On the Hardness of Approximating the Chromatic Number

Sanjeev Khanna
Stanford University

Nathan Linial
Hebrew University

Shmuel Safra*
IBM Almaden &
Stanford University

Abstract

The paper considers the computational hardness of approximating the chromatic number of a graph.

We first give a simple proof that approximating the chromatic number of a graph to within a constant power (of the value itself) is NP-hard.

We then consider the hardness of coloring a 3-colorable graph with as few as possible colors. We show that determining whether a graph is 3-colorable or any legal coloring of it requires at least 5 colors is NP-hard. Therefore, coloring a 3-colorable graph with 4 colors is NP-hard.

1 Introduction

Given a graph G, a *legal coloring* of G is an assignment of a color to each vertex of G, such that no edge of G is connecting two vertices that are assigned the same color. The chromatic-number of G, denoted $\chi(G)$ is the minimum number of colors necessary for a legal coloring of G.

Given as input a graph G, one may want to come up with a legal coloring of G with the fewest possible colors. A polynomial-time algorithm that achieves coloring with a set of colors whose size is not vastly larger than the optimum ($\chi(G)$) has many practical applications. 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 of the set of tasks such that no set of the partition contains two conflicting tasks. This problem is equivalent to the chromatic number problem: consider the graph whose vertices are all tasks, where two vertices are connected if the corresponding tasks are conflicting. A coloring of this graph is a non-conflicting partition of the set of tasks.

Coloring a graph G with the minimum $\chi(G)$ colors was shown to be NP-hard by Karp [Kar72]; the results there imply that coloring a 3-colorable graph with 3 colors is NP-hard (this implies the same hardness result for k-colorable graph for any $k \geq 3$).

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

On the other hand, Blum [B89, B90], following earlier work by Wigderson [W83] provided a polynomial-time algorithm which, on input a 3-colorable n-vertex graph, finds a legal coloring using at most $O(n^{3/8}\log^{\frac{5}{8}} n)$ colors. In general, for a k-colorable graph, the best known performance guarantee is due to an algorithm by Halldórsson [H93] that colors a graph G with $O(\frac{n(\log\log n)^2}{\log^3 n}) \cdot \chi(G)$ colors.

These results leave a huge gap between the achievable coloring and the apparently intractable. Where is the true value: it may be the case that a 3-colorable graph can be colored in polynomial-time by a constant number of colors; or is it the case that even coloring G with n^ϵ for some ϵ is an NP-hard problem?

In the last section of the paper, we slightly narrow this gap by strengthening Karp's hardness result. In particular, we show that separating between the case where $\chi(G) \leq 3$ and the case where $\chi(G) \geq 5$ is NP-hard. This implies that given a 3-colorable graph G it is NP-hard to color G with 4 colors. An immediate corollary of this 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.

Some results on the hardness of approximate coloring were obtained over the years. Garey and Johnson [GJ76] have used composition operations on graphs to show that if, for every k, there exists a polynomial time procedure to color a k-colorable graph with fewer than $2k-6$ colors then P=NP. However, they do not specify an integer k for which it is NP-hard to color a k-colorable graph with fewer than $2k-6$ colors. Furthermore, their result clearly says nothing about the hardness of coloring a 3-colorable graph. Linial and Vazirani [LV89] have used squares of graphs to investigate the best ratio of approximating the chromatic number in polynomial time. They give evidence that this ratio, for n-vertex graphs is either below $\log^{1+o(1)} n$ or above $n^{\Omega(1)}$.

Until recently, almost all the approximation versions of NP-optimization problems shared this state of affairs, in which there are only relatively simplistic hardness results for some factors of approximation, that were usually quite large compared to the factors for which a polynomial-time approximation algorithm is known.

In a recent development, a new method for proving NP-hardness of approximation problems was suggested

[FGLSS91, AS92]; using a new characterization of NP — as all languages whose membership-witness can be verified probabilistically while reading only a small number of bits — one may obtain an instance of an optimization problem, in which the optimized parameter is either large or small depending on the probability of error of the verification procedure. Therefore, even an approximated value of the optimized parameter determines whether the original input has a valid membership-witness (i.e., it is in the NP language) or not.

Continuing this direction, Lund and Yannakakis [LY92] showed that for certain two constants $0 < c_1 < c_2 < 1$, it is NP-hard to color an n^{c_1}-colorable graph with n^{c_2} colors. The authors go on to prove that for every constant h there exists a constant c_h such that it is NP-hard to color a c_h-colorable graph by $c_h \cdot h$ colors. However, c_h depends on h and is relatively large in comparison; this is again not applicable to small fixed values of the chromatic number.

Organization

In Sections 2 and 3 we present a simple proof to both the theorems of Lund and Yannakakis [LY92] concerning the hardness of approximating the chromatic number. In Section 4 we prove that, given a graph G, it is NP-Hard to distinguish between the case where $\chi(G) \leq 3$ and the case where $\chi(G) \geq 5$, and therefore, it is NP-hard to color a 3-colorable graph using only 4 colors.

2 Hardness of Approximating the Chromatic Number when the Value May be Large

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

Notation

We use $\omega(G)$ to denote the size of the largest clique in a graph G. The chromatic number of G is denoted by $\chi(G)$, while its clique-cover number is denoted by $\bar{\chi}(G)$. We use $\bar{G}$ to denote the complement graph of G. Since any clique-cover of a graph G (a partition of G into cliques) is a legal coloring of the graph $\bar{G}$ (a partition of $\bar{G}$ into independents sets), it follows that $\bar{\chi}(G) = \chi(\bar{G})$.

The Starting Point

We start with an r-partite graph G that results from the reduction of [FGLSS91], so it either has an r-clique (the maximum possible) or every clique of G is of small size.

Pictorially, one can think of the graph G as consisting of r rows, each being an independent set, and having some edges connecting vertices in different rows. Hence a clique of size r in G has exactly one representative in each row. Note that each row of G may have different number of vertices.

Let q be the maximum number of vertices in any row of G, and assume that either $\omega(G) = r$ or $\omega(G) < \frac{r}{q^{\epsilon}}$ for some absolute constant ϵ, $0 < \epsilon < 1$. For any $q = o(r)$, separating between these two cases is NP-hard [FGLSS91, AS92, AMSSS92, ALMSS92, PS92].

Our Approach

Given such a graph G with $q = r^{\delta}$ for some constant δ, $0 < \delta < 1$, we construct a graph H with rq' vertices for some $q' = (rq)^{O(1)}$ 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.

The graph H can be intuitively described as a transformation of G which makes H symmetric under certain rotations, so that an r-clique has several symmetric images which together cover H completely. On the other hand, the transformation also ensures that the size of the largest clique in H is at most that of G.

The graph H, like G, is an r-partite graph, and is constructed so that for any clique C in H there are $q' - 1$ additional cliques in H, that together cover all the vertices in all the rows in which C has a representative. It will be shown below that $\omega(G) = \omega(H)$. Therefore, if $\omega(G) = r$ then $\bar{\chi}(H) = q'$. On the other hand, if $\omega(G) < \frac{r}{q^{\epsilon}}$ then by simply dividing the number of vertices in H by the size of the largest clique in H, we obtain the claimed lower-bound on $\bar{\chi}(H)$.

The Structure of H

The graph H is an r-partite graph; each row of H corresponds to a distinct row of G. A row of H consists of q' vertices (arranged in q' columns, denoted $0, ..., q' - 1$). We consider a mapping (to be specified below) of the vertices in each row of G to the vertices in the corresponding row of H. We call the vertex v' of H, to which a vertex v of G is mapped, the *image vertex* of v.

We now describe how the edge set of H is constructed. For every edge (u, v) in G, we have 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 $(i + m \bmod q')^{\underline{\text{th}}}$ vertex in row k to the $(j + m \bmod q')^{\underline{\text{th}}}$ vertex in row l for every $m \in \{1, .., q' - 1\}$.

In order to fully describe H, it now only remains to describe the image function, which maps each vertex v of G to a unique vertex in H. However, from the above description alone we can conclude the following:

Lemma 1 *If $\omega(G) = r$ then $\bar{\chi}(H) = q'$.*

Proof: The graph H is *symmetric under rotation*; i.e., 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. □

The Image Function

For an integer $k > 0$ we denote by $[k]$ the set $\{0, ..., k-1\}$. As a preliminary for the description of the image function, from G to H, we start off with a simple lemma, showing the existence of an injection T, of some special structure, from any set of size n into the range $[m]$, for some $m = n^{O(1)}$.

Lemma 2 *For every positive integer n, there exists a function $T: [n] \rightarrow [m]$ 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 $T(i_1) + T(i_2) + T(i_3) \bmod m$ as well as its subsums (that is, $T(i_1) + T(i_2) \bmod m$ and $T(i_1)$) are distinct.*

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 $T(0), ..., T(l-1)$ 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 $T(i_1) + T(i_2) \bmod m$ 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. It is clearly the case that the mapping T is an injection. □

The Reduction

Let T be the injection shown to exist in Lemma 2 for the domain of all vertices of G, i.e., $n = rq$, and let $q' = m$, the size of the range of T. Hence,

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

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

We now show that our transformation of graph G into the graph H preserves the clique number, that is, G and H have the same clique number.

Lemma 3 *$\omega(G) = \omega(H)$.*

Note that this immediately yields our goal concerning the chromatic number of $\bar{H}$ since if $\omega(G) = r$ then $\bar{\chi}(H) = q'$ by Lemma 1. Otherwise, $\omega(G) < \frac{r}{q^\epsilon}$ and then simply counting the number of vertices of H and dividing by the size of the largest clique in H (using Lemma 3), we get

$$\bar{\chi}(H) > \frac{rq'}{\frac{r}{q^\epsilon}} = q' \cdot q^\epsilon.$$

Proof (of Lemma 3):

$\omega(H) \geq \omega(G)$: The image of a clique in G forms a clique in H.

$\omega(H) \leq \omega(G)$: We say that an edge (u, v) in G is the *origin* of an edge e in H if e is obtained via the images of u and v or by rotating the edge so obtained.

Using the special property 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

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

and it follows that

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

which contradicts the property 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, which implies our claim.

Assume, by way of contradiction, that the origins of the edges connecting vertices in C' are not consistent; then there would have been 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'), and 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

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

which by the special property of T, and the fact that the origins of the edges connecting vertices u', v' and w' are in three distinct parts of G, implies that $u_1 = u_2$ (as well as that $v_1 = v$ and $w_1 = w$).

Hence we conclude that $\omega(H) \leq \omega(G)$. □

3 Hardness of Approximating the Chromatic Number when the Value is (a Large) Constant

We now sketch how the argument in the previous section can be extended to eliminate the dependency of q' on r and thereby show that for any constant h there exists a constant c such that it is NP-hard to determine whether $\chi(G) \leq c$ or $\chi(G) \geq c \cdot h$.

We start with the graph G as described earlier, however, we now assume that q is a constant. Furthermore, the injection T which was previously defined over all the vertices in G, is now defined only over the vertices in any part of G, that is, 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.

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. For every row of G, we include a block of k rows in G' such that the $j^{\underline{\text{th}}}$ row in the block corresponding to the row i of G is simply the $i^{\underline{\text{th}}}$ 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 $\omega(G') = k \cdot \omega(G)$.

Once the graph G' is constructed, we transform it into the graph H in similar manner as before by applying the mapping T to the vertices in each row of G'. The edge set of H is also constructed in a similar manner. For every edge (u, v) in G', we have an edge (u', v') in H where u' and v' are the images of u and v respectively. We extend this edge set by including all the rotations of these edges. Let v be the $k^{\underline{\text{th}}}$ vertex in the $j^{\underline{\text{th}}}$ 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 $T(k+i \bmod q)$ column of H.

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 claim shows that this can happen only when u' and v' are in the same column in H.

Claim 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 origins edge and let (x_1, y_1) and (x_2, y_2) be any two such distinct origin edges. Thus 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 must have $\{x_1, y_2\} = \{x_2, y_1\}$. 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 $col(u') = T(x_1) = T(y_1) = col(v')$. □

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 4 *$k \cdot \omega(G) \leq \omega(H) \leq k \cdot \omega(G) + r$.*

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 $\bar{\chi}(H) > \frac{q' \cdot q^\epsilon}{2}$.

Proof (of Lemma 4):

$\omega(H) \geq k \cdot \omega(G)$: The image of a clique C in G' forms a clique in H and $\omega(G') = k\omega(G)$.

$\omega(H) \leq k \cdot \omega(G) + r$: 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 $\omega(G) \geq b$.

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.

We now show that for any edge $e_1 = (u'_1, w'_1)$ within a block of H, and an edge $e_2 = (u'_2, w'_2)$ also within a block of H, where 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 $x_1 = z_1$.

□

4 Hardness of Approximating the Chromatic Number for Small Values

In this section we describe our proof that distinguishing between the case that the chromatic number of a graph is at most 3 and the case that it is at least 5 is NP-hard.

The Starting Point

As before, 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. Separating between these two cases is NP-hard [FGLSS91, AS92].

Our Approach

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 $\omega(G) < \frac{r}{2}$. Since the chromatic number of the graph $\bar{H}$ is equal to the clique cover number of H, our result 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 easy to see that given the graph G, it can always be transformed into a graph G' such that G' satisfies both these assumptions and $\omega(G') = r$ iff $\omega(G) = r$ and $\omega(G') < \frac{r}{2}$ otherwise.

4.1 The Structure of H

The graph H is a multi-partite graph; the rows of H (partite sets) are divided into r *blocks*, one for each row (partite set) of G. The $i^{\underline{\text{th}}}$ block of H consists of $O(q_i)$ rows (more precisely, $5q_i - 7$ rows), where q_i is the number of vertices in the $i^{\underline{\text{th}}}$ row of G. Therefore, the number of rows in H is at most $O(rq)$, where q is the maximum number of vertices in any row of G.

Each of the rows in the $i^{\underline{\text{th}}}$ block of H consists of 3 vertices (arranged in 3 columns, denoted 0, 1 and 2) and is associated with a ordered 3-way partition of the set of vertices of the $i^{\underline{\text{th}}}$ row of G (however, not all possible such partitions will be used in constructing the graph H). Thus each vertex in a row in the $i^{\underline{\text{th}}}$ block of H is labeled by a single subset of the vertices of the $i^{\underline{\text{th}}}$ row of G, which we refer to as the *label set* of that vertex. These labels, as we will see next, determine the edge set of H.

A vertex of H, labeled by a set of vertices $X \subset V$, is connected to a vertex labeled by the set $Y \subset V$, if there exists a vertex $u \in X$ and a vertex $v \in Y$ 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:

$$X \qquad Y \qquad Z$$

In order to fully describe H, it now only remains to describe the ordered partitions associated with each of its rows. However, from the above description alone we can conclude the following.

Claim 2 *Let S be a subset of the rows of H, and let C_0 be a clique that has one representative in each row in S. Then the set of rows S is coverable by 3 cliques.*

Proof: The graph H is symmetric under rotation; i.e., let $s(H)$ be the graph that results by rotating all rows of H by 1 to the right (in a wraparound manner), then $s(H) = s(s(H)) = H$. Therefore, C_0 has two rotational images, C_1 and C_2, that together cover all vertices in all rows in S. □

The Ordered Partitions

Let $Q = \{v_1, ..., v_q\}$ denote the set of vertices in the $i^{\underline{\text{th}}}$ row of G, and for each $j \in [1..q-1]$, let $X_j = \{v_{j+1}, ..., v_q\}$, $Y_j = \{v_1, ..., v_j\}$ 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 putting all the vertices in one set:

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

- For each j, $1 \leq j < q$, H contains one row with labeling given by

  $$X_j \qquad Y_j \qquad \emptyset\,,$$

  and another row that corresponds to the following ordering of the same partition :

  $$Y_j \qquad X_j \qquad \emptyset\,.$$

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

  $$Z_j \qquad \{v_j\} \qquad \emptyset\,,$$

  and

  $$\{v_j\} \qquad Z_j \qquad \emptyset\,.$$

- And finally, for each j, $1 < j < q$, H contains a single row with labeling

  $$X_{j+1} \qquad Y_j \qquad \{v_j\}\,.$$

4.2 A Clique of Size r in G Implies a 3-Coloring of $\bar{H}$

The following lemma is a rather immediate consequence of the construction of the edge set of H and claim 2.

Lemma 5 *If $\omega(G) = r$ then $\bar{\chi}(H) = 3$.*

Proof: Let $C = \{v_1, ..., v_r\}$ 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. Using Claim 2, we conclude that H is coverable by 3 cliques. □

4.3 A 4-Coloring of $\bar{H}$ Implies a Large Clique in G

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

Assume H has a 4-clique cover. We use such a cover to identify a set of vertices of G, say $S_{CLQ} = \{v_1, ..., v_r\}$, where each v_i belongs to a distinct row of G such that S_{CLQ} is a union of two cliques S_1 and S_2 in G. The construction of such a set S_{CLQ} is clearly sufficient to conclude that $\omega(G) \geq \frac{r}{2}$.

The Critical Cliques

Consider the $i^{\underline{\text{th}}}$ block of rows in H. There are three cliques that contain a vertex in the row whose row label is of the form

[1] As an aside, it may be noted that by Lemma 5, this in turn implies that $\bar{\chi}(H) = 3$.

$$Q \qquad \emptyset \qquad \emptyset$$

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.

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. Nevertheless, we see that in such a case, we can replace all occurrences of one of the two cliques by the other in all rows of the above form. This follows from the structure of H and the assumption that any vertex in G is connected to at least one vertex in any row of G.

Therefore, we can assume from now on that all the rows of the above form are covered by precisely three cliques in the 4-clique cover of H. Let C_0, C_1 and C_2 denote these three cliques such that C_i is assigned shift value i in every block. 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.

The Voting Scheme

In any row of H, whose label is of the form

$$L_0 \qquad L_1 \qquad L_2,$$

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 labeling 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 they contain in the above row. Thus in a row whose label has the form

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

each of C_0, C_1 and C_2 vote for Q.

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 claims summarize some useful observations.

Claim 3 *In a row whose row label is of the form,*

$$L_0 \qquad L_1 \qquad \emptyset$$

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 row label is of the form

$$Q \qquad \emptyset \qquad \emptyset$$

Therefore, 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. Thus it can only vote for either the entry in column 0 or the entry in column 1 of the given row. The claim follows. $\square$

Claim 4 *In a pair of rows of H of the form*

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

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 Claim 3, 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. Thus property (a) follows. 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. Observe now that 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. $\square$

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 claims.

Claim 5 *In a pair of rows of H of the form*

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

majority is always defined.

Proof : If all three critical cliques appear in this pair of rows, majority is clearly defined by Claim 4(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 Claim 4(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.

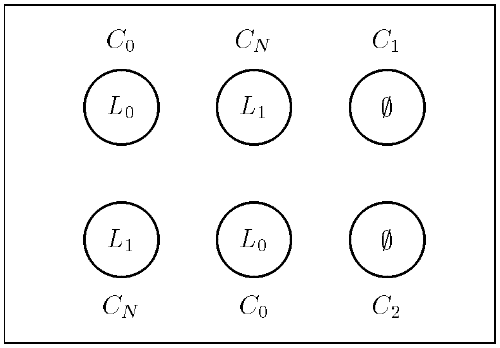

Figure 1: An example of critical cliques choosing L_1

Singling Out the Singletons

We have established so far that in every pair of rows which corresponds to two different orderings of a partition of the form $\langle L_0, L_1, \phi\rangle$, 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.

Lemma 6 *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 $Q = \{v_1, ..., v_q\}$ is the set of vertices in the $i^{\underline{\text{th}}}$ row of G, such that $\{v_j\}$ is elected by the majority of the critical cliques in this pair.*

Proof: Recall that we defined $X_j = \{v_{j+1}, ..., v_q\}$, $Y_j = \{v_1, ..., v_j\}$ and $Z_j = X_j \cup Y_{j-1}$, where $j \in [1..q-1]$. 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 $Y_1 = \{v_1\}$ 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 $X_{k-1} (= \{v_k, ..., v_q\})$ is elected in the $(k-1)^{\underline{\text{th}}}$ pair and $Y_k (= \{v_1, ..., v_k\})$ is elected in the $k^{\underline{\text{th}}}$ 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 show that this makes it impossible to cover all three vertices in the row R whose row label is:

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

We argue that at most one of the critical cliques can cover a vertex in R; this clearly suffices as there is only one non-critical clique.

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 which 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 the three sets, $Y_{k-1}, \{v_k\}$ 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 claim:

Claim 6 *In any block of H, consider a pair of rows, say P_1, with labels of the form:*

$$\begin{array}{lll} L_0 + S & L_1 & \emptyset \\ L_1 & L_0 + S & \emptyset \end{array},$$

and a clique C in H that contains 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 :

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

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 $\langle Z_k, \{v_k\}, \phi\rangle$. □

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.

Lemma 7 *Let $S_{CLQ} = \{v_1, ..., v_r\}$ be a set of vertices of G, such that v_i belongs to the $i^{\underline{\text{th}}}$ row of G, and $\{v_i\}$ is elected in the $i^{\underline{\text{th}}}$ block of H; then S_{CLQ} is a union of two cliques in G.*

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

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

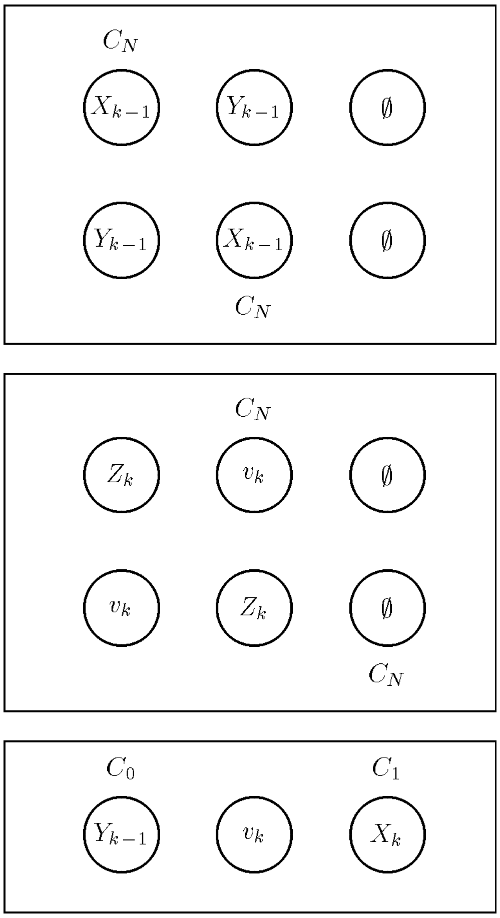

Figure 2: Critical cliques C_0 and C_1 vote for Y_{k-1} and v_k, respectively, in row r. This forces the non-critical clique C_N to cover 4 vertices, as indicated, which do not induce a clique in H.

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.

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 $v_{i'}$.

Now, regarding S_1, we are given $i \neq i'$ such that $v_i, v_{i'} \in S_1$, and such that C_N holds $\{v_i\}$ with shift j in P_i and it holds $\{v_{i'}\}$ with shift j' in $P_{i'}$. We show that, unless v_i is connected in G to $v_{i'}$, a 4-clique-cover, in which both v_i and $v_{i'}$ are elected, is not possible.

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). Hence, we have P_i that looks like:

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

and $P_{i'}$ that looks like:

$$\begin{array}{lll} \{v_{i'}\} & \emptyset & Q_{i'} - \{v_{i'}\} \\ \{v_{i'}\} & Q_{i'} - \{v_{i'}\} & \emptyset\ . \end{array}$$

The following claim is immediate now.

Claim 7 *Unless v_i is connected to $v_{i'}$ in G, H has no vertical edges connecting the first (second) row of P_i to 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 $P_{i'}$, i.e., C_N contains two vertices in the same column in both P_i and $P_{i'}$. If $j = j'$, by claim 7 we are done. Hence we assume, from now on, that $j \neq j'$; let $m \in \{0,1,2\}$ be such that $m \neq j$ and $m \neq j'$. Note that C_N does not contain any vertices in column m.

It must be the case that the critical cliques with shifts j' and m constitute the majority of votes in P_i, and the critical cliques with shifts j and m constitute the majority of votes in $P_{i'}$.

Since C_m votes for $\{v_i\}$ in P_i and $\{v_{i'}\}$ in $P_{i'}$, C_m contains vertices only in column m.

By claim 7, unless v_i is connected to $v_{i'}$, C_m 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, clique C_j in P_i and clique C'_j in P_i can contain either the vertex in the first row in column m in P_i and the second row in column m in $P_{i'}$, respectively, or vice versa. This yields a contradiction. □

4.4 Off with the 4th Clique

We can now conclude the following theorem:

Theorem 1 *Coloring a 3-colorable graph by 4 colors is NP-hard.*

Proof: From Lemma 7 we conclude that if $\bar{\chi}(H) \leq 4$ then $\omega(G) \geq \frac{r}{2}$, therefore $\omega(G) = r$. □

The following is a rather straightforward corollary of theorem 1:

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

Open Problems

There are still quite a few interesting open problems to consider concerning the hardness of chromatic-number. Is it NP-hard to color a 3-colorable graph with, say, 5 colors? Is it NP-hard for any constant? It may be the case that it is NP-hard to color a 3-colorable graph with $|G|^\epsilon$ colors, for some constant ϵ.

It would be also interesting to show hardness of coloring a k-colorable graph by k' colors, where k' is larger than any constant power of k.

On the other hand, it may be the case that these approximation problems have polynomial-time algorithm.

Acknowledgements

The authors would like to thank Michal Parnas, Steven Phillips and Carsten Lund for insightful discussions.

Oded Goldreich deserves an excellent plus grade for his most insightful comments on previous versions of the paper.

References

[AS92] S. Arora and S. Safra. Probabilistic Checking of Proofs; a New Characterization of NP. *Proc. 33rd IEEE Sym. on Foundation of Computer Science*, October 1992, pp. 1-13.

[AMSSS92] S. Arora, R. Motwani, S. Safra, M. Sudan and M. Szegedy. PCP and the hardness of approximation problems. Unpublished note, February 1992.

[ALMSS92] S. Arora, C. Lund, R. Motwani, M. Sudan and M. Szegedy, Proof verification and intractability of approximation problems. *Proc. 33rd IEEE Sym. on Foundation of Computer Science*, October 1992, pp. 14-23.

[B89] A. Blum. An $O(n^{0.4})$ approximation algorithm for 3-coloring. *Proc. 21st Annual Symposium on the Theory of Computing*, May 1989, pp. 535-542.

[B90] A. Blum. Some Tools for Approximate 3-Coloring. *Proc. 31st IEEE Sym. on Foundation of Computer Science*, October 1990, pp. 554-562.

[FGLSS91] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating clique is almost NP-complete. *Proc. 32nd IEEE Sym. on Foundations of Computer Science*, October 1991, pp. 2-12.

[GJ76] M. R. Garey and D. S. Johnson. The complexity of near-optimal graph coloring. *JACM*, vol. 23, 1976, pp. 43-49

[GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[H93] M. M. Halldórsson. A still better performance guarantee for approximate graph coloring. *Inform. Process. Lett.*, vol. 45, 1993, pp. 19-23

[Kar72] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

[LV89] N. Linial and U. Vazirani. Graph Products and Chromatic Numbers. *Proc. 30th IEEE Sym. on Foundations of Computer Science*, October 1989, pp. 124-128.

[LY92] C. Lund, M. Yannakakis. On the Hardness of Minimization Problems. Manuscript, July 1992.

[PS92] S. Phillips and S. Safra. PCP and tighter bounds for approximating MAX-SNP. Manuscript, April 1992.

[W83] A. Wigderson. Improving the performance guarantee for approximate graph coloring. *JACM*, vol. 30, No. 4, 1983, pp. 729-735