

Vertex Cover in Graphs with Locally Few Colors

Fabian Kuhn¹

Monaldo Mastrolilli²

fabian.kuhn@usi.ch monaldo@idsia.ch

¹Faculty of Informatics, University of Lugano (USI), 6904 Lugano, Switzerland

²Dalle Molle Institute for Artificial Intelligence (IDSIA), 6928 Manno, Switzerland

June 27, 2011

Abstract

In [14], Erdős et al. defined the local chromatic number of a graph as the minimum number of colors that must appear within distance 1 of a vertex. For any $\Delta \geq 2$, there are graphs with arbitrarily large chromatic number that can be colored so that (i) no vertex neighborhood contains more than Δ different colors (*bounded local colorability*), and (ii) adjacent vertices from two color classes induce a complete bipartite graph (*biclique coloring*).

We investigate the weighted vertex cover problem in graphs when a locally bounded coloring is given. This generalizes the vertex cover problem in bounded degree graphs to a class of graphs with arbitrarily large chromatic number. Assuming the Unique Game Conjecture (UGC), we provide a tight characterization. We prove that it is UGC-hard to improve the approximation ratio of $2 - 2/(\Delta + 1)$ if the given local coloring is not a biclique coloring. A matching upper bound is also provided. Vice versa, when properties (i) and (ii) hold, we present a randomized algorithm with approximation ratio of $2 - \Omega(1) \frac{\ln \ln \Delta}{\ln \Delta}$. This matches known inapproximability results for the special case of bounded degree graphs.

Moreover, we show that when both the above two properties (i) and (ii) hold, the obtained result finds a natural application in a classical scheduling problem, namely the precedence constrained single machine scheduling problem to minimize the total weighted completion time, denoted as $1|prec| \sum w_j C_j$ in standard scheduling notation. In a series of recent papers it was established that this scheduling problem is a special case of the minimum weighted vertex cover in graphs $G_{\mathcal{P}}$ of incomparable pairs defined in the dimension theory of partial orders. We show that $G_{\mathcal{P}}$ satisfies properties (i) and (ii) where $\Delta - 1$ is the maximum number of predecessors (or successors) of each job.

1 Introduction

Vertex Cover is one of the most studied problems in combinatorial optimization: Given a graph $G = (V, E)$ with weights w_i on the vertices, find a subset $V' \subseteq V$, minimizing the objective function $\sum_{i \in V'} w_i$, such that for each edge $\{u, v\} \in E$, at least one of u and v belongs to V' .

The related bibliography is vast and cannot be covered in this introductory note. We mention here that vertex cover cannot be approximated within a factor of 1.3606 [12], unless $P=NP$. Moreover, if the Unique Game Conjecture (UGC) [28] holds, Khot and Regev [29] show that vertex cover is hard to approximate within any constant factor better than 2. On the other side several simple 2-approximation algorithms are known (see e.g. [36, 25]). Hochbaum [25] uses the natural linear program (LP) relaxation and a threshold rounding approach to obtain better than 2 approximation algorithms when a k -coloring of the graph is given as input. An optimal solution to the LP assigns a non-negative real value to each vertex of the input graph G . It is well known that such a solution is half-integral [36]. Let S_1 be the vertices of the input graph which attain value 1 in this solution; and let S_2 be the vertices which attain value $1/2$. The vertices S_1 together with a cover of the subgraph induced by S_2 are sufficient to cover the whole graph G . A k -coloring of the subgraph induced by S_2 gives an independent set of value at least $w(S_2)/k$, where $w(S_2)$ is the sum of the vertex weights in S_2 . This yields a $(2 - 2/k)$ -approximation for the minimum weighted vertex cover problem. For graphs with degree bounded by d , this directly leads to a $(2 - 2/d)$ -approximation.

This basic approach has been considerably improved (for sufficiently large d) by Halperin [23]. The improvement is obtained by replacing the LP relaxation with a stronger semidefinite program (SDP) relaxation, and using a fundamental result by Karger, Motwani and Sudan [27]. The algorithm in [23] achieves a performance ratio of $2 - (1 - o(1)) \frac{2 \ln \ln d}{\ln d}$, which improves the previously known [22] ratio of $2 - \frac{\ln d + O(1)}{d}$. Under the UGC [28], Austrin, Khot and Safra [6] have recently proved that it is NP-hard to approximate vertex cover in bounded degree graphs to within a factor $2 - (1 + o(1)) \frac{2 \ln \ln d}{\ln d}$. This exactly matches the algorithmic result of Halperin [23] up to the $o(1)$ term. For general vertex cover, the currently best approximation ratio is due to Karakostas [26], achieving a performance of $2 - \Theta(1/\sqrt{\ln n})$.

The famous Brooks' theorem states that a graph in which every vertex has at most d neighbors, the vertices may be colored with only d colors, except for two cases, complete graphs and cycle graphs of odd length, which require $d + 1$ colors. In this paper we consider the vertex cover problem in graphs with bounded local chromatic number, a generalization of the bounded degree case with arbitrarily large chromatic number.

For an undirected graph $G = (V, E)$ and a vertex $u \in V$, we use $N(u) := \{v \in V : \{u, v\} \in E\}$ to denote the set of neighbors of u . Given a graph G , a valid vertex coloring of G is a function $\varphi : V \rightarrow \mathbb{N}$ such that $\varphi(u) \neq \varphi(v)$ whenever $\{u, v\} \in E$. In [14], Erdős et. al introduced the notion of local colorings.

Definition 1.1 (Local Coloring). Let k be a positive integer. A k -local coloring of a graph G is a valid vertex coloring φ such that for every vertex $u \in V$, $|\{\varphi(v) : v \in N(u)\}| < k$.

Note that the definition requires the number of colors in each neighborhood $N(u)$ to be strictly less than k . This guarantees that in each closed neighborhood $\{u\} \cup N(u)$, the number of different colors is bounded by k . The *local chromatic number* $\psi(G)$ of a graph G is the minimum k such that G admits a k -local coloring [14]. Since any valid coloring with k colors also is a local k -coloring, clearly, $\psi(G) \leq \chi(G)$. Interestingly, it can also be shown that the local chromatic number is always at least as large as the fractional chromatic number, i.e., $\psi(G) \geq \chi_f(G)$ [31]. For the definition and basic properties of the fractional chromatic number we refer to [38].

Given a valid vertex coloring φ and an integer i , let $C_i := \{v \in V : \varphi(v) = i\}$ be the set of vertices with color i . Further, for integers $i \neq j$, we use $N_j(C_i) := C_j \cap \bigcup_{v \in C_i} N(v)$ to denote the vertices with color j that have a neighbor with color i . We consider colorings with the following density condition.

Definition 1.2 (Biclique Coloring). A coloring φ of a graph G is called a biclique coloring if for any two colors i and j , the subgraph induced by $N_i(C_j)$ and $N_j(C_i)$ is either empty or a complete bipartite graph.

We will consider vertex cover in graphs for which a local biclique coloring φ is given. For any fixed $k \geq 3$, it is shown in [14] that there are n -vertex graphs with local chromatic number k and chromatic number $\Theta(\log \log n)$ and also that the gap between the local chromatic number and the chromatic number cannot be larger than this. This is shown using biclique colorings (cf. Def. 1.3 and Lemma 1.1 in [14]). Consequently, there are graphs that have chromatic number $\Theta(\log \log n)$ and admit a 3-local biclique coloring.

1.1 Our Contribution

In this paper we study the vertex cover problem in graphs with bounded local colorings. Assuming the UGC [28], the provided results give a tight characterization of the problem. The two main results are summarized as follows.

Theorem 1.1. *The vertex cover problem in graphs $G = (V, E)$ for which a $(\Delta+1)$ -local biclique coloring φ of G is given as input admits a randomized polynomial-time algorithm with approximation ratio $2 - \Omega(1) \frac{\ln \ln \Delta}{\ln \Delta}$.*

Theorem 1.2. *Assuming the UGC, it is NP-hard to approximate the vertex cover problem in graphs for which a $(\Delta + 1)$ -local coloring is given as input, within any constant factor better than $2 - 2/(\Delta + 1)$.*

The result stated in Theorem 1.1 matches (up to the constant factor in the lower order term) a known inapproximability result [6]. In Section 3 we provide a matching upper bound for the inapproximability result of Theorem 1.2.

Besides generalizing the bounded degree case to a class of graphs with arbitrarily large chromatic number, we show that Theorem 1.1 finds a natural application in a classical scheduling problem, namely the precedence constrained single machine scheduling problem to minimize the weighted sum of completion times, known as $1|prec|\sum w_j C_j$ in standard scheduling notation (see Section 1.3 for the definition and a review of the literature). In a series of papers [2, 10, 11] it was established that this scheduling problem is a special case of minimum weighted vertex cover in graphs $G_{\mathcal{P}}$ of incomparable pairs defined in the dimension theory of partial orders. We prove the following in this paper.

Theorem 1.3. *For any graph of incomparable pairs $G_{\mathcal{P}}$, a $(\Delta + 1)$ -local biclique coloring of $G_{\mathcal{P}}$ can be computed in polynomial time, where $\Delta - 1$ is the maximum number of predecessors (or successors) of each job.*

Together with Theorem 1.1, this result improves the previously best $(2 - \frac{2}{\max\{\Delta, 2\}})$ -approximation algorithm described in [3].

1.2 Review of the SDP Approach for Bounded Degree Graphs

In [27] the authors consider the problem of coloring k -colorable graphs as a function of the graph maximum degree using semidefinite programming. Given a graph $G = (V, E)$ on n vertices, and a real number $k \geq 2$, a *vector k -coloring* [27] of G is an assignment of unit vectors $v_i \in \mathbb{R}^n$ to each vertex $i \in V$, such that for any two adjacent vertices i and j the dot product of their vectors satisfies the inequality $v_i \cdot v_j \leq -1/(k - 1)$. They show that it is possible to check if a graph admits a vector k -coloring by using semidefinite programming. Moreover, they prove that vector k -colorable graphs have a “large” independent set when k is “small”.

Theorem 1.4 ([27]). *For every integral $k \geq 2$, a vector k -colorable graph $G = (V, E)$ with maximum degree d has an independent set I of value $\Omega(\frac{w(V)}{d^{1-2/k}\sqrt{\ln d}})$.*

The result in Theorem 1.4 has been obtained by generalizing the technique of *Rounding via Hyperplane Partitions* [19] to the technique of *Rounding via Vector Projections* [27]. More precisely, the independent set I claimed in Theorem 1.4 is obtained as follows. Suppose we have a vector k -coloring assigning unit vectors v_i to the vertices. Fix a parameter $c = c(k, d)$ and choose a random n dimensional

vector r from the n -dimensional standard normal distribution. Consider the set I' of vertices i with $v_i \cdot r \geq c$. The independent set I consists of all the isolated vertices in the subgraph induced by the vertex set I' . Intuitively, for small k the vector k -colorability implies that endpoints of an edge have vectors pointing away from each other. If the vector associated with a vertex has a large dot product with r , then the vector corresponding to an adjacent vertex will not have such a large dot product with r and hence will not be selected. Thus, only a relatively small fraction of the edges are likely to be in the induced subgraph on the selected set of vertices. The independent set I is large iff in the induced subgraph the expected number of edges is a small fraction of n . The latter is guaranteed by the rounding technique and analysis in [27] assuming that the graph has “few”, i.e., $nd/2$, edges (a property that does not hold for the problem addressed in the present paper).

In [30], it is proved that the following program is a semidefinite relaxation of the vertex cover problem. Moreover, it can be solved within an additive error of $\varepsilon > 0$ in polynomial time in $\ln \frac{1}{\varepsilon}$ and n using the ellipsoid method.

$$\begin{aligned} \min \quad & \sum_{u=1}^n w_u \frac{1+v_0 \cdot v_u}{2} \\ \text{s.t.} \quad & (v_i - v_0)(v_j - v_0) = 0, \quad \{i, j\} \in E \\ & \|v_u\| = 1, \quad u \in V \cup \{0\}, v_u \in \mathbb{R}^{n+1} \end{aligned} \quad (1)$$

Note that in an “integral” solution of (1) (corresponding to a vertex cover), vectors for vertices that are picked coincide with v_0 , while the other vectors coincide with $-v_0$. For general graphs, it is shown in [30] that the integrality gap is $2 - \varepsilon$, for any $\varepsilon > 0$, i.e., for every $\varepsilon > 0$ there is a graph G_ε such that $vc(G_\varepsilon)/sd(G_\varepsilon)$ is at least $2 - \varepsilon$, where $vc(G_\varepsilon)$ and $sd(G_\varepsilon)$ denote, respectively, the minimum vertex cover value and the optimum value of (1).

Halperin [23] uses (1) to provide an efficient randomized algorithm that approximates vertex cover in graphs with maximum degree d . The improvement is obtained as follows by using the following threshold rounding approach. Solve relaxation (1) and let

$$S_1 = \{u \in V \mid v_0 \cdot v_u \geq x\} \quad (2)$$

$$S_2 = \{u \in V \mid -x \leq v_0 \cdot v_u < x\}, \quad (3)$$

where x is a small positive number. As in Hochbaum’s approach [25], it holds that the vertices S_1 together with a cover of the subgraph induced by S_2 are sufficient to cover the whole graph G . Moreover, for any two adjacent vertices i and j in the graph $G[S_2]$ induced by S_2 , the edge constraint of (1) implies that $v_i \cdot v_j \leq -1 + 2x$. This has the important consequence that $G[S_2]$ is a vector k -colorable graph, where $k = \frac{2-2x}{1-2x}$ is close to 2 for small x . We can then use¹ Theorem 1.4 to obtain a large

¹In Theorem 1.4 the integrality assumption on k is not technically necessary, and it can be easily

valued independent set I of $G[S_2]$. The returned vertex cover is $S_1 \cup (S_2 \setminus I)$ and the result of [23] follows by choosing a suitable value for x .

1.3 Literature Review of the Scheduling Problem

The problem we consider in this paper is a classical problem in scheduling theory, known as $1|prec|\sum w_j C_j$ in standard scheduling notation (see e.g. Graham et al. [20]). It is defined as the problem of scheduling a set $N = \{1, \dots, n\}$ of n jobs on a single machine, which can process at most one job at a time. Each job j has a processing time p_j and a weight w_j , where p_j and w_j are nonnegative integers. Jobs also have precedence constraints between them that are specified in the form of a *partially ordered set (poset)* $\mathbf{P} = (N, P)$, consisting of the set of jobs N and a partial order, i.e., a reflexive, antisymmetric, and transitive binary relation P on N , where $(i, j) \in P$, whenever $i \neq j$, implies that job i must be completed before job j can be started. The goal is to find a non-preemptive schedule which minimizes $\sum_{j=1}^n w_j C_j$, where C_j is the time at which job j completes in the given schedule.

The described problem was shown to be strongly NP-hard already in 1978 [32, 34] by Lawler [32] and Lenstra and Rinnooy Kan [34]. For the general version of $1|prec|\sum w_j C_j$, several polynomial time 2-approximation algorithms are known [39, 21, 10, 9, 35]. Until recently, no inapproximability results were known, and closing the approximability gap has been listed as one of ten outstanding open problems in scheduling theory (e.g., [40]). Ambühl, Mastrolilli and Svensson [5] proved that the problem does not admit a PTAS, assuming that NP-complete problems cannot be solved in randomized subexponential time. Moreover, if a fixed cost present in all feasible schedules is ignored then the problem is as hard to approximate as vertex cover [5]. Recently, Bansal and Khot [7] showed that the gap for the general problem indeed closes assuming a variant of the UGC [28], by providing a 2-inapproximability result based on that assumption.

In a series of papers [2, 10, 11] it was proved that $1|prec|\sum w_j C_j$ is a special case of minimum weighted vertex cover in some special graphs $G_{\mathbf{P}}$ that depend on the input poset \mathbf{P} . More precisely, it is shown that any feasible solution to the vertex cover problem in graphs $G_{\mathbf{P}}$ can be turned in polynomial time into a feasible solution to $1|prec|\sum w_j C_j$ without deteriorating the objective value. This result was achieved by investigating different integer LP formulations and relaxations [37, 10, 11] of $1|prec|\sum w_j C_j$, using linear ordering variables δ_{ij} such that the variable δ_{ij} has value 1 if job i precedes job j in the corresponding schedule, and 0 otherwise.

Dushnik and Miller [13] introduced dimension as a parameter of partial orders in 1941. The dimension is one of the most heavily studied parameters of partial

generalized to fractional k . As remarked in [23], by using exactly the same analysis and rounding technique as in [27], it is possible to compute an independent set of value at least $\Omega(\frac{w(S_2)}{d^{x/(1-x)} \sqrt{x \ln d}})$ for $k = \frac{2-2x}{1-2x}$.

orders, and many beautiful results have been obtained (see e.g. [41]). There is a natural way to associate with a poset \mathbf{P} a hypergraph $\mathbf{H}_{\mathbf{P}}$, called the *hypergraph of incomparable pairs*, so that the dimension of \mathbf{P} is the chromatic number of $\mathbf{H}_{\mathbf{P}}$ [17]. Furthermore, the fractional dimension of \mathbf{P} , a generalization due to Brightwell and Scheinerman [8] is equal to the fractional chromatic number of $\mathbf{H}_{\mathbf{P}}$. It turns out [4] that graph $G_{\mathbf{P}}$ is the (ordinary) graph obtained by removing from $\mathbf{H}_{\mathbf{P}}$ all edges of cardinality larger than two. This allows to apply the rich vertex cover theory to $1|prec|\sum w_j C_j$ together with the dimension theory of partial orders. One can, e.g., conclude that the scheduling problem with two-dimensional precedence constraints is solvable in polynomial time, as $G_{\mathbf{P}}$ is bipartite in this case [17, 11], and the vertex cover problem is well-known to be solvable in polynomial time on bipartite graphs. This considerably extends Lawler’s result [32] from 1978 for series-parallel precedence constraints. Further, these connections between the $1|prec|\sum w_j C_j$ and the vertex cover problem on $G_{\mathbf{P}}$, and between dimension and coloring, yield a framework for obtaining $(2 - 2/f)$ -approximation algorithms for classes of precedence constraints with bounded (fractional) dimension f [3, 2]. The framework is inspired by Hochbaum’s approach [25] for the vertex cover problem on “easily” colorable graphs. It yields the best known approximation ratios for all previously considered special classes of precedence constraints, like semi-orders, convex bipartite orders, interval orders, interval dimension 2, bounded in-degree posets.

2 Vertex Cover Using Bounded Local Biclique Colorings

In this section we provide and analyze an approximation algorithm for the vertex cover problem in graphs $G = (V, E)$ for which a $(\Delta + 1)$ -local biclique coloring φ of G is given.

The presented approximation algorithm follows the threshold rounding approach used for the bounded-degree vertex cover problem [23]: first solve the SDP (1) and, based on this solution and a parameter x that will be determined later, two vertex sets, S_1 and S_2 , are computed as follows.

$$S_1 = \{u \in V \mid v_0 \cdot v_u \geq x\} \tag{4}$$

$$S_2 = \{u \in V \mid -x \leq v_0 \cdot v_u < x, \} \tag{5}$$

The cover is obtained by picking the vertices in S_1 together with a cover of the subgraph induced by S_2 .

However, unlike in [23], Theorem 1.4 does not apply to graph $G[S_2]$. Indeed, the graph degree is not bounded and the theorem does not generalize to these graphs since the rounding procedure and analysis in [27] are strongly based on the assumption that the graph has “few”, i.e. $O(n)$ edges.

We show that graph $G[S_2]$ has a “large” independent set by using a new rounding procedure to compute it that works as follows. The vertices in S_2 are first grouped into overlapping *clusters*. For every two colors i and j (of the given coloring φ) such that the number of edges connecting vertices with the two colors is non-zero, there are two clusters $N_i(C_j)$ and $N_j(C_i)$. Note that $N_j(C_i) \neq \emptyset$ iff $N_i(C_j) \neq \emptyset$ in $G[S_2]$. Because we assume that the coloring φ is a local $(\Delta + 1)$ -coloring, each vertex $u \in V$ belongs to at most Δ clusters. Furthermore, for all i and j , clusters $N_i(C_j)$ and $N_j(C_i)$ are connected by complete bipartite sub-graphs (note that this property also holds when restricting the graph G to the vertex set S_2). It can be easily proved, that all vectors v_i corresponding to vertices from the same cluster almost point in the same direction. This essentially follows from having a biclique coloring and because, by the definition of the set S_2 , vectors corresponding to adjacent vertices in S_2 are almost antipodal. For each cluster $N_i(C_j)$, we arbitrarily choose one *representative* vertex. Let R be the set of representatives of all clusters $N_i(C_j)$. Further, for a vertex u , let R_u be the set of representatives of clusters $N_i(C_j)$ for which $u \in N_i(C_j)$. By only using the vectors of the representatives, we compute a subset $I' \subseteq S_2$ as follows: vertex $u \in S_2$ belongs to I' if and only if every representative $a \in R_u$ of u satisfies $v_a \cdot r \geq c$, where c is a parameter and r is a random $(n + 1)$ -dimensional vector r from the $(n + 1)$ -dimensional standard normal distribution, i.e., the components of r are independent Gaussian random variables with mean 0 and variance 1. We show that the set I' has a large independent set. Formally, we have

$$I' = \left\{ u \in S_2 : \bigwedge_{a \in R_u} v_a \cdot r \geq c \right\}. \quad (6)$$

The complete procedure is summarized in Algorithm 1.

2.1 Analysis

We first show that for all $u \in S_2$, the vectors corresponding to representatives of u 's clusters point in almost the same direction as the vector v_u corresponding to vertex u .

Lemma 2.1. *Let $u \in S_2$ be a vertex, $a \in R_u$ be the representative vertex of any cluster to which u belongs, and $x \in (0, 1)$. We have $v_u \cdot v_a \geq 1 - 8x + o(x)$.*

Proof. See Appendix A.1. □

The expected size of the set I' is $\sum_{u \in S_2} \Pr[u \in I']$. Using (6), we have

$$\Pr[u \in I'] = \Pr \left[\bigwedge_{a \in R_u} v_a \cdot r \geq c \right]. \quad (7)$$

1. Solve SDP (1)
2. Let $S_1 = \{u \in [n] \mid v_0 \cdot v_u \geq x\}$, $S_2 = \{u \in [n] \mid -x \leq v_0 \cdot v_u < x\}$.
3. Find an IS I in $G[S_2]$ as follows:
 - (a) Compute the set R of representatives
 - (b) Let $R_u = \{a \in R : \{u, a\} \subseteq N_i(C_j)\}$
for any two colors i and j of the given coloring φ
 - (c) Choose a random vector r and define

$$I' = \{u \in S_2 \mid \bigwedge_{a \in R_u} v_a \cdot r \geq c\}$$

- (d) Get I by removing one vertex of every edge in $G[I']$ from I'
4. Output the constructed vertex cover $S_1 \cup (S_2 \setminus I)$

Algorithm 1: Vertex Cover Approximation Algorithm

In the following we compute a lower bound on the above probability. The subsequent analysis uses some basic properties of the normal distribution. Let X be a standard normal random variable (X has mean 0 and variance 1). We use $\mathcal{N}(x)$ to denote the probability that X is at least x . Hence, $\mathcal{N}(x) := \Pr(X \geq x) = \int_x^\infty \phi(t) dt$, where $\phi(t) = e^{-t^2/2}/\sqrt{2\pi}$ is the density of X (see e.g. [15]).

Lemma 2.2. *For any $u \in S_2$ and any constant $\gamma > 0$*

$$\Pr \left[\bigwedge_{a \in R_u} v_a \cdot r \geq c \right] \geq \mathcal{N}(\alpha c) - \Delta \cdot \mathcal{N} \left(\frac{c}{\beta \sqrt{x}} \right),$$

where $\alpha = \gamma + 1 + o(1)$ and $\beta = \frac{4+o(1)}{\gamma}$.

Proof. Assume, without loss of generality, that $v_u = (1, 0, \dots, 0)$. By Lemma 2.1, for every $a \in R_u$ $v_u \cdot v_a \geq 1 - \delta$, where $\delta = 8x + o(x)$. Then the first component $v_{a,1}$ of $v_a = (v_{a,1}, v_{a,2}, \dots, v_{a,n+1})$ must be at least $1 - \delta$. Let $A = (1 - \delta, 0, \dots, 0)$ and $B_a = (0, v_{a,2}, \dots, v_{a,n+1})$ for any $a \in R_u$.

Since v_a is a unit length vector, we have that $(1 - \delta)^2 + \sum_{j=2}^{n+1} v_{a,j}^2 \leq 1$, which gives an upper bound on the length of B_a :

$$\|B_a\|^2 = \sum_{j=2}^{n+1} v_{a,j}^2 \leq 2\delta - \delta^2 < 2\delta.$$

We then get

$$\begin{aligned}
\Pr \left[\bigwedge_{a \in R_u} v_a \cdot r \geq c \right] &\geq \Pr \left[\bigwedge_{a \in R_u} (A \cdot r \geq (\gamma + 1)c \wedge B_a \cdot r \geq -\gamma c) \right] \\
&= \Pr \left[A \cdot r \geq (\gamma + 1)c \wedge \bigwedge_{a \in R_u} B_a \cdot r \geq -\gamma c \right] \\
&= 1 - \Pr \left[A \cdot r < (\gamma + 1)c \vee \bigvee_{a \in R_u} B_a \cdot r < -\gamma c \right] \\
&\geq 1 - \Pr [A \cdot r < (\gamma + 1)c] - \sum_{a \in R_u} \Pr [B_a \cdot r < -\gamma c] \\
&= \Pr \left[\frac{A}{\|A\|} \cdot r \geq \frac{(\gamma + 1)c}{\|A\|} \right] - \sum_{a \in R_u} \Pr \left[-\frac{B_a}{\|B_a\|} \cdot r > \frac{\gamma c}{\|B_a\|} \right] \\
&\geq \mathcal{N} \left(\frac{(\gamma + 1)c}{1 - \delta} \right) - \Delta \cdot \mathcal{N} \left(\frac{\gamma c}{\sqrt{2\delta}} \right) \\
&= \mathcal{N}(\alpha c) - \Delta \cdot \mathcal{N} \left(\frac{c}{\beta \sqrt{x}} \right).
\end{aligned}$$

The last inequality follows because the sum of k independent Gaussian random variables with variances $\sigma_1^2, \dots, \sigma_k^2$ is a Gaussian random variable with variance $\sum_{i=1}^k \sigma_i^2$ (see e.g. [15]). Consequently, the the dot product of a unit vector with r is a standard normal random variable. \square

The total weight of the vertices that are removed from I' is upper bounded by the weight of vertices of the edges in the graph $G[I']$ induced by I' . The next lemma bounds the probability that a vertex $u \in S_2$ is in I' and that u has a neighbor $u' \in S_2$ that is also in I' .

Lemma 2.3. *Consider a vertex $u \in S_2$. The probability that u as well as some neighbor u' of u are in I' is upper bound by*

$$\Pr [u \in I' \wedge \exists u' \in S_2 : \{u, u'\} \in E \wedge u' \in I'] \leq \Delta \cdot \mathcal{N} \left(\frac{c}{\sqrt{x}} \right).$$

Proof. Recall that R_u is the set of representative vertices for u 's clusters. Because we assume that we are given a local $(\Delta + 1)$ -coloring, u can only have neighbors in at most Δ other clusters. Actually, for each cluster $N_i(C_j)$, u belongs to, u has at least one neighbor in cluster $N_j(C_i)$ and these are the only neighboring clusters.

Let Q_u be the representative vertices of the at most Δ neighbor clusters of u . Further, let A be the event that $u \in I'$ and let B be the event that some neighbor u' of u is in I' . Given some vertex $b \in Q_u$, we define $a_b \in R_u$ such that $\{b, a_b\} \in E$. Such a vertex exists because if b is the representative of a cluster $N_i(C_j)$, we have $u \in N_j(C_i)$ and therefore we can choose a_b to be the representative of cluster $N_j(C_i)$. We have

$$\begin{aligned}
\Pr[A \cap B] &\leq \Pr \left[\left(\bigwedge_{a \in R_u} v_a \cdot r \geq c \right) \wedge \left(\bigvee_{b \in Q_u} v_b \cdot r \geq c \right) \right] \\
&\leq \sum_{b \in Q_u} \Pr [(v_{a_b} \cdot r \geq c) \wedge (v_b \cdot r \geq c)] \\
&\leq \sum_{b \in Q_u} \Pr \left[\frac{v_{a_b} + v_b}{\|v_{a_b} + v_b\|} \cdot r \geq \frac{2c}{\|v_{a_b} + v_b\|} \right] \\
&\leq \Delta \cdot \mathcal{N} \left(\frac{c}{\sqrt{x}} \right).
\end{aligned}$$

The last inequality follows because the definition of SDP (1), $a_b, b \in S_2$, and $\{a_b, b\} \in E$ together imply that $\|v_{a_b} + v_b\| \leq 2\sqrt{x}$ (as observed in [23]). \square

Based on Lemmas 2.2 and 2.3, we can now lower bound the expected size of the computed independent set I and we can thus obtain a bound on the expected approximation ratio of Algorithm 1.

Theorem 2.4. *Choosing $x = \frac{1-o(1)}{25} \cdot \frac{\ln \ln \Delta}{\ln \Delta}$, $c = (1 + o(1)) \cdot \sqrt{\frac{2x}{1-25x}} \ln \Delta$, and $\gamma = 4$, Algorithm 1 has an expected approximation ratio of $2 - \frac{2-o(1)}{25} \cdot \frac{\ln \ln \Delta}{\ln \Delta}$.*

Proof. See Appendix A.2. \square

Remark: Repeating the construction of the independent set I (i.e., choosing the random vector r) sufficiently often, we can obtain the same approximation ratio (up to a slightly worse $o(1)$ term) with high probability.

3 Vertex Cover in Graphs with Bounded Local Chromatic Number

Consider the vertex cover problem in graphs G_Δ for which a local $(\Delta+1)$ -coloring is given. Theorem 3.1 shows that the approximation ratio achievable from relaxation (1) is no better than $2 - \frac{2}{\Delta+1}$ if only the bounded local colorability, but not the biclique condition holds.

Theorem 3.1. *For any fixed $\Delta \geq 2$ and $\varepsilon > 0$, there is a local $(\Delta + 1)$ -colorable graph $G_{\Delta, \varepsilon}$ for which*

$$\frac{vc(G_{\Delta, \varepsilon})}{sd(G_{\Delta, \varepsilon})} \geq 2 - \frac{2}{\Delta + 1} - \varepsilon,$$

where $vc(G)$ and $sd(G)$ denote the size of a minimum weighted vertex cover of G and the solution value for the corresponding SDP relaxation (1), respectively.

Proof. See Appendix A.3. □

Under the Unique Game Conjecture, Khot and Regev [29] proved that vertex cover is NP-hard to approximate better than $2 - \varepsilon$, for any $\varepsilon > 0$.

Theorem 3.2 ([29]). *Assuming the Unique Game Conjecture, for arbitrarily small constants $\varepsilon, \delta > 0$, there is a polynomial time reduction mapping a SAT formula ϕ to an n -vertex graph G such that if ϕ is satisfiable then G has an independent set of size $(\frac{1}{2} - \varepsilon)n$ and if ϕ is unsatisfiable, then G has no independent set of size δn .*

By using Theorem 3.2, the reduction and the analysis in the proof of Theorem 3.1 can be easily adapted to obtain the following conditional hardness.

Theorem 3.3. *Assuming the Unique Game Conjecture, it is NP-hard to approximate the vertex cover problem in graphs for which a local $(\Delta + 1)$ -coloring is given as input, within any constant factor better than $2 - 2/(\Delta + 1)$.*

A matching upper bound can be obtained by showing that an independent set of value at least $\frac{\sum_{i=1}^n w_i}{\Delta + 1}$ is computable in polynomial time. Indeed, the following result shows that Turan's theorem (as proved by Caro and Wei, see e.g. [1]) for bounded degree graphs can be generalized to graphs with bounded local colorings.

Theorem 3.4. *For any weighted graph $G = (V, E)$, there exists an independent set of value at least $\sum_{v \in V} \frac{w_v}{\Delta_v + 1}$, where $\Delta_v = |\{c(u) : u \in N(v)\}|$ is the number of different colors in the neighborhood of $v \in V$ of any proper coloring c of G .*

Proof. Consider a uniformly chosen total ordering $<$ of the color classes $\{C_1, \dots, C_k\}$, where $C_i = \{v \in V : c(v) = i\}$. Define $I := \{v \in V : \{v, w\} \in E \wedge v \in C_i \wedge w \in C_j \rightarrow C_i < C_j\}$. Let X_v be the indicator random variable for $v \in I$ and $X = \sum_{v \in V} X_v = |I|$. For each v , $E[X_v] = Pr[v \in I] = \frac{1}{\Delta_v + 1}$. Hence, there exists a specific ordering of the classes so that the corresponding independent set I has value at least $\sum_{v \in V} \frac{w_v}{\Delta_v + 1}$. □

By using standard techniques, an upper bound that matches the lower bound of Theorem 1.2 can be obtained in deterministic polynomial time.

Corollary 3.5. *There exists a $(2 - \frac{2}{\Delta + 1})$ -approximation algorithm for the vertex cover problem in graphs for which a local $(\Delta + 1)$ -coloring is given as input.*

4 The Scheduling Application

Problem $1|prec|\sum w_j C_j$ is a classical and fundamental problem in scheduling theory [33, 40]. Its complexity certainly depends on the poset complexity: indeed, it can be efficiently solved when there are no precedence constraints or when the precedence constraints are not “very complicated”, namely when the dimension of the poset is at most two. More generally, the dimension of the input poset has been established to be an important parameter for the approximability of the problem [3, 2], with the lower the (fractional) dimension the better is the approximation ratio. Unfortunately, in the general case, recognizing the (fractional) dimension of a poset is hard even to approximate [24].

Another natural parameter of partial orders is given by the poset in- or out-degree [16], namely the job maximum number of predecessors or successors, respectively. One of the first NP-complete proofs [34] for $1|prec|\sum w_j C_j$ shows that the problem remains strongly NP-hard even if every job has at most two predecessors (or successors) in the poset. In [3], the authors present a $(2 - \frac{2}{\max\{\Delta, 2\}})$ -approximation algorithm, where $\Delta - 1$ is the minimum between the in- and the out-degree of the input poset. This gives a “good” approximation algorithm when the poset has a “small” Δ .

In this section we show two important properties of $G_{\mathbf{P}}$: (i) graph $G_{\mathbf{P}}$ can be efficiently colored so that no vertex neighborhood contains more than Δ different colors (*bounded local colorability*); and (ii) the coloring is a *biclique coloring*. This result together with Theorem 1.1 improves the previously known $(2 - 2/\max\{\Delta, 2\})$ -approximation algorithm in [3].

A *partially ordered set (poset)* $\mathbf{P} = (N, P)$, consists of a set N and a partial order P on N , i.e., a reflexive, antisymmetric, and transitive binary relation P on N . For $x, y \in N$, we write $x \leq y$ when $(x, y) \in P$, and $x < y$ when $(x, y) \in P$ and $x \neq y$. When neither $(x, y) \in P$ nor $(y, x) \in P$, we say that x and y are incomparable, denoted by $x \parallel y$. We call $\text{inc}(\mathbf{P}) = \{(x, y) \in N \times N : x \parallel y \text{ in } P\}$ the set of *incomparable pairs* of \mathbf{P} .

For any integer $k \geq 2$, a subset $S = \{(x_i, y_i) : 1 \leq i \leq k\} \subset \text{inc}(\mathbf{P})$ is called an *alternating cycle* when $x_i \leq y_{i+1}$ in P , for all $i = 1, 2, \dots, k$, and where $y_{k+1} = y_1$. An alternating cycle $S = \{(x_i, y_i) : 1 \leq i \leq k\}$ is *strict* if $x_i \leq y_j$ in P if and only if $j = i + 1$, for all $i, j = 1, 2, \dots, k$.

Definition 4.1 (Hypergraph and Graph of Incomparable Pairs [17]). For any give poset \mathbf{P} , the *hypergraph of incomparable pairs of \mathbf{P}* , denoted $\mathbf{H}_{\mathbf{P}} = (V, E)$, is the hypergraph that satisfies the following conditions:

1. The vertex set V is the set $\text{inc}(\mathbf{P})$ of incomparable pairs of \mathbf{P} .
2. The edge set E consists of those subsets of V that form strict alternating cycles.

The graph $G_{\mathbf{P}}$ of incomparable pairs of \mathbf{P} is the ordinary graph determined by all edges of size 2 in $\mathbf{H}_{\mathbf{P}}$.

Hence, in $G_{\mathbf{P}}$, there is an edge between incomparable pairs (i, j) and (k, ℓ) if and only if $(i, \ell), (k, j) \in P$.

In a series of papers [2, 11, 10], it was proved that $1|prec| \sum w_j C_j$ is equivalent to a weighted vertex cover problem on the graph of incomparable pairs $G_{\mathbf{P}}$ of the poset \mathbf{P} characterizing the precedence constraints of the scheduling problem. More precisely, given a scheduling instance S with precedence constraints \mathbf{P} , let $G_{\mathbf{P}}^S$ be the *weighted* version of graph $G_{\mathbf{P}}$ with the following weights. For all incomparable pairs $(i, j) \in \text{inc}(\mathbf{P})$, the weight of vertex (i, j) in $G_{\mathbf{P}}$ is $p_i \cdot w_j$, where p_i is the processing time of job i and w_j is the weight of process j .

Theorem 4.1 ([2, 11, 10]). *For any given instance S of problem $1|prec| \sum w_j C_j$ with precedence constraints \mathbf{P} , an α -approximate solution to the weighted vertex cover in graphs $G_{\mathbf{P}}^S$ can be turned in polynomial time into an α -approximate solution to S .*

Let $\mathbf{P} = (N, P)$ be a poset. For any $j \in N$, define the *degree of j* , denoted $\text{deg}(j)$, as the number of elements comparable (but not equal) to j in \mathbf{P} . Given $j \in N$, let $D(j)$ denote the set of all elements which are less than j , and $U(j)$ those which are greater than j in P . Let $\text{deg}_D(j) := |D(j)|$ be the *in-degree* of j and the *maximum in-degree* $\Delta_D(\mathbf{P}) := \max\{\text{deg}_D(j) : j \in N\}$. The *out-degree* of j $\text{deg}_U(j)$ and the *maximum out-degree* $\Delta_U(\mathbf{P})$ are defined analogously (see also [16]). The maximum vertex degree in the graph of incomparable pairs $G_{\mathbf{P}}$ is bounded by $(\Delta_D(\mathbf{P}) + 1) \cdot (\Delta_U(\mathbf{P}) + 1)$. Hence, if both $\Delta_D(\mathbf{P})$ and $\Delta_U(\mathbf{P})$ are bounded, $G_{\mathbf{P}}$ has bounded degree and therefore, the bounded degree vertex cover approximation of [23] can be used to approximate the scheduling problem $1|prec| \sum w_j C_j$ with precedence constraints \mathbf{P} . If only either the in-degree or the out-degree of \mathbf{P} is bounded, $G_{\mathbf{P}}$ does not have bounded degree. However, we will now show that in this case $G_{\mathbf{P}}$ has a good local biclique coloring.

Theorem 4.2. *Let $\mathbf{P} = (N, P)$ be a any poset and let $\Delta = 1 + \min\{\Delta_D(\mathbf{P}), \Delta_U(\mathbf{P})\}$. Then, we can efficiently compute a $(\Delta + 1)$ -local biclique coloring of $G_{\mathbf{P}} = (V, E)$.*

Proof. We assume that $\Delta - 1 = \Delta_D(\mathbf{P})$ is the largest in-degree. The case $\Delta - 1 = \Delta_U(\mathbf{P})$ can be proven analogously. We first show how to compute a $(\Delta + 1)$ -local coloring of $G_{\mathbf{P}}$ (cf. Def. 1.1). Partition the incomparable pairs into $|N|$ color classes: $C_i = \{(i, j) \in \text{inc}(\mathbf{P})\}$ for $i \in [|N|]$. It is easy to check that every C_i forms an independent set. Moreover any incomparable pair $(i, j) \in \text{inc}(\mathbf{P})$ is adjacent to $(k, \ell) \in \text{inc}(\mathbf{P})$ if (necessary condition) $(k, j) \in P$. Since the in-degree of j is bounded by $\Delta - 1$, it follows that the number of distinct pairs (k, j)

such that $(k, j) \in P$ is bounded by Δ (it is Δ and not $\Delta - 1$ because we also have to consider the pair $(j, j) \in P$). Therefore any incomparable pair (i, j) has neighbors in at most Δ clusters and thus the coloring is $(\Delta + 1)$ -local.

In order to show that the coloring is a biclique coloring (cf. Def. 1.2), assume that $\{(i, a), (j, c)\} \in E$ and $\{(i, b), (j, d)\} \in E$. The claim follows by proving that $\{(i, a), (j, d)\} \in E$ and $\{(i, b), (j, c)\} \in E$. By the assumption we have: $(i, c) \in P$, $(j, a) \in P$, $(i, d) \in P$ and $(j, b) \in P$. Since $(j, a) \in P \wedge (i, d) \in P$ we have $\{(i, a), (j, d)\} \in E$, and $(i, c) \in P \wedge (j, b) \in P$ implies $\{(i, b), (j, c)\} \in E$. \square

5 Final Remarks

This paper is about one of the very classic problem in combinatorial optimization. More precisely, we consider the minimum weighted vertex cover problem in graphs with bounded local chromatic number, a natural generalization of the bounded degree case with arbitrarily large chromatic number that naturally applies to scheduling with precedence constraints. Assuming the Unique Game Conjecture, a tight characterization of the weighted vertex cover problem in graphs with bounded local colorings is provided. We show that that one needs precisely the two defined properties, namely bounded local colorability and biclique coloring, in order to obtain Halperin’s bound.

Acknowledgments. The second author is indebted to Nikos Mutsanas and George Karakostas for many interesting discussions on the subject. We would also like to thank Jiří Sgall for suggesting the term “biclique coloring.” The research is supported by the Swiss National Science Foundation project 200020-122110/1 “Approximation Algorithms for Machine Scheduling Through Theory and Experiments III” and by Hasler Foundation Grant 11099.

References

- [1] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley, New York, 1992.
- [2] C. Ambühl and M. Mastrolilli. Single machine precedence constrained scheduling is a vertex cover problem. *Algorithmica*, 53(4):488–503, 2009.
- [3] C. Ambühl, M. Mastrolilli, N. Mutsanas, and O. Svensson. Scheduling with precedence constraints of low fractional dimension. In *IPCO*, pages 130–144, 2007.
- [4] C. Ambühl, M. Mastrolilli, and O. Svensson. Approximating precedence-constrained single machine scheduling by coloring. In *Proceedings of APPROX + RANDOM*, volume 4110 of *Lecture Notes in Computer Science*, pages 15–26. Springer, 2006.

- [5] C. Ambühl, M. Mastrolilli, and O. Svensson. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. In *FOCS*, pages 329–337, 2007.
- [6] P. Austrin, S. Khot, and M. Safra. Inapproximability of vertex cover and independent set in bounded degree graphs. In *IEEE Conference on Computational Complexity*, pages 74–80, 2009.
- [7] N. Bansal and S. Khot. Optimal Long-Code test with one free bit. In *Foundations of Computer Science (FOCS)*, pages 453–462, 2009.
- [8] G. R. Brightwell and E. R. Scheinerman. Fractional dimension of partial orders. *Order*, 9:139–158, 1992.
- [9] C. Chekuri and R. Motwani. Precedence constrained scheduling to minimize sum of weighted completion times on a single machine. *Discrete Applied Mathematics*, 98(1-2):29–38, 1999.
- [10] F. A. Chudak and D. S. Hochbaum. A half-integral linear programming relaxation for scheduling precedence-constrained jobs on a single machine. *Operations Research Letters*, 25:199–204, 1999.
- [11] J. R. Correa and A. S. Schulz. Single machine scheduling with precedence constraints. *Mathematics of Operations Research*, 30(4):1005–1021, 2005.
- [12] I. Dinur and S. Safra. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162(1):439–485, 2005.
- [13] B. Dushnik and E. Miller. Partially ordered sets. *American Journal of Mathematics*, 63:600–610, 1941.
- [14] P. Erdős, Z. Füredi, A. Hajnal, P. Komjáth, V. Rödl, and Á. Seress. Coloring graphs with locally few colors. *Discrete Mathematics*, 59(1-2):21–34, 1986.
- [15] W. Feller. *An Introduction to Probability Theory and Its Applications*. Wiley, 1968.
- [16] Felsner and Trotter. On the fractional dimension of partially ordered sets. *DMATH: Discrete Mathematics*, 136:101–117, 1994.
- [17] S. Felsner and W. T. Trotter. Dimension, graph and hypergraph coloring. *Order*, 17(2):167–177, 2000.
- [18] P. Frankl and V. Rödl. Forbidden intersections. *Transactions of the American Mathematical Society*, 300(1):259 – 286, 1987.

- [19] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995.
- [20] R. Graham, E. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling: A survey. In *Annals of Discrete Mathematics*, volume 5, pages 287–326. North–Holland, 1979.
- [21] L. A. Hall, A. S. Schulz, D. B. Shmoys, and J. Wein. Scheduling to minimize average completion time: off-line and on-line algorithms. *Mathematics of Operations Research*, 22:513–544, 1997.
- [22] M. M. Halldórsson and J. Radhakrishnan. Greed is good: Approximating independent sets in sparse and bounded-degree graphs. *Algorithmica*, 18(1):145–163, 1997.
- [23] E. Halperin. Improved approximation algorithms for the vertex cover problem in graphs and hypergraphs. *SIAM J. Comput.*, 31(5):1608–1623, 2002.
- [24] R. Hegde and K. Jain. The hardness of approximating poset dimension. *Electronic Notes in Discrete Mathematics*, 29:435–443, 2007.
- [25] D. S. Hochbaum. Efficient bounds for the stable set, vertex cover and set packing problems. *Discrete Applied Mathematics*, 6:243–254, 1983.
- [26] G. Karakostas. A better approximation ratio for the vertex cover problem. *ACM Transactions on Algorithms*, 5(4), 2009.
- [27] D. R. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2):246–265, 1998.
- [28] S. Khot. On the power of unique 2-prover 1-round games. In *STOC*, pages 767–775, 2002.
- [29] S. Khot and O. Regev. Vertex cover might be hard to approximate to within 2-epsilon. *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.
- [30] J. M. Kleinberg and M. X. Goemans. The lovász theta function and a semidefinite programming relaxation of vertex cover. *SIAM J. Discrete Math.*, 11(2):196–204, 1998.
- [31] J. Körner, C. Pilotto, and G. Simonyi. Local chromatic number and Sperner capacity. *Journal on Combinatorial Theory, Series B*, 95(1):101–117, 2005.

- [32] E. L. Lawler. Sequencing jobs to minimize total weighted completion time subject to precedence constraints. *Annals of Discrete Mathematics*, 2:75–90, 1978.
- [33] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys. Sequencing and scheduling: Algorithms and complexity. In S. C. Graves, A. H. G. Rinnooy Kan, and P. Zipkin, editors, *Handbooks in Operations Research and Management Science*, volume 4, pages 445–552. North-Holland, 1993.
- [34] J. K. Lenstra and A. H. G. Rinnooy Kan. The complexity of scheduling under precedence constraints. *Operations Research*, 26:22–35, 1978.
- [35] F. Margot, M. Queyranne, and Y. Wang. Decompositions, network flows and a precedence constrained single machine scheduling problem. *Operations Research*, 51(6):981–992, 2003.
- [36] G. L. Nemhauser and L. E. Trotter. Vertex packings: Structural properties and algorithms. *Mathematical Programming*, 8:232–248, 1975.
- [37] C. N. Potts. An algorithm for the single machine sequencing problem with precedence constraints. *Mathematical Programming Study*, 13:78–87, 1980.
- [38] E. R. Scheinerman and D. H. Ullman. *Fractional Graph Theory*. John Wiley and Sons Inc., 1997.
- [39] A. S. Schulz. Scheduling to minimize total weighted completion time: Performance guarantees of LP-based heuristics and lower bounds. In *Proceedings of the 5th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 301–315, 1996.
- [40] P. Schuurman and G. J. Woeginger. Polynomial time approximation algorithms for machine scheduling: ten open problems. *Journal of Scheduling*, 2(5):203–213, 1999.
- [41] W. T. Trotter. *Combinatorics and Partially Ordered Sets: Dimension Theory*. Johns Hopkins Series in the Mathematical Sciences. The Johns Hopkins University Press, 1992.

Appendix

A Deferred Proofs

A.1 Proof of Lemma 2.1

From the local completeness condition, it follows that there is a vertex $b \in S_2 \cap N_i(C_j)$ that is a common neighbor of u and a . From the definition of the SDP (1) and the definition of S_2 in (5), we have $v_a \cdot v_b \geq -1 + 2x$ and $v_b \cdot v_u \geq -1 + 2x$. Hence, the angle between v_a and v_b and the angle between v_b and v_u are both not smaller than $\theta := \arccos(-1 + 2x)$. The angle between v_a and v_u is therefore at most $2\pi - 2\theta$, that is, $\arccos(v_u \cdot v_a) \leq 2\pi - 2\theta$. The claim now follows by observing that $v_u \cdot v_a \geq \cos(2\pi - 2\theta) = \cos(2\theta) = 1 - 8x - o(x)$ by using Taylor expansion. \square

A.2 Proof of Theorem 2.4

For a set of vertices $U \subseteq V$, let $w(U)$ be the sum of the weights of the vertices in U . We start the proof by lower bounding the expected size of the computed independent set I of $G[S_2]$. By linearity of expectation and Lemmas 2.2 and 2.3, the expected size of I can be bounded as

$$\mathbb{E}[w(I)] \geq w(S_2) \cdot \left(\mathcal{N}(\alpha c) - \Delta \cdot \mathcal{N}\left(\frac{c}{\beta\sqrt{x}}\right) - \Delta \cdot \mathcal{N}\left(\frac{c}{\sqrt{x}}\right) \right), \quad (8)$$

where $\alpha = \gamma + 1 + o(1)$, $\beta = (4 + o(1))/\gamma$ and $\gamma > 0$ is some constant. We choose $\gamma = 4$, yielding $\alpha = 5 + o(1)$ and $\beta = 1 + o(1)$. Using Lemma B.1 and the choice of x and c , the right-hand side of (8) can then be lower bounded by $w(S_2) \cdot \mu x$, for any constant $\mu > 0$.

For a set $U \subseteq V$ of vertices, we define $sd(U) := \sum_{u \in U} w_u \cdot \frac{1 + v_0 \cdot v_u}{2}$ to be the contribution of U to the objective value of SDP (1). For $u \in S_1$, we have $v_0 \cdot v_u \geq x$ and thus, following the analysis in [23],

$$sd(S_1) \geq \sum_{u \in S_1} w_u \frac{1 + x}{2}. \quad (9)$$

Similarly, for $u \in S_2$, it holds that $v_0 \cdot v_u \geq -x$ and therefore

$$sd(S_2) \geq \sum_{u \in S_2} w_u \frac{1 - x}{2}. \quad (10)$$

The expected approximation ratio of the algorithm can be upper bounded by

$$\begin{aligned}
\max \left\{ \frac{w(S_1)}{sd(S_1)}, \frac{w(S_2) - w(I)}{sd(S_2)} \right\} &\stackrel{(9),(10)}{\leq} \max \left\{ \frac{2}{1+x}, \frac{w(S_2) - w(I)}{w(S_2)(1-x)/2} \right\} \\
&= \max \left\{ 2 - 2x - o(1), \left(2 - \frac{2w(I)}{w(S_2)} \right) (1+x + o(1)) \right\} \\
&\stackrel{(8)}{\leq} 2 - 2x - o(1).
\end{aligned}$$

The claim now follows from the choice the parameter x . \square

A.3 Proof of Theorem 3.1

The proof follows by reducing the graphs used by Kleinberg and Goemans [30] to show an integrality gap of $2 - \varepsilon$ of relaxation (1) for general vertex cover and any $\varepsilon > 0$, to local $(\Delta + 1)$ -colorable graphs $G_{\Delta, \varepsilon} = (V, E)$.

Let γ and m be two parameters that depend on ε and Δ . The graph consists of $\Delta + 1$ classes $C_1, \dots, C_{\Delta+1}$ of vertices. Every class is an independent set of 2^m vertices corresponding to the 2^m many m bit sequences of zeroes and ones. Therefore, $|V| = (\Delta + 1) \cdot 2^m$. Two vertices from two different classes are joined by an edge if they differ in exactly $(1 - \gamma)m$ bits. Every vertex has unit weight. This defines a $(\Delta + 1)$ -colorable graph $G_{\Delta, \varepsilon}$.

For any $i \in V \cup \{0\}$, a feasible solution of relaxation (1) for graph $G_{\Delta, \varepsilon}$ is obtained by setting vector v_i as in [30]. Let

$$\alpha = \frac{\varepsilon'}{2}, \quad \beta = \sqrt{1 - \alpha^2}, \quad \text{and} \quad \gamma = \frac{1}{2} - \frac{(1 - \alpha)^2}{2\beta^2},$$

where $\varepsilon' \leq \varepsilon$ is a rational number. Note that this also implies that γ is a rational number and we can therefore choose m such that $(1 - \gamma)m$ is an integer. The corresponding unit vectors in \mathbb{R}^{m+1} , for $i \in V \cup \{0\}$, are as follows: v_0 has all zeros, but the $(m + 1)^{st}$ bit that is 1; For $i \in V$, $v_i^{(p)} = \beta/\sqrt{m}$ if the p^{th} bit of i is 1, $v_i^{(p)} = -\beta/\sqrt{m}$ if it is zero, and $v_i^{(m+1)} = \alpha$.

It is easy to check that the defined vectors satisfy the constraints of (1) and

$$sd(G_{\Delta, \varepsilon}) \leq \frac{1}{2}(1 + \alpha) \cdot (\Delta + 1) \cdot 2^m \quad (11)$$

As in [30], we make use of the following theorem by Frankl and Rödl [18].

Theorem A.1 ([18]). *Let \mathcal{C} be a collection of m -bit strings, ξ a constant satisfying $0 < \xi < 1/2$, and d an even integer satisfying $\xi m < d < (1 - \xi)m$. Then for some constant δ depending only on ξ , if $|\mathcal{C}| > (2 - \delta)^m$, then \mathcal{C} contains two strings with Hamming distance exactly d .*

Choose $\xi < \gamma$ and let δ denote the constant obtained by applying the above theorem. Let $d = (1 - \gamma)m$, and choose m large enough so that d is an even integer and

$$(2 - \delta)^m \leq \frac{\alpha}{\Delta + 1} \cdot 2^m$$

Now, consider two different classes, C_i and C_j , of vertices. We say that two vertices x and y , with $x \in C_i$ and $y \in C_j$, are *twins* if they have associated the same string of bits. Let $I = I_1 \cup I_2 \cup \dots \cup I_{\Delta+1}$ be some maximum independent set of $G_{\Delta, \varepsilon}$, where $I_i \subseteq C_i$, for $i \in [\Delta + 1]$. For each $i \in [\Delta + 1]$, let T_i denote the set of nodes from I_i that have some twin in one of the sets I_j for $j \neq i$. By Theorem A.1, in $G_{\Delta, \varepsilon}$ if for some i , set T_i is larger than $\frac{\alpha}{\Delta+1} \cdot 2^m$, one of the vertices in T_i must have an edge to its twin in one of the sets I_j , $j \neq i$. Therefore, since I is an independent set, we obtain that, $|T_i| \leq \frac{\alpha}{(\Delta+1)} \cdot 2^m$ for all $i \in [\Delta + 1]$, and we have that

$$|I| = \sum_{i \in [\Delta+1]} |S_i| + \sum_{i \in [\Delta+1]} |T_i| \leq 2^m + \alpha \cdot 2^m.$$

This implies that the following bound on the optimal vertex cover holds.

$$vc(G_{\Delta, \varepsilon}) = |V| - |I| \geq (\Delta + 1) \cdot 2^m \left(1 - \frac{1 + \alpha}{\Delta + 1}\right)$$

Therefore,

$$\frac{vc(G_{\Delta, \varepsilon})}{sd(G_{\Delta, \varepsilon})} \geq \frac{2}{1 + \alpha} - \frac{2}{\Delta + 1} \geq 2 - \frac{2}{\Delta + 1} - \varepsilon$$

□

B On the independent set size

Let α , β , λ and μ be positive parameters. The goal of this section is to compute suitable values for x and c such that the following is guaranteed:

$$\left(\mathcal{N}(\alpha \cdot c) - \lambda \cdot \Delta \cdot \mathcal{N}\left(\frac{c}{\beta \cdot \sqrt{x}}\right) \right) \geq \mu x \quad (12)$$

The proof of the following lemma is a different rewriting of the analysis appeared in [27, 23] where the parameters α , β , μ and λ are all equal to one.

Lemma B.1. *For any fixed $\alpha, \beta, \mu, \lambda > 0$ and sufficiently large Δ , it is possible to satisfy (12) by setting:*

$$c = \sqrt{\frac{2x\beta^2}{1 - (\alpha\beta)^2x} \ln \Delta} \quad (13)$$

$$x = \frac{1 - o(1)}{(\alpha\beta)^2} \cdot \frac{\ln \ln \Delta}{\ln \Delta} \quad (14)$$

Proof. The claim follows by verifying that the following sufficient conditions hold when c and x are chosen as in (13) and (14), respectively:

$$\mathcal{N}(\alpha \cdot c) \geq 2\mu x \quad (15)$$

$$\Delta \cdot \mathcal{N}\left(\frac{c}{\beta \cdot \sqrt{x}}\right) \leq \frac{\mu x}{\lambda} \quad (16)$$

We will use the following bounds [15]:

$$\phi(c)\left(\frac{1}{c} - \frac{1}{c^3}\right) < \mathcal{N}(c) < \frac{\phi(c)}{c} \quad (17)$$

where

$$\phi(c) = \frac{1}{\sqrt{2\pi}} e^{-\frac{c^2}{2}} \quad (18)$$

Condition (15): By (17) and since, for large Δ , we have $\alpha^2 c^2 > 2$, it follows that

$$\mathcal{N}(\alpha c) > \frac{\phi(\alpha c)}{\alpha c} \left(1 - \frac{1}{\alpha^2 c^2}\right) \geq \frac{\phi(\alpha c)}{2\alpha c}$$

Therefore a sufficient condition to satisfy Condition (15) is

$$\frac{\phi(\alpha c)}{\alpha c} = 4\mu x \quad (19)$$

Condition (16): We can guarantee Condition (16) by satisfying the following condition:

$$\Delta \cdot 4\alpha\beta\lambda\sqrt{x} \leq \frac{\phi(\alpha c)}{\phi\left(\frac{c}{\beta\sqrt{x}}\right)} \quad (20)$$

For large Δ we have $x \leq 1/(4\alpha\beta\lambda)^2$ and by (20) it is sufficient that

$$\Delta \leq \frac{\phi(\alpha c)}{\phi\left(\frac{c}{\beta\sqrt{x}}\right)} = e^{-\frac{(\alpha c)^2}{2} + \frac{(c/\beta)^2}{2x}}$$

and Condition (16) (using $x \leq 1/(\alpha\beta)^2$) is satisfied by setting c as in (13). By (13), it follows that Condition (19) holds by choosing a value for x that satisfies the following equality.

$$\frac{\frac{1}{\sqrt{2\pi}} e^{-\frac{2(\alpha\beta)^2 x \ln \Delta}{1 - (\alpha\beta)^2 x}}}{\sqrt{\frac{2(\alpha\beta)^2 x \ln \Delta}{1 - (\alpha\beta)^2 x}}} = 4\mu x \quad (21)$$

By using $1/(1 - (\alpha\beta)^2 x) = 1 + o(1)$ we have:

$$\frac{e^{-x(\alpha\beta)^2(1+o(1)) \ln \Delta}}{\sqrt{4\pi x(\alpha\beta)^2(1+o(1)) \ln \Delta}} = 4\mu x$$

Let $y := x \ln \Delta$ and rewrite the equality:

$$e^{(\alpha\beta)^2(1+o(1))y^{3/2}} = \frac{\ln \Delta}{C}$$

where $C = 4\mu\sqrt{4\pi(\alpha\beta)^2(1+o(1))}$. By applying the logarithm:

$$(\alpha\beta)^2(1+o(1))y + \frac{3}{2} \ln y = \ln \frac{\ln \Delta}{C}$$

Note that for sufficiently large Δ , it is $(\alpha\beta)^2(1+o(1))y + \frac{3}{2} \ln y = (\alpha\beta)^2(1+o(1))y$. Therefore, there is an x that satisfies (21) such that (14) holds. □