H and thus $\bar{\chi}(H)=q'$. Otherwise, $\omega(G)<\frac{r}{q^\epsilon}$ and therefore,

$$\omega(H) < \frac{k \cdot r}{q^\epsilon} + r \ .$$

By choosing $k = \lceil q^\epsilon \rceil$, we get $\omega(H) < 2\frac{k \cdot r}{q^\epsilon}$. Now simply dividing the total number of vertices in H by the size of the largest clique in H, we get $\bar{\chi}(H) > \frac{q' \cdot q^\epsilon}{2}$.

4. Hardness result for fixed constant values

The hardness result of the preceding section tells us that chromatic number is hard to approximate to within any constant when the optimal value itself is allowed to be a sufficiently large constant. But what happens when the graphs is k-colorable for some small integer $k \geq 3$? Specifically, we now focus on the hardness of approximating the chromatic number on 3-colorable graphs. We will show the following theorem:

Theorem 1. **(Main)** *It is NP-hard to color a 3-colorable graph with 4 colors.*

A straightforward corollary of the above theorem is as follows.

Corollary 1. *For any $k \geq 3$, it is NP-hard to color a k-chromatic graph with at most $k+2\lfloor\frac{k}{3}\rfloor-1$ colors.*

To establish our main theorem, we start with an r-partite graph $G=(V,E)$ that either has an r-clique or every clique of G contains less than $\frac{r}{2}$ vertices. It is NP-hard to distinguish between these two cases [9, 2]. Given such a graph G, we construct a graph H such that $\bar{\chi}(H)=3$ when $\omega(G)=r$ and $\bar{\chi}(H)\geq 5$ when $\omega(G)<\frac{r}{2}$. Clearly, Theorem 1 follows from the construction of such a graph H.

For the purpose of this reduction, we assume that every vertex of G is connected to itself, and that every vertex is connected to at least one vertex in each row of G. It is easily verified that any given G can always be transformed into a graph G' such that G' satisfies both these assumptions, and $\omega(G')=r$ if and only if $\omega(G)=r$ and $\omega(G')<\frac{r}{2}$ otherwise.

4.1. The structure and properties of H

The graph H is a multi-partite graph with exactly 3 vertices in each partite set (arranged in 3 columns, denoted 0, 1 and 2). Each row of graph G maps to a block of $O(q)$ rows in H where q is the maximum number of vertices in any row of G. Thus the total number of rows as well as the total number of vertices in H is $O(rq)$. As in our previous constructions, the graph H is symmetric under rotations. H satisfies the following two properties:

1. if $\omega(G)=r$ then $\bar{\chi}(H)=3$, and
2. if $\bar{\chi}(H)\leq 4$ then $\omega(G)\geq r/2$.

As indicated above, this would immediately imply Theorem 1. In what follows, we first decribe how H is constructed and then establish each of the above properties.

4.2. Construction of H

The vertex set. A row in the $i^{\underline{\text{th}}}$ block of H is associated with an ordered 3-way partition of the set of vertices of the $i^{\underline{\text{th}}}$ row of G. Each vertex in a row in the $i^{\underline{\text{th}}}$ block of H can be labeled by a subset of the vertices of the $i^{\underline{\text{th}}}$ row of G, which we refer to as the *label set* of that vertex. The $i^{\underline{\text{th}}}$ block of H consists of $5q_i-7$ rows where q_i is the number of vertices in the $i^{\underline{\text{th}}}$ row of G.

The edge set. The edges of H are determined by the label sets of the vertices. A vertex labeled by a set X is connected to a vertex labeled by a set Y if there exists $u\in X$ and $v\in Y$ such that u and v are connected in G. Note that due to our assumption that each vertex of G is connected to itself, two vertices in the same block of H whose label sets have a non-empty intersection are connected in H. The edges inserted in this manner are referred to as *direct* edges.

In addition, the edge set of H is extended to include all the *rotations* of the above edges as follows: If H contains a direct edge connecting the $i^{\underline{\text{th}}}$ vertex in row k to the $j^{\underline{\text{th}}}$ vertex in row l, then we also add to H edges connecting the $(i+1 \bmod 3)^{\underline{\text{th}}}$ vertex in row k to the $(j+1 \bmod 3)^{\underline{\text{th}}}$ vertex in row l, and similarly between the $(i+2 \bmod 3)^{\underline{\text{th}}}$ and $(j+2 \bmod 3)^{\underline{\text{th}}}$ vertices in these rows.

For conciseness, a row whose vertices 0, 1 and 2 have labels X, Y and Z respectively, is said to have a *row label* of the form

$$X \qquad Y \qquad Z\ ,$$

and we often abbreviate it as $\langle X,Y,Z\rangle$.

In order to fully describe H, it now only remains to describe the ordered partitions associated with each of its rows.

The ordered partitions. Let $Q=\{v_1,\ldots,v_q\}$ denote the set of vertices in the $i^{\underline{\text{th}}}$ row of G, and for $1\leq j<q$, let $X_j=\{v_{j+1},\ldots,v_q\}$, $Y_j=\{v_1,\ldots,v_j\}$ and $Z_j=X_j\cup Y_{j-1}=Q-\{v_j\}$). We now describe how the $i^{\underline{\text{th}}}$ block of rows in H is constructed.

- The first row is the trivial partition, corresponding to having all vertices in one set:

$$Q \qquad \emptyset \qquad \emptyset\,.$$

- For $1\leq j<q$, H contains the following pair of rows:

$$\begin{array}{lll} X_j & Y_j & \emptyset \\ Y_j & X_j & \emptyset\,. \end{array}$$

- In addition, for $1<j<q$, H contains a pair of rows whose labeling corresponds to a partition that singles out the unique vertex not included in Z_j :

$$\begin{array}{lll} Z_j & \{v_j\} & \emptyset \\ \{v_j\} & Z_j & \emptyset\,. \end{array}$$

- Finally, for $1<j<q$, H contains a single row with the label

$$X_j \qquad Y_{j-1} \qquad \{v_j\}\,.$$

4.3. A clique of size r in G implies a 3-coloring of $\bar{H}$

The following lemma directly follows from our construction of the graph H.

Lemma 4.1. *If $\omega(G)=r$ then $\bar{\chi}(H)=3$.*

Proof. Let $C=\{v_1,\ldots,v_r\}$ be the set of vertices forming a clique of size r in G. Clearly, each row of H contains one vertex whose labeling contains some $v_i\in C$; let C_0 be the set of those vertices in H. By the construction of the edge set of H, C_0 constitutes a clique with a representative in every row of H. Since the graph H is symmetric under rotation, C_0 has two rotational images, C_1 and C_2, that together cover all vertices in all rows in H. ∎

4.4. A 4-coloring of $\bar{H}$ implies a large clique in G

Our objective now is to show that $\bar{\chi}(H) \le 4$ implies $\omega(G) = r$. We will show that if $\bar{\chi}(H) \le 4$ then $\omega(G) \ge \frac{r}{2}$, which, by the constraint on the values taken by $\omega(G)$, implies that $\omega(G)$ must be r.[1]

We now give an overview of the main ingredients of our proof. Our starting point is a 4-clique cover of H. We will show that any 4-clique cover of H can be used to identify a *representative* vertex in each row of G such that these representative vertices induce a union of two clique graphs in G. Clearly, one of the two cliques must be $r/2$ in size, giving us the desired result. Broadly speaking, there are four main steps in our proof:

1. The first step is to modify any given 4-clique cover of H into one satisfying a certain property. Once this is done, we identify three cliques in the cover as the *critical* cliques and the fourth one is refered to as the *non-critical* clique (Section 4.4.1).
2. Next we introduce a notion of *voting* and *election* of vertices of H by the critical cliques. In particular, we show that under our notion of election, in every pair of rows of the form $\{\langle L_0, L_1, \emptyset\rangle, \langle L_1, L_0, \emptyset\rangle\}$, the critical cliques elect either L_0 or L_1 (Section 4.4.2). This notion of election plays a central role in identifying the representative vertices in the rows of G.
3. The third step is to show that in every block of rows in H, there exists a pair of the above form such that the set elected is a singleton (Section 4.4.3). These singletons in H have a one-to-one correspondence to the representative vertices in G described above.
4. Finally, let $S_{CLQ} = \{v_1, \ldots, v_r\}$ be set of elected singletons as indentified in the preceding step; each v_i correponds to a vertex from a distinct row in G. We now show that S_{CLQ} can be partitioned into two sets S_1 and S_2 such that each set induces a clique in G. This suffices to conclude that $\omega(G) \ge \frac{r}{2}$ (Section 4.4.4).

4.4.1. The critical cliques

We start by showing how to transform the given 4-clique cover to one that satisfies a simple property. Consider the $i^{\underline{\text{th}}}$ block of rows in H. There are three cliques that contain a vertex in the row whose label is of the form $\langle Q, \emptyset, \emptyset\rangle$, where Q denotes the set of vertices in the $i^{\underline{\text{th}}}$ row of G. To each of these cliques we assign a *shift*, which is either 0,1 or 2, according to whether the clique contains the first, second or the third vertex respectively.

[1] As an aside, it may be noted that by Lemma 4.1, this in turn implies that $\bar{\chi}(H) = 3$.

Observe that the vertices in two rows (in different blocks) with row labels of the above form, are connected only if they appear in the same column. Hence a clique is assigned in this manner at most one shift value over all the blocks in H. However, one of the shift values may be assigned to two cliques, say C and C'. If this were the case, we replace all occurrences of one of the two cliques, say C', with the clique C in each row of the form $\langle Q, \emptyset, \emptyset \rangle$. It follows easily from the structure of H and the assumption that any vertex in G is connected to at least one vertex in every row of G that the modified C still induces a clique in H.

So we can assume from now on that all the rows of the above form are covered by a fixed set of three cliques in the 4-clique cover of H. Let C_0, C_1 and C_2 denote these three cliques; C_i is the clique that is assigned a shift value of i in every block of H. We refer to these three cliques as the *critical* cliques while the remaining clique (which may be empty), denoted by C_N, is referred to as the *non-critical* clique.

4.4.2. The voting and election scheme

The voting scheme. In any row of H, whose label is of the form $\langle L_0, L_1, L_2 \rangle$, we say that the critical clique C_s *votes for* L_i if C_s contains the vertex with label L_j where $i \equiv (j-s) \bmod 3$. That is, C_0 votes for the set that labels the vertex it contains, while C_1 and C_2 vote for the set labeling the vertex to the immediate left and right (in a wraparound manner) respectively, of the vertex that they contain in the above row. Thus in a row whose label has the form $\langle Q, \emptyset, \emptyset \rangle$, each of C_0, C_1 and C_2 vote for Q. Roughly speaking, this voting procedure is helping us identify one of the critical cliques in the cover as a direct image of some clique in G (the clique C_0), and the remaining two critical cliques as shifted images of it.

Let us now consider a pair of rows in the same block of H which have the form

$$\begin{array}{ccc} L_0 & L_1 & \emptyset \\ L_1 & L_0 & \emptyset \end{array}$$

and consider the votes that may be casted by the critical cliques in that pair. The following propositions summarize some useful observations.

Proposition 2. *In a row whose row label is of the form $\langle L_0, L_1, \emptyset \rangle$, no critical clique can vote for the empty set.*

Proof. By definition, the critical clique C_i appears in the column i of the row whose label is of the form $\langle Q, \emptyset, \emptyset \rangle$. By our construction of the edge set of H, C_i can only appear in the columns i and $(i+1) \bmod 3$ of the given row.

So it can only vote for either the entry in column 0 or the entry in column 1 of the given row. The proposition follows. ∎

Proposition 3. *In a pair of rows of H of the form $\{\langle L_0, L_1, \emptyset\rangle, \langle L_1, L_0, \emptyset\rangle\}$, the following two properties are always satisfied :*

(a) a critical clique which appears in both rows either votes for L_0 in both rows or votes for L_1 in both rows, and

(b) two critical cliques which appear in both rows, either together vote for L_0 or together vote for L_1.

Proof. By Proposition 2, we know that a critical clique only votes for L_0 or L_1 in either of the two rows. Since there are no vertical edges between the two rows, it cannot vote for L_0 in one row and L_1 in the other row of the pair. This gives us property (a).

To see property (b), consider a critical clique, say C_i, which appears in both rows. By property (a), it either votes for L_0 or L_1 in both rows. Without loss of generality, assume it votes for L_0 in both rows. Then the vertex which corresponds to the critical clique $C_{(i+1) \bmod 3}$ voting for L_1 in the second row of the pair, is taken by C_i. Similarly, the vertex which corresponds to $C_{(i+2) \bmod 3}$ voting for L_1 in the first row of the pair, is also taken by C_i. Thus if either $C_{(i+1) \bmod 3}$ or $C_{(i+2) \bmod 3}$ also appears in both rows of the pair, it must also vote for L_0. ∎

The election scheme. We say that the label L_i is *elected* in the above pair of rows if the majority (i.e. two) of the critical cliques vote for L_i, $i = \{0, 1\}$. The following is a straightforward consequence of the preceding two propositions.

Lemma 4.2. *In a pair of rows of H of the form $\{\langle L_0, L_1, \emptyset\rangle, \langle L_1, L_0, \emptyset\rangle\}$, majority is always defined.*

Proof. If all three critical cliques appear in this pair of rows, majority is clearly defined by Proposition 3(a). On the other hand, if only two of the critical cliques appear in this pair of rows (i.e., the pair of rows is covered by only three cliques), then each of these two critical cliques appears in both the rows and therefore by Proposition 3(b), both of them must vote for either L_0 or L_1. ∎

For example, in Figure 1., C_0 votes for L_0 while C_1 and C_2 vote for L_1. Hence L_1 is elected in this pair of rows.

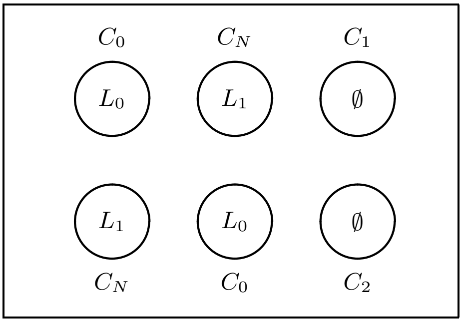

Fig. 1. An example of critical cliques electing L_1

4.4.3. Every block elects a singleton

We have established so far that in every pair of rows of the form $\{\langle L_0, L_1, \emptyset\rangle, \langle L_1, L_0, \emptyset\rangle\}$, either L_0 or L_1 must get majority of the votes. Our next goal is to show that among all such pairs of rows in a block, there exists one such that the set elected is a singleton set.

Lemma 4.3. *In every block B_i of rows in H, there exists a pair of rows corresponding to the partition of the form $\langle Q - \{v_j\}, \{v_j\}, \emptyset\rangle$ where $Q = \{v_1, \ldots, v_q\}$ is the set of vertices in the $i^{\underline{th}}$ row of G, such that $\{v_j\}$ is elected by the majority of the critical cliques in this pair.*

Proof: Recall that $X_j = \{v_{j+1}, \ldots, v_q\}$, $Y_j = \{v_1, \ldots, v_j\}$ and $Z_j = X_j \cup Y_{j-1}$, where $1 \leq j < q$. Now consider the sequence of pairs of rows with the row labels of the form:

$$\begin{array}{ccc} X_j & Y_j & \emptyset \\ Y_j & X_j & \emptyset \ . \end{array}$$

If either $Y_1 = \{v_1\}$ is elected in the first pair ($j = 1$), or $X_{q-1} = \{v_q\}$ is elected in the last pair ($j = q-1$), we are done. Otherwise, there must be a *switch point* $k \in [2..q-1]$ such that $X_{k-1} = \{v_k, \ldots, v_q\}$ is elected in the $(k-1)^{\underline{th}}$ pair and $Y_k = \{v_1, \ldots, v_k\}$ is elected in the $k^{\underline{th}}$ pair. We show that then it must be the case that $\{v_k\}$ is elected in the pair of rows with row label of the form:

$$\begin{array}{ccc} Z_k & \{v_k\} & \emptyset \\ \{v_k\} & Z_k & \emptyset \ . \end{array}$$

Suppose by way of contradiction, the vertex with label Z_k is elected in the above pair of rows. We then show that it is impossible to cover all three vertices in the row R below:

$$R: \qquad X_k \qquad Y_{k-1} \qquad \{v_k\}\ .$$

Specifically, we will argue that at most one critical clique can be present in R. Since a non-critical clique can cover at most one vertex in any row, this will contradict the fact that we started with a 4-clique cover of H.

Consider the three pairs of rows which corresponds to the following three partitions:

- the partition just before the switch: $\langle X_{k-1}, Y_{k-1}, \emptyset\rangle$,
- the partition just after the switch: $\langle X_k, Y_k, \emptyset\rangle$, and
- the partition that singles out the vertex that causes the switch: $\langle Z_k, \{v_k\}, \emptyset\rangle$.

A critical clique C_i that contains a vertex in row R votes for one of $Y_{k-1}, \{v_k\}$ or X_k. Let W denote this set. Clearly, one of the above three partitions has the form $\langle W, Q-W, \emptyset\rangle$. Let P be the pair of rows corresponding to this partition. By our assumption, the label $Q-W$ is elected in P. The critical clique C_i cannot vote for $Q-W$ in P since C_i votes for W in R and the edges connecting the vote for W in R to the votes for $Q-W$ in P do not exist in H (there are three direct edges connecting a row in P to R, however, none is a rotation of the edge that if existed in H were to connect W in R to $Q-W$ in P). Therefore, it must be the case that the two remaining critical cliques, namely $C_{(i+1) \bmod 3}$ and $C_{(i+2) \bmod 3}$, vote for $Q-W$ in P (otherwise $Q-W$ would not have been elected in P).

Let $j_1 \neq j_2 \in \{0,1\}$ be the columns, of the vertices in the first and second rows of P respectively, that are labeled by $Q-W$. Consider the vertex in the first row and the $(j_1+i) \bmod 3$ column of P, and the vertex in the second row and $(j_2+i) \bmod 3$ column of P (if the critical clique C_i were to vote for $Q-W$, it would cover at least one of these two vertices). These two vertices cannot be contained in the critical cliques $C_{(i+1) \bmod 3}$ or $C_{(i+2) \bmod 3}$ as this would mean these cliques do not vote for $Q-W$, which then could not have been elected. Therefore, these two vertices must be contained in the remaining non-critical clique C_N (*see Figure 2.*). Thus C_N contains vertices in both column i and column $(i+1) \bmod 3$ (this follows because $j_1 \neq j_2 \in \{0,1\}$).

Now if another critical clique, say $C_{i'}$ ($i' \neq i$), appears in row R voting for some set W', there exists a different pair of rows, say P', such that it corresponds to a partition of the form $\langle W', Q-W', \emptyset\rangle$ and by our assumption, the label $Q-W'$ is elected in P'. By applying the same argument as before, we can conclude that C_N must contain a vertex in column i' and a vertex in column $(i'+1) \bmod 3$ of this pair. This means C_N contains vertices in all three columns in the pairs P and P'. However, taking into account the edges connecting pairs P and P', this contradicts the following simple proposition:

Proposition 4. *In any block of H, consider a pair of rows, say P_1, with labels of the form:*

$$\begin{matrix} L_0+S & L_1 & \emptyset \\ L_1 & L_0+S & \emptyset \end{matrix} .$$

Suppose a clique C in H contains a vertex i, for $i \in \{0,1,2\}$, in the first row and vertex $(i+1) \bmod 3$ in the second row (i.e., some shift of a choice that corresponds to label L_0+S). Now consider the pair of rows, say P_2, with row labels of the form :

$$\begin{matrix} L_0 & L_1+S & \emptyset \\ L_1+S & L_0 & \emptyset \end{matrix} .$$

Then C cannot contain in this pair of rows a vertex in the $(i+2) \bmod 3$ column.

Proof. Simply observe that the vertex i in the first row of P_1, is not connected to the vertex $(i+2) \bmod 3$ in the first row of P_2 and similarly, the vertex $(i+1) \bmod 3$ in the second row of P_1, is not connected to the vertex $(i+2) \bmod 3$ in the second row of P_2. ∎

Consequently, only one critical clique can appear in R and thus all the vertices of R could not have been covered by this clique cover. This is a contradiction. We therefore conclude that the vertex with label $\{v_k\}$ is elected in the pair of rows corresponding to the partition $\langle Z_k, \{v_k\}, \phi\rangle$. ∎

4.4.4. The singletons form at most two cliques in G

Our final goal now is to show that the set of vertices formed by the singleton sets elected in each block (we associate exactly one singleton set with each block), is indeed the set S_{CLQ} we described earlier. This clearly will complete the proof of our main theorem.

Lemma 4.4. *Let $S_{CLQ} = \{v_1, \ldots, v_r\}$ be a set of vertices of G, such that v_i belongs to the $i^{\underline{th}}$ row of G and $\{v_i\}$ is elected in the $i^{\underline{th}}$ block of H. Then S_{CLQ} is a union of two cliques in G.*

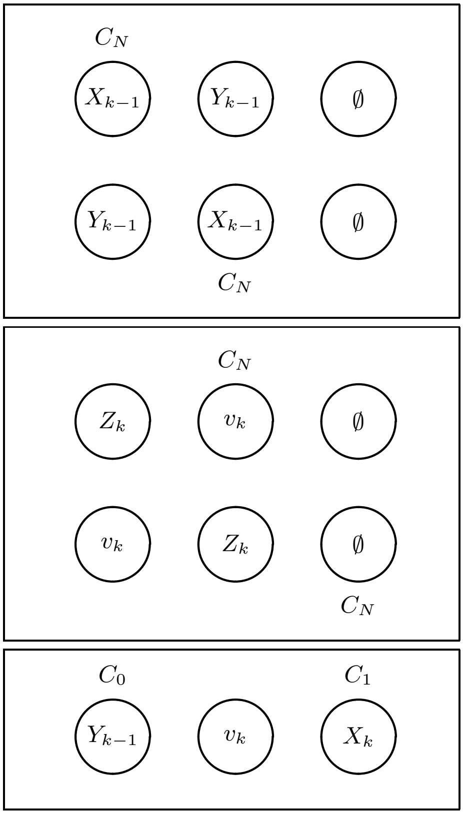

Fig. 2. Critical cliques C_0 and C_1 vote for Y_{k-1} and v_k, respectively, in row R. This forces the non-critical clique C_N to cover 4 vertices, as shown, that do not induce a clique in H.

Proof. Let Q_i denote the set of vertices in the $i^{\underline{\text{th}}}$ row of G, where $i \in \{1,\dots,r\}$ and let P_i denote a pair of rows of the form

$$\begin{array}{lll} Q_i - \{v_i\} & \{v_i\} & \emptyset \\ \{v_i\} & Q_i - \{v_i\} & \emptyset, \end{array}$$

where $v_i \in Q_i$. We say that C_N *holds* $\{v_i\}$ *with shift* $j \in \{0,1,2\}$ *in* P_i if C_N contains both the vertex in the $(j+1 \bmod 3)^{\underline{\text{th}}}$ column of the first row, and the vertex in the $j^{\underline{\text{th}}}$ column in the second row.

We define $S_1 \subseteq S_{CLQ}$ to be the set of vertices v_i such that, in the pair of rows P_i, C_N holds $\{v_i\}$ with some shift $j \in \{0,1,2\}$ and let $S_2 = S_{CLQ} - S_1$. We claim that S_1 and S_2 each form a clique in G.

S_2 induces a clique in G. Let us first look at the easier case which is that of S_2. In each P_i, such that $v_i \in S_2$, C_N does not hold $\{v_i\}$ with any shift j. Therefore, every critical clique that contains a vertex in P_i votes for $\{v_i\}$ and hence it must be the case that one of the critical cliques votes for $\{v_i\}$ in P_i and $\{v_{i'}\}$ in $P_{i'}$, and contains 3 vertices in these four rows. This is not possible unless v_i is connected in G to $v_{i'}$.

S_1 induces a clique in G. We now focus on the set S_1. Consider any pair of vertices $v_i, v_{i'} \in S_1$. By our construction of S_1, there exists $j, j' \in \{0,1,2\}$ such that C_N holds $\{v_i\}$ with shift j in P_i and $\{v_{i'}\}$ with shift j' in $P_{i'}$. We argue that unless v_i and $v_{i'}$ are connected in G, both v_i and $v_{i'}$ could not have been elected.

For clarity of exposition, let us represent the pair of rows P_i and $P_{i'}$ such that their first row is shifted one column to the left (note that we are not changing the edge set of H).

$$
\begin{array}{llll}
P_i : & \{v_i\} & \emptyset & Q_i - \{v_i\} \\
 & \{v_i\} & Q_i - \{v_i\} & \emptyset
\end{array}
$$

$$
\begin{array}{llll}
P_{i'} : & \{v_{i'}\} & \emptyset & Q_{i'} - \{v_{i'}\} \\
 & \{v_{i'}\} & Q_{i'} - \{v_{i'}\} & \emptyset
\end{array}
$$

The following proposition can be easily verified now.

Proposition 5. *Unless v_i is connected to $v_{i'}$ in G, H has no vertical edges connecting a vertex in the first (second) row of P_i to a vertex in the second (first) row of $P_{i'}$.* ∎

By our assumption, C_N holds $\{v_i\}$ with some shift j in P_i and holds $\{v_{i'}\}$ with some shift j' in $P_{i'}$. Note that this implies that C_N contains both vertices in column j in P_i as well as both vertices in column j' in $P_{i'}$. If $j = j'$, there exists a column such that C_N contains a vertex in this column in the first (second) row of P_i as well as in the second row of $P_{i'}$. In this case, we are done by Proposition 5. So we assume from now on that $j \neq j'$. Let $\alpha \in \{0,1,2\}$ be such that $\alpha \neq j$ and $\alpha \neq j'$. Observe that C_N does not contain any vertices in column α.

Since v_i and $v_{i'}$ are elected, it must be the case that (a) the critical cliques with shifts j' and α constitute the majority of votes in P_i, and (b) the critical cliques with shifts j and α constitute the majority of votes in $P_{i'}$. As C_α votes for $\{v_i\}$ in P_i and $\{v_{i'}\}$ in $P_{i'}$, it follows that C_α contains vertices only in column α.

Assume by way of contradiction that v_i is not connected to $v_{i'}$. Then by Proposition 5, C_α either contains vertices only from the first row of P_i and the first row of $P_{i'}$, or the second row of P_i and second row of $P_{i'}$. However, using the facts (a) and (b) above, a straightforward case analysis shows that C_j and $C_{j'}$ can cover the vertices in column α of P_i and $P_{i'}$ in only one of the following two ways:

– clique C_j contains the vertex in the first row in column α of P_i and the clique C'_j contains the vertex in the second row in column α of $P_{i'}$, or
– clique C_j contains the vertex in the second row in column α of P_i and the clique C'_j contains the vertex in the first row in column α of $P_{i'}$.

In either case, we are left with one vertex in column α of the rows in P_i and $P_{i'}$ that is not covered by any of the 4 cliques. This contradicts the fact that these cliques constitute a clique cover of H. Hence v_i must be connected to $v_{i'}$ in G. ∎

Concluding remarks. The obvious open problem is to tighten the gap between the known positive and negative approximability results for the 3-coloring problem. Another interesting problem is to show the hardness of coloring a k-colorable graph with $f(k)$ colors where $f(k)$ grows faster than any constant power of k.

Acknowledgements. The authors would like to thank Michal Parnas, Steven Phillips and Carsten Lund for many useful discussions. A very special thanks to Oded Goldreich for his insightful comments on an earlier version of this paper.

References

[1] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy: Proof verification and hardness of approximation problems, *Journal of the ACM*, **45** (3) (1998), 501–555. Preliminary version in *Proceedings of the 33rd IEEE Symposium on Foundation of Computer Science*, 1992.

[2] S. Arora and S. Safra: Probabilistic checking of proofs: A new characterization of NP, *Journal of the ACM*, **45** (1) (1998), 70–122. Preliminary version in *Proceedings of the 33rd IEEE Symposium on Foundation of Computer Science*, 1992.

[3] A. Blum: An $O(n^{0.4})$ approximation algorithm for 3-coloring, *Proceedings of the 21st ACM Symposium on the Theory of Computing*, (1989), 535–542.

[4] A. Blum: Some tools for approximate 3-coloring, *Proceedings of the 31st IEEE Symposium on Foundation of Computer Science*, (1990), 554–562.

[5] A. Blum: *Algorithms for Approximate Graph Coloring*, Ph.D. thesis, Laboratory for Computer Science, MIT, Cambridge, MA, 1991.

[6] A. Blum: New approximation algorithms for graph coloring, *Journal of the ACM*, **41** (1994), 470–516.

[7] A. Blum and D. R. Karger: An $\tilde{O}(n^{3/14})$-coloring algorithm for 3-colorable graphs, *Information Processing Letters*, **61** (1) (1997), 49–53.

[8] M. Bellare and M. Sudan: Improved non-approximability results, *Proceedings of the 26th ACM Symposium on Theory of Computing*, (1994), 184–193.
[9] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy: Interactive proofs and the hardness of approximating cliques, *Journal of the ACM*, **43** (2) (1996), 268–292. Preliminary version in *Proceedings of the 32nd IEEE Sym. on Foundations of Computer Science*, 1991.
[10] U. Feige and J. Kilian: Zero knowledge and the chromatic number, *J. Comput. System Sci.*, **57** (2) (1998), 187–199. Preliminary version in *Proceedings of the 11th IEEE Conference on Computational Complexity*, 1996.
[11] M. Fürer: Improved hardness results for approximating the chromatic number, *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, (1995), 414–421.
[12] M. R. Garey and D. S. Johnson: The complexity of near-optimal graph coloring, *Journal of the ACM*, **23** (1976), 43–49.
[13] M. R. Garey and D. S. Johnson: *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.
[14] V. Guruswami and S. Khanna: On the hardness of 4-coloring a 3-colorable graph, *in preparation*, 1999.
[15] M. M. Halldórsson: A still better performance guarantee for approximate graph coloring, *Inform. Process. Letters*, **45** (1993), 19–23.
[16] J. Håstad: Clique is hard to approximate within $n^{1-\epsilon}$, *Proceedings of the 36th IEEE Symposium on Foundations of Computer Science*, (1996), 627–636.
[17] Richard M. Karp: Reducibility among combinatorial problems, in: *Complexity of Computer Computations* (eds.: Raymond E. Miller and James W. Thatcher), 85–103, Plenum Press, 1972.
[18] D. R. Karger, R. Motwani and M. Sudan: Approximate graph coloring using semidefinite programming. *Journal of the ACM*, **45** (2) (1998), 246–265.
[19] N. Linial and U. Vazirani: Graph products and chromatic numbers, *Proc. 30th IEEE Sym. on Foundations of Computer Science*, (1989), 124–128.
[20] C. Lund and M. Yannakakis: On the hardness of approximating minimization problems, *Journal of the ACM*, **41** (1994), 960–981. Preliminary version in *Proceedings of the 25th ACM Symposium on Theory of Computing*, 1993.
[21] A. Wigderson: Improving the performance guarantee for approximate graph coloring, *Journal of the ACM*, **30** (4) (1983), 729–735. Preliminary version in *Proceedings of the 14th ACM Symposium on Theory of Computing*, 1982.

Sanjeev Khanna

Department of Computer
and Information Science
University of Pennsylvania
Philadelphia, PA 19104, USA
sanjeev@cis.upenn.edu

Nathan Linial

Institute of Computer Science
Hebrew University of Jerusalem
Jerusalem 91904, Israel
nati@cs.huji.ac.il

Shmuel Safra

Department of Computer Science
Tel Aviv University
Tel Aviv 69978, Israel
safra@math.tau.ac.il