Approximating the Achromatic Number Problem on Bipartite Graphs

Guy Kortsarz and Sunil Shende*

Department of Computer Science, Rutgers University, Camden, NJ 08102
{guyk,shende}@camden.rutgers.edu

Abstract. The achromatic number of a graph is the largest number of colors needed to legally color the vertices of the graph so that adjacent vertices get different colors and for every pair of distinct colors c_1, c_2 there exists at least one edge whose endpoints are colored by c_1, c_2. We give a greedy $O(n^{4/5})$ ratio approximation for the problem of finding the achromatic number of a bipartite graph with n vertices. The previous best known ratio was $n \cdot \log\log n / \log n$ [12]. We also establish the first non-constant hardness of approximation ratio for the achromatic number problem; in particular, this hardness result also gives the first such result for bipartite graphs. We show that unless NP has a randomized quasi-polynomial algorithm, it is not possible to approximate achromatic number on bipartite graph within a factor of $(\ln n)^{1/4-\epsilon}$. The methods used for proving the hardness result build upon the combination of one-round, two-provers techniques and zero-knowledge techniques inspired by Feige et.al. [6].

1 Introduction

A *proper coloring* of a graph $G(V, E)$ is an assignment of colors to V such that adjacent vertices are assigned different colors. It follows that each color class (i.e. the subset of vertices assigned the same color) is an independent set. A k-coloring is one that uses k colors. A coloring is said to be *complete* if for every pair of distinct colors, there exist two adjacent vertices which are assigned these two colors. The *achromatic number* $\psi^*(G)$ of a graph G is the *largest* number k such that G has a complete k-coloring.

A large body of work has been devoted to studying the achromatic number problem which has applications in clustering and network design (see the surveys by Edwards [4] and by Hughes and MacGillivray [11]). Yannakakis and Gavril [15] proved that the achromatic number problem is NP-hard. Farber et.al. [5] show that the problem is NP-hard on bipartite graphs. Bodlaender [1] established that the problem is NP-hard on graphs that are simultaneously co-graphs and interval graphs. Cairnie and Edwards [2] show that the problem is NP-hard on trees.

* Research supported in part under grant no. 9903240 awarded by the National Science Foundation.

G. Di Battista and U. Zwick (Eds.): ESA 2003, LNCS 2832, pp. 385–396, 2003.

Given the intractability of solving the problem optimally (assuming, of course, that $P \neq NP$), the natural approach is to seek a *guaranteed* approximation to the achromatic number. An *approximation algorithm* with ratio $\alpha \geq 1$ for the achromatic number problem is an algorithm takes as input a graph G and produces a a complete coloring of G with at least $\psi^*(G)/\alpha$ colors in time polynomial in the size of G. Let n denote the number of vertices in graph G. We will use the notation ψ^* for $\psi^*(G)$ when G is clear from the context.

Chaudhary and Vishwanathan [3] gave the first sublinear approximation algorithm for the problem, with an approximation ratio $O(n/\sqrt{\log n})$. Kortsarz and Krauthgamer [12] improve this ratio slightly to $O(n \cdot \log\log n/\log n)$. It has been conjectured [3] that the achromatic number problem *on general graphs* can be approximated within a ratio of $O(\sqrt{\psi^*})$. The conjecture is partially proved in [3] with an algorithm that gives a $O(\sqrt{\psi^*}) = O(n^{7/20})$ ratio approximation for graphs with *girth* (length of the shortest simple cycle) at least 7. Krysta and Loryś [13] give an algorithm with approximation ratio $O(\sqrt{\psi^*}) = O(n^{3/8})$ for graphs with girth at least 6. In [12], the conjecture is proved for graphs of girth 5 with an algorithm giving an $O(\sqrt{\psi^*})$ ratio approximation for such graphs. In terms of n, the best ratio known for graphs of girth 5 is $O(n^{1/3})$ (see [12]).

From the direction of lower bounds on approximability, the first (and only known) hardness of approximation result for general graphs was given in [12], specifically that the problem admits no $2-\epsilon$ ratio approximation algorithm, unless P=NP. It could be that no $n^{1-\epsilon}$ ratio approximation algorithm (for any constant $\epsilon > 0$) is possible for general graphs (unless, say, P=NP). An $\Omega(n^{1-\epsilon})$ inapproximability result does exist for the maximum independent set problem [8] and the achromatic number problem and the maximum independent set problem are, after all, closely related.

On another negative note, consider the minimum maximal independent set problem. A possible "greedy" approach for finding an achromatic partition is to iteratively remove from the graph maximal independent sets of small size (maximality here is with respect to containment). However, the problem of finding a *minimum maximal independent set* cannot be approximated within ratio $n^{1-\epsilon}$ for any $\epsilon > 0$, unless P=NP [7].

To summarize, large girth (i.e. girth greater than 4) is known to be a sufficient condition for a relatively low ratio approximation to exist. It is not known if the absence of triangles helps in finding a good ratio approximation algorithm for the problem. All the current results thus point naturally to the next frontier: the family of bipartite graphs.

1.1 Our Results

We give a combinatorial greedy approximation algorithm for the achromatic number problem on bipartite graphs achieving a ratio of $O(n^{4/5})$ and hence breaking the $\tilde{O}(n)$ barrier (the upper bound for general graphs [12]). We also give a hardness result that is both the first non-constant lower bound on approximation for the problem on general graphs, and the first lower bound on approximation for bipartite graphs. We prove that unless $NP \subseteq RTIME(n^{\text{polylog } n})$, the problem does not admit an $(\ln n)^{1/4-\epsilon}$ ratio approximation, for any constant $\epsilon > 0$.

This improves the hardness result of 2 on *general graphs* [12]. Note that the result in [12] constructs a graph with large cliques, which therefore is not bipartite. The best previous hardness result for bipartite graphs was only the NP-hardness result [5].

2 Preliminaries

We say that a vertex v is *adjacent* to a set of vertices U if v is adjacent to at least one vertex in U. Otherwise, v is *independent* of U. Subsets U and W are adjacent if for some $u \in U$ and $w \in W$, the graph has the edge (u, w). U *covers* W if every vertex $w \in W$ is adjacent to U. We note that in a complete coloring of G, every pair of distinct color classes are adjacent to each other. For any subset U of vertices, let $G[U]$ be the subgraph of G induced by U. A *partial* complete coloring of G is a complete coloring of some induced subgraph $G[U]$, namely, a coloring of U such that all color classes are pairwise adjacent. Lemmas 1 and 2 below are well known [3,4,13,14]:

Lemma 1. *A partial complete coloring can be extended greedily to a complete coloring of the entire graph.*

Lemma 2. *Consider v, an arbitrary vertex in G, and let $G \setminus v$ denote the graph resulting from removing v and its incident edges from G. Then, $\psi^*(G) - 1 \leq \psi^*(G \setminus v) \leq \psi^*(G)$.*

A collection M of edges in a graph G is called a *matching* if no two edges in M have a common endpoint. The matching M is called *semi-independent* if the edges (and their endpoints) in M can be indexed as $M = \{(x_1, y_1), \ldots, (x_k, y_k)\}$ such that both $X = \{x_1, \ldots, x_k\}$ and $Y = \{y_1, \ldots, y_k\}$ are independent sets, and for all $j > i$, it holds that x_i is *not* adjacent to y_j. As a special case, if x_i is not adjacent to y_j for all i, j, then the matching is said to be *independent*. A semi-independent matching can be used to obtain a partial complete coloring, as demonstrated in the next lemma; a weaker version, based on an independent matching, is used in [3].

Lemma 3. [14] *Given a semi-independent matching of size $\binom{t}{2}$ in a graph, a partial complete t-coloring of the graph (i.e. with t color classes) can be computed efficiently.*

Now, consider a presentation of a bipartite graph $G(U, V, E)$ with independent sets U and V forming the bipartition, and with edges in E connecting U and V. Assume that U has no isolated (degree 0) vertices. If $\Delta(V)$, the largest degree of a vertex in V, is suitably small, then by repeatedly removing a *star* formed by the current largest degree vertex in V and its neighbors in U, we can obtain a collection of at least $|U|/\Delta(V)$ stars. By picking a single edge out of every star, we get a semi-independent matching of size at least $|U|/\Delta(V)$. Applying Lemmas 1 and 3, we get the following result.

Lemma 4. *Let $G(U, V, E)$ be a bipartite graph with no isolated vertices in U. Then, the star removal algorithm produces an achromatic partition of size at least $\Omega(\sqrt{|U|/\Delta(V)})$.*

Hell and Miller [9,10] define the following equivalence relation (called the reducing congruence) on the vertex set of G (see also [4,11]). Two vertices in G are *equivalent* if and only if they have the same set of neighbors in the graph. We denote by $S(v, G)$, the subset of vertices that are equivalent to v under the reducing congruence for G; we omit G when it is clear from the context. Assume that the vertices of G are indexed so that $S(v_1), \ldots, S(v_q)$ denote the equivalence classes of vertices, where q is the number of distinct equivalence classes. Note that two equivalent vertices cannot be adjacent to each other in G, so $S(v_i)$ forms an independent set in G. The *equivalence graph* (also called the *reduced graph*) of G, denoted G^*, is a graph whose vertices are the equivalence classes $S(v_i)$ $(1 \leq i \leq q)$ and whose edges connect $S(v_i), S(v_j)$ whenever the set $S(v_i)$ is adjacent to the set $S(v_j)$.

Lemma 5. [12] *A partial complete coloring of G^* can be extended to a complete coloring of G. Hence, $\psi^*(G) \geq \psi^*(G^*)$.*

Theorem 1. [12] *Let G be a bipartite graph with q equivalence classes of vertices. Then, there is an efficient algorithm to compute an achromatic partition of G of size at least $\min\{\psi^*/q, \sqrt{\psi^*}\}$. Thus, the achromatic number of a bipartite graph can be approximated within a ratio of $O(\max\{q, \sqrt{\psi^*}\})$.*

Let the *reduced* degree $d^*(v, G)$ be the degree of the vertex $S(v)$ in the reduced graph G^*; equivalently, this is the maximum number of pairwise non-equivalent neighbors of v. Then, we have:

Lemma 6. *Let v, w be a pair of vertices of G such that $S(v) \neq S(w)$ and $d^*(w) \geq d^*(v)$. Then there is a vertex z adjacent to w but not to v.*

Proof. Assume that every neighbor of w is also a neighbor of v. Since $d^*(w) \geq d^*(v)$, it follows that v and w have exactly the same set of neighbors contradicting $S(w) \neq S(v)$. □

3 The Approximation Algorithm

Let $\psi^*(G)$ denote the maximum number of parts in an achromatic partition of a graph G (we omit G in the notation when the graph is clear from the context). We may assume that ψ^* is known, e.g. by exhaustively searching over the n possible values or by using binary search. Throughout, an algorithm is considered efficient if it runs in polynomial time. Let $G(U, V, E)$ be a bipartite graph, and consider subsets $U' \subseteq U$ and $V' \subseteq V$. The (bipartite) subgraph induced by U' and V' is denoted by $G[U', V']$ where the (implicit) edge set is the restriction of E to edges between U' and V'.

Our approach is to iteratively construct an achromatic partition of an induced subgraph of $G[U, V]$. Towards this end, we greedily remove a *small*, *independent* set of vertices A_i in each iteration while also deleting some other vertices and edges. The invariant maintained by the algorithm is that A_i always has a U-vertex and the subset of U vertices that survive the iteration is *covered* by A_i. This ensures that the collection $\mathcal{A} = \{A_i : \ i \geq 1\}$, forms a partial complete coloring of G.

To obtain such a collection $\mathcal{A}$ with large cardinality, we need to avoid deleting too many *non-isolated* vertices during the iterations since the decrease in achromatic number may be as large as the number of deleted, non-isolated vertices (by Lemma 2, while noting that the achromatic number remains unchanged under deletions of isolated vertices). For every $i \geq 1$, consider the sequence of induced subgraphs of G over the first $(i-1)$ iterations, viz. the sequence $G_0 \supset G_1 \ldots \supset G_{i-1}$ where G_k, $0 \leq k < i$, is the surviving subgraph at the beginning of the $(k+1)^{th}$ iteration. The algorithm uses the following notion of *safety* for deletions in the i^{th} iteration:

Definition 1. *During iteration i, the deletion of some existing set of vertices S from G_{i-1} is said to be* ***safe*** *for G_{i-1} if the number of non-isolated vertices (including those in S) cumulatively removed from the initial subgraph G_0 is at most $\psi^*(G)/4$.*

3.1 Formal Description of the Algorithm

We first provide a few notational abbreviations that simplify the formal description and are used in the subsequent analysis of the algorithm. A set is called **heavy** if it contains at least $n^{1/5}$ vertices. Otherwise, it is called **light**. Given a subset of vertices U belonging to graph G, we denote by $d^*(v, U, G)$ the maximum number of pairwise non-equivalent neighbors that v has in U. v is called U**-heavy** if $d^*(v, U, G) \geq n^{1/5}$.

The approximation algorithm ABIP described below produces an achromatic coloring of G. It invokes the procedure PARTITION (whose description follows that of ABIP) twice, each time on a different induced subgraph of G. The procedure returns a partial complete achromatic partition of its input subgraph.

Algorithm ABIP.
Input: $G(U, V)$, a bipartite graph.

1. Let $\mathcal{A}_1 = \text{PARTITION}(G[U, V])$, and let $G^{[1]} = G[U^{[1]}, V^{[1]}]$ be the induced subgraph that remains when the procedure halts.
2. Let $\mathcal{A}_2 = \text{PARTITION}(G[V^{[1]}, U^{[1]}])$; note that the roles of the bipartitions are *interchanged*. Let $G^{[2]} = G[U^{[2]}, V^{[2]}]$ be the induced subgraph that remains when this second application halts.
3. If either of the achromatic partitions $\mathcal{A}_1$ or $\mathcal{A}_2$ is of size at least $\psi^*/(16 \cdot n^{1/5})$, then the corresponding partial complete coloring is extended to a complete achromatic coloring of G which is returned as final output.
4. Otherwise, apply the algorithm of Theorem 1 on the subgraph $G^{[2]}$. A partial complete coloring is returned which can then be extended to a complete achromatic coloring of G returned as final output.

Procedure PARTITION
Input: $G_0(U_0, V_0)$, an induced subgraph of a bipartite graph $G(U, V)$.

1. **if** $\psi^* < 8 \cdot n^{4/5}$, **return** an arbitrary achromatic partition.
2. $\mathcal{A} \leftarrow \{\}$ /* $\mathcal{A}$ contains the collection of A_i's computed so far */
3. **for** $i = 1, 2, \ldots$ /* Iteration i */
   a) **if** there are no light U_{i-1}-equivalence classes in G_{i-1}, then **break**
   b) Choose a vertex $u \in U_{i-1}$ with smallest equivalence class size and smallest reduced degree in G_{i-1} (break ties arbitrarily).
   c) Remove $S(u)$ from U_{i-1}, the neighbors of u from V_{i-1} and let $G^{'} = G[U^{'}, V^{'}]$ be the resulting induced subgraph.
   $C_i \leftarrow \emptyset$
   d) **while** $U^{'} \neq \emptyset$ and there exists a $U^{'}$-heavy vertex in $V^{'}$ **do**
      i. Choose v with largest reduced degree $d^*(v, U^{'}, G^{'})$ in the current graph $G^{'}$.
      ii. Add v to C_i
      iii. Remove v from $V^{'}$ and its neighbors from $U^{'}$
   e) Let $q^{'}$ be the number of $U^{'}$-equivalence classes in $G^{'*}$.
   f) **if** $q^{'} > n^{3/5}$, let $\mathcal{A}$ be the partition obtained by applying the star removal algorithm to $G^{'}$ (see Lemma 4).
   break
   g) Let D_i be the vertices in $U^{'}$ with light equivalence classes in $G^{'}$
   h) **for** every heavy $U^{'}$-equivalence class $S(w)$ **do**
   add an arbitrary neighbor of $S(w)$ to C_i.
   i) $A_i \leftarrow S(u, G_{i-1}) \cup C_i$;
   Let $L_i \subseteq V_{i-1}$ be the set of isolated vertices in the graph $G[U_{i-1} \setminus A_i, V_{i-1} \setminus A_i]$
   j) **if** it is **not safe** to delete $(A_i \cup D_i)$ from G_{i-1} then **break**
   k) add A_i to $\mathcal{A}$;
   Remove $S(u, G_{i-1})$ and D_i from U_{i-1} leaving U_i
   Remove C_i and L_i from V_{i-1} leaving V_i
   $G_i \leftarrow G[U_i, V_i]$
4. **return** $\mathcal{A}$

3.2 The Approximation Ratio

We now analyze the approximation ratio of ABIP; detailed proofs have been omitted due to space considerations. Our goal is to show that the approximation ratio is bounded by $O(n^{4/5})$. The analysis is conducted under the assumption that $\psi^*(G) \geq 8 \cdot n^{4/5}$. Otherwise, returning an arbitrary achromatic partition (say, the original bipartition of size 2), as done in line 1 of PARTITION, trivially gives an $O(n^{4/5})$ ratio.

We start by observing that the loop on line 3 in PARTITION could exit in one of *three mutually exclusive ways* during some iteration $(k + 1) \geq 1$. We say that PARTITION takes

- **exit 1** if the star removal algorithm can be applied (at line 3f) during the iteration,
- **exit 2** if just prior to the end of the iteration, it is found that the current deletion of $(A_{k+1} \cup D_{k+1})$ is not safe for G_k (at line 3j), or
- **exit 3** if at the beginning of the iteration,there are no light U_k-equivalence classes in G_k (at line 3a).

Note that the induced subgraphs G_i $(i \geq 1)$ form a monotone decreasing chain so if the star removal algorithm (exit 1) is not applicable at any intermediate stage, then PARTITION will eventually take one of the latter two exits, i.e. the procedure *always* terminates. We say that iteration $i \geq 1$ in an execution of PARTITION is *successful* if none of the exits are triggered during the iteration, i.e. the procedure continues with the the next iteration. Let $(k+1) \geq 1$ be the first unsuccessful iteration.

Lemma 7. *If* PARTITION *takes exit 1 during iteration* $(k+1)$*, then the achromatic partition returned has size at least* $n^{1/5}$*. As* $\psi^* \leq n$*, an* $O(n^{4/5})$*−ratio is derived.*

Proof. Consider $U^{'}$ and $V^{'}$ when exit 1 is triggered. It is easy to show that every vertex $w \in U^{'}$ is adjacent to $V^{'}$. Furthermore, the inner loop condition (line 3d) guarantees that every vertex in $V^{'}$ is adjacent to at most $n^{1/5}$ $U^{'}$-equivalence classes. When the star removal algorithm is applied, $q^{'}$ (the number of $U^{'}$-equivalence classes) is at least $n^{3/5}$. From the discussion preceding Lemma 4, it is easy to see that the star removal algorithm will produce a collection of at least $\sqrt{n^{3/5}/n^{1/5}} = n^{1/5}$ stars. Thus, an achromatic partition of size at least $n^{1/5}$ is returned as claimed. □

Next, we show that if the procedure takes exit 2 during iteration $(k+1)$ because an unsafe deletion was imminent, then it must be the case that k is large and hence, that we have already computed a large achromatic partition $\mathcal{A} = \{A_1, A_2, \ldots A_k\}$. To this end, we establish a sequence of claims.

Claim 1 *For all* i *such that* $1 \leq i \leq k$*, the set* A_i *is an independent set and is adjacent to* A_j *for every* $j \in [i+1, k]$*. Equivalently,* $\mathcal{A}$ *is an achromatic partition of the subgraph* $G[\cup_{1 \leq i \leq k} A_i]$*.*

Proof. We first verify that at the end of a successful iteration i, the set of vertices A_i is an independent set. By construction, the vertices retained in U_i at the end of the iteration are all covered by C_i. The set A_j, for $i < j \leq k$, contains at least one vertex in $U_{j-1} \subset U_i$. Hence there is always an edge between A_i and A_j. □

Claim 2 *For all* i *such that* $1 \leq i \leq k$*, the size of the set* $(A_i \cup D_i)$*, just prior to executing the safety check on line 3j, is bounded by* $4n^{4/5}$*.*

Proof. By construction, $A_i = S(u) \cup C_i$ prior to executing line 3j. We know that $S(u)$ is a light equivalence class and hence, $\mid S(u) \mid < n^{1/5}$. A vertex $v \in V_{i-1}$ is added to C_i either during the inner loop (line 3d) or later, if it happens to be adjacent to a heavy U'-equivalence class (line 3h). In both cases, we can show that no more than $n^{4/5}$ vertices could have been added to C_i. Together, we have at most $3 \cdot n^{4/5}$ being added to A_i. Now, U' has at most $n^{3/5}$ light equivalence classes when control reaches line 3g. Since the vertices in D_i just prior to executing the safety check are simply those belonging to such light U'-equivalence classes, there are at most $n^{1/5} \cdot n^{3/5} = n^{4/5}$ vertices in D_i. □

Claim 3 *If the first k iterations are successful, then the difference,* $\psi^*(G_0) - \psi^*(G_k)$*, is at most* $4k \cdot n^{4/5}$.

Proof. Follows from Lemma 2 and Claim 2. □

Lemma 8. *If* PARTITION *takes exit 2 during iteration* $(k+1)$*, then the achromatic partition returned has size at least* $\lfloor \psi^*(G)/16n^{4/5} \rfloor$ *thus giving an* $O(n^{4/5})$ *ratio approximation.*

Proof. Since the first k iterations were successful, it follows that for each $i \in [1, k]$, it is safe to delete $(A_i \cup D_i)$. However, it is unsafe to delete $(A_{k+1} \cup D_{k+1})$ and by Definition 1 and Claim 3, this can only happen if

$$4(k+1)n^{4/5} \;>\; \psi^*(G)/4.$$

Hence $\mathcal{A} = \{A_1, A_2, \ldots A_k\}$, which is an achromatic partition of the subgraph $G[\cup_{1 \le i \le k} A_i]$ by Claim 1, has size $k \ge \lfloor \psi^*(G)/(16n^{4/5}) \rfloor$. Applying Lemma 1, we conclude that a complete achromatic coloring of G with at least $\lfloor \psi^*(G)/(16n^{4/5}) \rfloor$ colors can be computed thus giving an $O(n^{4/5})$ ratio approximation. □

Finally, if PARTITION takes exit 3 in iteration $(k+1)$, then we have two possibilities. If $k \ge \lfloor \psi^*(G)/(16n^{4/5}) \rfloor$, a sufficiently large partition has been found and we are done. Otherwise, $k < \lfloor \psi^*(G)/(16n^{4/5}) \rfloor$ and we may not necessarily have a good approximation ratio. However, note that G_k, the graph at the beginning of iteration $(k+1)$, has no light U_k-equivalence classes (this triggers the exit condition). Hence, U_k has no more than $n^{4/5}$ equivalence classes that are all heavy, since each heavy class has at least $n^{1/5}$ vertices and $|U_k| \le n$.

Claim 4 *Assume that both applications of* PARTITION *on lines 1 and 2 of algorithm* ABIP *respectively take exit 3. Let* q_1 *(respectively,* q_2*) be the number of light* $U^{[1]}$*-equivalence classes in* $G^{[1]}$ *(respectively, the number of light* $U^{[2]}$*-equivalence classes in* $G^{[2]}$*). Then, the graph* $G^{[2]}$ *has achromatic number at least* $\psi^*(G)/2$ *and at most a total of* $(q_1 + q_2) \le 2n^{4/5}$ *equivalence classes.*

Proof. Observe that the removal of vertices (along with all their incident edges) from a graph cannot increase the number of equivalence classes: two vertices that were equivalent before the removal, remain equivalent after. Hence, the number of $V^{[2]}$ equivalence classes is at most q_1 (note that the partitions are interchanged before the second application of PARTITION on line 2). Thus $G^{[2]}$

has at most a total of $(q_1 + q_2)$ equivalence classes. The discussion preceding the claim shows that $(q_1 + q_2)$ is bounded above by $2n^{4/5}$. Since neither application of PARTITION took exit 2, the vertices deleted during both applications were safe for deletion. Hence, the net decrease in the achromatic number is at most $2\psi^*(G)/4 = \psi^*(G)/2$. □

Theorem 2. *The algorithm* ABIP *has an approximation ratio of* $O(n^{4/5})$.

Proof. By Lemmas 7 and 8, if either of the two applications of PARTITION take exits 1 or 2, then we are guaranteed a ratio of $O(n^{4/5})$. If both applications of PARTITION on lines 1 and 2 of ABIP halt on exit condition 3, then an application of the algorithm of Theorem 1 on graph $G^{[2]}$ (see line 4 of ABIP) provides an $O(q)$ approximation ratio for $G^{[2]}$ where q is the number of equivalence classes of $G^{[2]}$. By claim 4, this achromatic coloring is a partial complete coloring of G with ratio $O(n^{4/5})$. □

4 A Lower Bound for Bipartite Graphs

Let $G(U, V, E)$ be a bipartite graph. A *set-cover* (of V) in G is a subset $S \subseteq U$ such that S covers V, i.e. every vertex in V has a neighbor in S. Throughout, we assume that the intended bipartition $[U, V]$ is given explicitly as part of the input, and that every vertex in V can be covered. A *set-cover packing* in the bipartite graph $G(U, V, E)$ is a collection of pairwise-disjoint set-covers of V. The *set-cover packing problem* is to find in an input bipartite graph (as above), a maximum number of pairwise-disjoint set-covers of V. Our lower bound argument uses a modification of a construction by Feige et.al. [6] that creates a set-cover packing instance from an instance of an NP-complete problem. Details of the construction are omitted due to space limitations and will appear in the full version of the paper. The main result obtained is the following:

Theorem 3. *For every* $\epsilon > 0$*, if NP does not have a (randomized) quasi-polynomial algorithm then the achromatic number problem on bipartite graphs admits no approximation ratio better than* $(\ln n)^{1/4-\epsilon}/16$.

4.1 The Set-Cover Packing Instance [6]

Our lower bound construction uses some parts of the construction in [6]. That paper gives a reduction from an arbitrary NP-complete problem instance I to a set-cover packing instance[1] $G(U, V, E)$, with $|U| + |V| = n$.

The idea is to use a collection of disjoint sets of vertices $\{A_i : 1 \leq i \leq q\}$ and $\{B_i : 1 \leq i \leq q\}$; all these sets have the same size N where N is a parameter. In the construction, $U = (\bigcup_{i=1}^{q} A_i) \cup (\bigcup_{i=1}^{q} B_i)$. Also, the set $V = \bigcup M(A_i, B_j)$ with the union taken over certain pre-defined pairs (A_i, B_j) that arise from the NP-complete instance. The set $M(A_i, B_j)$ is called a *ground-set*. The reduction uses randomization to specify the set of edges E in the bipartite graph with the following properties:

[1] The construction described here corresponds to the construction in [6] for the special case of two provers.

1. If I is a yes-instance of the NP-complete problem, then U can be decomposed into N vertex-disjoint set-covers $S_1, \ldots, S_N$ of V. Each set-cover contains a unique A vertex and a unique B vertex for every A, B. Each S_i is an *exact cover* of V in the sense that every V-vertex has a unique neighbor in S_i.
2. In the case of a no-instance, the following properties hold:
   a) **The A, B property:** Only the $A_i \cup B_j$ vertices are connected in G to $M(A_i, B_j)$. **Comment:** The next properties concern the induced subgraph $G[(A_i \cup B_j), M(A_i, B_j)]$.
   b) **The random half property:** Each $a \in A_i$ and $b \in B_j$ is connected in $M(A_i, B_j)$ to a random half of $M(A_i, B_j)$.
   c) **The independence property:** For every $a \in A_i$ and $b \in B_j$, the collection of neighbors of a in $M(A_i, B_j)$ and the collection of neighbors of b in $M(A_i, B_j)$ are mutually independent random variables.
   d) **The equality or independence property:** The neighbors of two vertices $a, a' \in A_i$ in $M(A_i, B_j)$ are either the same, or else their neighbors in $M(A_i, B_j)$ are mutually independent random variables. A similar statement holds for a pair of vertices $b, b' \in B_j$.

Thus, vertices $a \in A_i$ and $b \in B_j$ have, on average, $|M(A_i, B_j)|/4$ common neighbors in $M(A_i, B_j)$ because a and b are joined to two independent random halves in $M(A_i, B_j)$.

4.2 Our Construction

The basic idea is similar to the above construction, namely that we wish to convert a **yes** instance of the NP-complete problem to a bipartite graph with a "large" achromatic partition and a **no** instance to a bipartite graph with a "small" achromatic partition. Towards this end, we extend the construction in [6] as follows.

Construction of a yes instance: A *duplication* of a vertex $u \in U$ involves adding to U a new copy of u connected to the neighbors of u in V. By appropriately duplicating the original vertices in U, we can make the number of vertex-disjoint set-covers larger. Specifically, we can duplicate vertices in U to ensure that every A and B set has $|V|$ elements and hence, the number of set-covers in the packing for a **yes** instance becomes $|V|$ as well (recall, from the previous discussion, that for a **yes** instance, each set-cover contains exactly one A vertex and exactly one B vertex, and so $|A| = |B| = |V|$ is the number of set-covers in the packing).

Using some technical modifications, we can also make G regular. Hence, G admits a perfect matching where each $v \in U$ is matched to some corresponding vertex $m(v) \in V$. Observe that for the case of a yes instance, the number of $m(v)$ vertices, namely, $|V|$, is equal to the number of set-covers in the set-cover packing.

The idea now is to form a collection of $|V|$ sets, one for each $v \in U$, by adding the matched vertex $m(v)$ to an (exact) set-cover S_i. However, the resulting sets are not independent sets because each S_i is an exact set-cover of V and hence contains a neighbor of $m(v)$. But this problem can be fixed by ensuring that during the duplication process, a *special* copy of v is inserted; specifically, the special copy gets all the neighbors of v *except* $m(v)$ as its neighbors. With this modification, the collection consists of independent sets which form an achromatic partition because each of the S_i's are exact. This implies that in the case of a yes instance, the corresponding bipartite graph admits a size $|V|$ complete coloring.

Construction of a no instance: The main technical difficulty is showing that in a case of a no-instance, the maximum size achromatic partition is "small". Let $X_1, X_2, X_3, \ldots$ be the color classes in a maximum coloring in the case of a no-instance. Consider the contribution of A, B to the solution, i.e. how many vertices from the A and B sets belong to any X_i. Observe that in the case of a yes instance, each color contains one vertex from every A and every B. If we could color the graph corresponding to a no instance with "many" colors, this would mean that each X_i has to contain only "few" A and "few" B vertices. Similarly, for a yes instance, each color contains exactly one V vertex. Therefore, each X_i must contain only "few" $M(A, B)$ vertices.

Say, for example, that for every i, X_i satisfies the conditions: $|X_i \cap M(A, B)| = 1$ and $|X_i \cap (A \cup B)| = 1$ as in a yes instance. Let $v_2 \in (X_i \cap M(A, B))$ and $v_1 \in (X_i \cap (A \cup B))$. Observe that the random half property implies that the edge (v_1, v_2) exists only with probability 1/2.

On the other hand, we note that if a coloring has close to $|V|$ colors, events as the one above should hold true for $\Omega(|V|^2)$ pairs. since every X_i and X_j must by definition share at least one edge. The equality or independence property ensures that "many" (but not all) of events such as the ones above are independent. Therefore, it is unlikely that polynomially many such independent events can occur simultaneously.

The above claim has its limits. If we take subsets of size, say, $\sqrt{2 \log n}$ from $A \cup B$ and $M(A, B)$ into every X_i, X_j, then the number of "edge-candidates" between X_i and X_j is now $(\sqrt{2 \log n})^2 = 2 \log n$. Namely, each one of the $\log n$ pairs is a possible candidate edge, so that if it is chosen by the randomized choice, it guarantees at least one edge between X_i and X_j as required by a legal achromatic partition. Every candidate edge exists with probability 1/2. Thus the probability for at least one edge between X_i and X_j could be as high as $1 - 1/n^2$, and it is not unlikely that "many" (like $|V|^2 < n^2$) of these events happen simultaneously.

Note this if each X_i contains roughly $\sqrt{\log n}$ vertices from every A, B and from every $M(A, B)$, the number of colors in the solution could be as high as $|V|/\sqrt{\log n}$. This gives a limitation for this method (namely, we can not expect a hardness result beyond $\sqrt{\log n}$). In fact, the hardness result we are able to prove is only $(\log n)^{1/4-\epsilon}$ due to the fact that the events described above are not really totally independent.

Acknowledgments. The first author would like to thank Robert Krauthgamer and Magnus M. Halldórsson for useful discussions and comments.

References

1. H. L. Bodlaender. Achromatic number is NP-complete for cographs and interval graphs. *Inform. Process. Lett.*, 31(3):135–138, 1989.
2. N. Cairnie and K. Edwards. Some results on the achromatic number. *J. Graph Theory*, 26(3):129–136, 1997.
3. A. Chaudhary and S. Vishwanathan. Approximation algorithms for the achromatic number. In *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 558–563, 1997.
4. K. Edwards. The harmonious chromatic number and the achromatic number. In *Surveys in combinatorics, 1997 (London)*, pages 13–47 .
5. M. Farber, G. Hahn, P. Hell, and D. Miller. Concerning the achromatic number of graphs. *J. Combin. Theory Ser. B*, 40(1):21–39, 1986.
6. U. Feige, M. Halldórsson, G. Kortsarz and A. Srinivasan. Approximating the domatic number. Accepted to Siam J. on Computing conditioned on a revision.
7. M. M. Halldórsson. Approximating the minimum maximal independence number. *Inform. Process. Lett.*, 46(4):169–172, 1993.
8. J. Hastad, Clique is Hard to Approximate within n to the power 1-epsilon, *Acta Mathematica*, Vol 182, 1999, pp 105–142.
9. P. Hell and D. J. Miller. On forbidden quotients and the achromatic number. In *Proceedings of the 5th British Combinatorial Conference (1975)*, pages 283–292. Congressus Numerantium, No. XV. Utilitas Math., 1976.
10. P. Hell and D. J. Miller. Achromatic numbers and graph operations. *Discrete Math.*, 108(1-3):297–305, 1992.
11. F. Hughes and G. MacGillivray. The achromatic number of graphs: a survey and some new results. *Bull. Inst. Combin. Appl.*, 19:27–56, 1997.
12. G. Kortsarz and R. Krauthgamer. On approximating the achromatic number. *Siam Journal on Discrete Mathematics*, vol 14, No. 3, pages: 408–422
13. P. Krysta and K. Loryś. Efficient approximation algorithms for the achromatic number. In *ESA '99 (Prague)*, pages 402–413. Springer, 1999.
14. A. Máté. A lower estimate for the achromatic number of irreducible graphs. *Discrete Math.*, 33(2):171–183, 1981.
15. M. Yannakakis and F. Gavril. Edge dominating sets in graphs. *SIAM J. Appl. Math.*, 38(3):364–372, 1980.