Robust Cost Colorings

Takuro Fukunaga* Magnús M. Halldórsson† Hiroshi Nagamochi*

Abstract

We consider graph coloring problems where the cost of a coloring is the sum of the costs of the colors, and the cost of a color is a monotone concave function of the total weight of the class. This models resource allocation problems where the cost of a resource depends on the use of the resource. The specific case of interval graphs is of special interest as multi-criteria interval scheduling. We give an algorithm for all perfect graphs that yields a robust coloring: a particular solution that simultaneously approximates all concave functions. For graphs with uniform weights, we show how to modify the solution to approximate any monotone cost function. We complement these results with a number of hardness results and some exact algorithms on restricted classes of graphs.

1 Introduction

In interval scheduling, jobs are given with fixed start and finish times, and the problem is to find a schedule of all the jobs that is of minimum cost. The standard cost function is the number of machines needed to process the jobs, where a given machine can be assigned a collection of disjoint intervals. This corresponds to minimizing the number of colors in a proper coloring of the corresponding interval graph. While the chromatic number is arguably the most important parameter of a coloring, it cannot possibly capture all practical cost accounting situations. The goal of this paper is to initiate the study of colorings where the cost of a color is a monotone function of the sum of the vertex weights.

Interval scheduling is a form of resource allocation problem, in which the machines are the resource. As argued by Kolen et al. [**?**], operations management has undergone a "transition in the last decennia from resource oriented logistics (where the availability of resources dictated the planning and completion of jobs) to demand oriented logistics (where the jobs and their completion are more or less fixed and the appropriate resources must be found)." The transition is caused both by increased competition and greater client-oriented focus, as well current logistics developments where production is organized by supply chains. They suggest that this implies a move from traditional scheduling to interval scheduling. The need to marshal resources to service exogenous requests also suggest a need to allow for varying cost functions.

We consider here coloring problems, where the cost of a color class is a function of its weight and the cost of a coloring is a sum of the costs of the color classes. When given an interval graph, this corresponds to an interval schedule where the cost of a resource is a function of the total weights of jobs assigned to that resource. We are interested in *concave* cost functions: the more we use a resource, the less each use of it costs. This corresponds to the economic principle of volume discounts. A common weaker assumption is for a function to be *sub-additive*, which is when the cost of a whole is never more than the sum of its parts; i.e., breaking a set into pieces will not make it cheaper. This correspond to purchasing in chunks, i.e. "cheaper by the dozen".

Cost coloring applies to many of the innumerable practical applications of graph coloring. For instance, the cost of a classroom in a timetabling application is not really a unit; different classrooms may have different costs, depending on size and on the amount of use. The cost of a frequency in frequency allocation may depend on time- or space-limitations of the usage. The cost of fulfilling server requests, e.g., for bandwidth allocation in networks, may depend on the willingness to deploy servers, outsource some of the traffic (at a volume-dependent cost), or to pay the indirect cost of refusing service.

Our results Our main result is algorithm that finds an interval schedule that approximates any concave cost function within a factor of 4. More generally, we obtain a constant factor for all perfect graphs under any sub-additive cost function. We actually get a *simultaneous* approximation: the algorithm is oblivious of the cost function, and approximates all sub-additive cost functions simultaneously. For graphs with uniform weights, we can modify the colorings (in a function-specific way, necessarily) to approximate any monotone cost function. We complement these results with a num-

*Dept. Applied Math. and Physics, Grad. School of Informatics, Kyoto University

†School of Computer Science, Reykjavik University. Work done while visiting Kyoto University.

ber of hardness results and exact algorithms on special classes of graphs.

Some specific cost functions In a companion paper [?], we consider a few specific cost coloring problems. In each case, the cost of a coloring C is a sum of the costs of the individual colors: $Cost(C) = \sum_{I \in C} f(I)$.

One is *"rent-or-buy"* coloring (RBC), where $f(I) = \min(1, w(I))$, which models the simplest case of either buying a machine, or renting one at fixed price per job. An extension would be to form a piecewise linear function. If the slopes of the lines are decreasing, the combined function is concave. The case of two linear functions (TTR) was solved optimally on interval graphs in [?] and a 2-approximation given for perfect graphs.

Threshold colorings correspond to the case of a discrete set of rates, depending on the size of the color class. E.g., we can either rent a taxi, a minivan, or a bus. Namely, we have a set of thresholds $0 = x_1, x_2, \ldots, x_t = \infty$, such that $f(x) = f(x')$ whenever $x_i \leq x \leq x' < x_{i+1}$, for some i. This results in a non-concave function (within each rate bracket, the price-per-item is decreasing), but can be sub-additive if $f(x_{i+1}) \leq f(x_i) \cdot \lceil x_{i+1}/x_i \rceil$, for each i. A special case of the single-threshold case is when the above-threshold cost is too high to be ever cost-effective, e.g. n-fold the below-threshold cost. We have then the *bounded-coloring* problem, which models the case of scheduling conflicting unit-size jobs with bounded number of machines. It is NP-hard on bipartite and interval graphs [?]. The case of a single threshold is approximated within a factor of 4.78 on 3.7 on perfect graphs [?].

The information-theoretic measure of a graph coloring known as *chromatic entropy* was introduced by Alon and Orlitsky [?] in 1986. This relates to the minimum achievable rate for (zero-error) coding with side information, where a random variable X must be transmitted to a receiver that already has some partial information about X. Such entropy colorings also have application to compression of digital image partitions created by segmentation algorithms [?]. In our notation above, this measure is $f_{Ent}(C) = \sum_{I \in C} w(I) \ln(1/w(I))$. It is assumed that $w(G) = 1$. Observe that this is a strictly concave function, but not monotone. Cardinal, Fiorini and Joret [?] gave a number of results for Entropy coloring, including NP-hardness on weighted interval graphs.

Related work Gijswijt, Jost and Queyranne [?] recently introduced a general framework for cost coloring problems that they call *value-polymatroidal*, that includes all concave separable cost functions. It contains monotone problems where moving vertices from a smaller class to a bigger class does not increase the total cost, i.e., when $f(I \cup \{v\}) + f(J) \leq f(I) + f(J \cup \{v\})$, for any independent sets I, J with $f(I) \leq f(J)$. In the case of uniform weights, value-polymatroidality coincides with concavity, but for non-uniform weights it allows for some non-separable or non-concave functions. This includes the Probabilistic coloring, as well as the *max coloring* problem [?, ?], which has the non-separable, monotone cost function $f(I) = \max_{v \in I} w_v$. It does not hold for Threshold coloring, or sub-additive functions. Gijswijt et al [?] give a polynomial time algorithm for any value-polymatroidal cost coloring on co-interval graphs (complements of interval graphs).

Our approach is heavily based on finding large induced subgraphs with small chromatic number. Weighted k-MCS is known to be polynomially solvable in interval graphs [?, ?] and more generally in comparability and co-comparability graphs [?, ?]. It is also easily seen to be solvable in bipartite graphs, and partial k-trees. The problem is NP-hard in chordal (and split graphs), but polynomial for fixed k. A bicriteria approximation algorithm for perfect graphs was recently given in [?], which we shall be using.

To deal with uncertainty, researchers have been increasingly turning to robust optimization, where the solutions are evaluated under a collection of scenarios. Some directions that focus on the data uncertainty are stochastic optimization [?] and demand-robustness [?]. The uncertainty can also involve the measure, as the objective may be poorly characterized. With multiple measures available, it is preferable to satisfy most or all of them with a single solution, than to give individual solutions for each measure. An example of such *robust* approximations are 2-approximate solutions for the classical scheduling problem under all norms [?, ?], and $\sqrt{2}$-approximation for weighted partial matchings of all cardinalities [?]. The work of Goel and Meyerson [?] is perhaps the most closely related to ours. In particular, they give a logarithmic approximations for maximizing all symmetric concave functions. This corresponds to maximizing profit (or throughput) in resource allocation, such as bandwidth allocation in networks. Their scenario is very general and captures throughput versions of our coloring problems. Additionally, their technique of majorization is very similar to our main approach. The approximate minimization of concave functions has also been applied to facility location [?], albeit in a function specific way.

Organization of the paper In Section 2, we give a general schema for approximating cost coloring problems. We show this to give robust approximations of all concave, and more generally sub-additive, separable cost functions. In Section 3, we show how to modify these colorings to handle more general non-concave cost functions, in the case of uniform weights. The resulting colorings are necessarily function-specific, but hold for any monotone separable cost function. In Section 4, we give upper and lower bounds on approximability of cost coloring general graphs. Finally, in Section 5 we give some hardness and limitation results, as well as algorithms for some specific classes of graphs. All of these results are given for abstract cost functions with specified properties.

Note: We shall consider only monotone cost functions in this paper, without explicit indication.

1.1 Notation Let $G = (V, E)$ be a graph given with vertex weights w_v. Let n denote the number of vertices. For a subset $S \subset V$, $G[S]$ denotes the subgraph of G induced by S. For a set S, let $w(S) = \sum_{v \in S} w_v$, and let $w(G) = w(V)$.

A coloring is a partition of V into independent sets. A *k-subgraph* is an induced k-colorable subgraph. We may overload the notation and refer to a vertex subset $S \subset V$ as a k-subgraph if $G[S]$ is k-colorable. *k-MCS* refers to the problem of finding a k-subgraph of maximum total weight, and *Graph Coloring* refers to the classical vertex coloring problem, using the minimum number $\chi(G)$ of colors.

For a set S and vertex u, let $S + u$ denote $S \cup \{u\}$.

A set function $f : 2^V \to \mathbf{R}$ is *monotone* if $f(S+u) \geq f(S)$, for any set U and vertex u; *separable*, if $f(S) = f(T)$ whenever $w(S) = w(T)$; and *value-polymatroidal*, if $f(S+u) - f(S) \geq f(T+u) - f(T)$, for any S, T with $f(T) \geq f(S)$. A separable set function f is *represented* by a function $f' : \mathbf{R} \mapsto R$ if $f(S) = f'(w(S))$ for any S; we shall blur the distinction between a separable set function and its representation. A separable function f is *sub-additive* if $f(x) = \min_{x' \leq x} f(x') \cdot \lceil \frac{x}{x'} \rceil$; *concave* if $f((x + y)/2) = (f(x) + f(y))/2$ for any x, y. Concave separable functions are also value-polymatroidal; the converse is not true in general, but holds in the case of uniform weights.

2 Robust Colorings of Concave and Sub-Additive Cost Functions

We consider a general purpose algorithm that can be applied to a wide variety of cost functions.

Consider the following algorithm. Let $\ell_0 = \lceil \lg \chi(G) \rceil$.

ACS(G)
  $G_0 \leftarrow G$
  for $i \leftarrow 0$ to $\ell_0 - 1$ do
    Let B_i be a maximum weight 2^i-subgraph in G_i
    Color $G[B_i]$ with 2^i colors
    $G_{i+1} \leftarrow G_i \setminus B_i$
  Color the remaining graph with at most $\chi(G)$ colors.

Observe that the last iteration of the loop involves at most $2^{\ell_0 - 1} \leq \chi(G) - 1$ colors, and the total number of colors used up to that point was one less. Hence, adding the last $\chi(G)$ colors used, we have the following.

Observation 2.1. *ACS uses at most* $3(\chi(G) - 1)$ *colors.*

Theorem 2.1. *Given a weighted interval graph, ACS finds a coloring that simultaneously approximates any concave cost function within a factor of 4.*

We shall be needing the following lemma. A much more general result was originally given by Hardy, Littlewood and Polya [**?**] (see [**?**]).

Lemma 2.2. *Let f be a concave function. Let $x_0, x_1, \ldots, x_t, y_0, y_1, y_2, \ldots, y_t$ be non-negative values satisfying*

$$\sum_{i=j}^{t} y_i \leq \sum_{i=j}^{t} x_i, \quad \textit{for each } j = 0, 1, \ldots, t.$$

Furthermore, assume $x_0 \geq x_1 \geq \cdots \geq x_t$. Then,

$$\sum_{i=0}^{t} f(y_i) \leq \sum_{i=0}^{t} f(x_i).$$

Proof. Proof by contradiction. Assume that we have a counterexample that is minimal in that none exists for smaller t or for $y_0 - x_0$ smaller. Then, $y_0 > x_0$, as otherwise a counterexample holds for a shorter sequence. There must be a value $\epsilon > 0$ and index $r \in \{1, 2, \ldots, t\}$ such that

$$y_0 - x_0 \geq \epsilon, \quad \text{and} \quad x_r - y_r \geq \epsilon.$$

Set

$$y_i' = \begin{cases} y_0 - \epsilon & \text{if } i = 0 \\ y_r + \epsilon & \text{if } i = r \\ y_i & \text{otherwise} \end{cases}$$

We have that $y_0 > y_0' \geq x_0 \geq x_r \geq y_r' > y_r$. Thus, since f is concave and $y_0' + y_r' = y_0 + y_r$,

$$f(y_0') + f(y_r') \geq f(y_0) + f(y_r).$$

Thus,

$$\sum_{i=0}^{t} f(y_i') \geq \sum_{i=0}^{t} f(y_i) \geq \sum_{i=0}^{t} f(x_i).$$

Then, we have a counterexample with smaller difference $y_0' - x_0$ in the first coordinate, which is a contradiction.

Proof of Theorem 2.1: Let f be a concave cost function. Concavity implies that $f(x) \le rf(x/r)$, for any $r \ge 1$ and any x.

Let B_i be the set of the color classes used in iteration i of ACS, for $i = 0, 1, \ldots, \ell_0$, whose number $|B_i|$ is $\chi(G)$, for $i = \ell_0$, and 2^i, otherwise. Let b_i be the total weight of B_i, i.e., the sum of the vertex weights.

Consider an optimal coloring of G with respect to f, and assume without loss of generality that the colors are numbered $0, 1, \ldots, \chi_{opt}$ in order of nonincreasing cost (and, by monotonicity, by nonincreasing weight). Let $\ell = \lceil \lg \chi_{opt} \rceil$, and note that $\ell \ge \ell_0 = \lceil \lg \chi(G) \rceil$. For $i \ge 1$, let A_i be the set of the colors numbered $2^{i-1}, \ldots, 2^i - 1$ in the optimal solution, and let A_0 contain color class 0. Let a_i be the weight $w(A_i)$ of A_i, and let $|A_i|$ be the number of color classes in A_i. Since we find a maximum weight b_j-subgraph in each iteration j, we have that

$$a_0 + a_1 + \cdots + a_j \le b_0 + b_1 + \cdots + b_j, \quad \forall j.$$

Thus,

(2.1) $b_j + b_{j+1} + \cdots + b_\ell \le a_j + a_{j+1} + \cdots + a_\ell, \forall j.$

Let $ACS(B_i)$ be the cost of the solution of the algorithm on subgraph B_i. By concavity, we have

$$ACS(B_i) \le |B_i| f(b_i/|B_i|) \le 2^i \cdot f(b_i/2^i).$$

Let $t = 2^\ell = 2^{\lceil \lg \chi_{opt} \rceil}$. For $h = 0, \ldots, t$, let $s = s_h = \lceil \lg(h+1) \rceil$, and define $y_h = b_s/2^s$ as representing a lower bound the average size of a color class i of ACS. Also, let $x_h = a_s/2^s$. From (2.1), we have, for each h that is a power of 2, that

(2.2) $y_h + y_{h+1} + \cdots + y_t \le x_h + x_{h+1} + \cdots + x_t.$

Since it holds for $h = 2^f$ and $h = 2^{f+1}$, for some f, and since $y_{2^f} = y_{2^f+1} = \cdots = y_{2^{f+1}-1}$ and $x_{2^f} = x_{2^f+1} = \cdots = x_{2^{f+1}-1}$, Equation (2.2) does hold for all values of h. Thus, we can apply Lemma 2.2 to bound the cost of the solution of ACS by

$$\begin{aligned}\sum_{i=0}^{\ell} ACS(B_i) &\le \sum_{i=0}^{\ell} 2^i \cdot f(b_i/2^i) = \sum_{h=0}^{t} f(y_h) \\ &\le \sum_{h=0}^{t} f(x_h) = \sum_{i=0}^{\ell} 2^i \cdot f(a_i/2^i).\end{aligned}$$

The weight of any class in A_i is at least that of any class in A_{i+1}, and thus at least their average size $a_{i+1}/2^i$. Thus, the cost of the optimal solution is at least

$$\begin{aligned} OPT &= \sum_i \sum_{I \in A_i} f(w(I)) \\ &\ge f(a_0) + \sum_{i=1}^{\ell-1} |A_i| f(a_{i+1}/|A_{i+1}|) \\ &= f(a_0) + \sum_{i=2}^{\ell} 2^{i-2} f(a_i/2^{i-1}). \qquad (2.3)\end{aligned}$$

By comparing with the bound on ACS, and using that $f(a_0) \ge f(a_1)$, we easily get a ratio of at most 4. □

Using algorithms for k-MCS on comparability graphs [?] and co-comparability graphs [?], we can obtain the same bounds on robustness for these graph classes.

If we are willing to forego of polynomial computability, our method and analysis extends to all graphs.

Corollary 2.1. *For any graph, there exists a coloring that simultaneously approximates every concave cost function within a factor of 4.*

We can also use the strategy of ACS even in some cases where the k-MCS problem is NP-hard. We say that a set S is an (s,t)-*approximation* to k-MCS if S is $s \cdot k$-colorable, and $w(V \setminus S) \le t \cdot w(V \setminus S^*)$, where S^* is a maximum k-MCS. Namely, it gives a subgraph that requires t times as many colors, and leaves behind up to s times the weight left by the optimal solution.

We shall use the following recent result.

Theorem 2.2. ([?]) *There is an algorithm that, given a perfect graph G and integers k and t, yields a $(t, \frac{t}{t-1})$-approximation to k-MCS.*

Theorem 2.3. *Let G be a perfect graph and consider the application of ACS on G with a bicriteria approximation for k-MCS. Then, the coloring produced by ACS simultaneously approximates any concave cost function within a factor of 6.*

Proof. Our argument mimics the proof of Theorem 2.1. Let f be a concave cost function. We find a (z,t)-approximation to k-MCS, where $z = 3/2$ and $t = 3$.

We now have, in place of (2.1), that

$$\sum_{i=j}^{\ell} b_j \le t \cdot \sum_{i=j}^{\ell} a_j, \quad j = 0, 1, \ldots, \ell.$$

Also, the number of colors in B_i, $|B_i|$, is at most $z2^i = 3 \cdot 2^{i-1}$. By concavity, we bound the ACS solution by

$$\begin{aligned} ACS(B_i) &\le |B_i| f(b_i/|B_i|) \le z2^i \cdot f(b_i/(z2^i)) \\ &= 3 \cdot 2^{i-1} \cdot f(b_i/(3 \cdot 2^{i-1})).\end{aligned}$$

We now apply Lemma 2.2 with $x_h = ta_s/(z2^s) = a_s/2^{s-1}$ and $y_h = b_s/(z2^s) = b_s/(3 \cdot 2^{s-1})$, for $h = 0, \ldots, 2^\ell$ and $s = s_h\lceil \lg(h/3) \rceil$, to bound the cost of the solution of ACS by

$$\begin{aligned} \sum_{i=0}^{\ell} ACS(B_i) &\leq \sum_{i=0}^{\ell} 3 \cdot 2^{i-1} \cdot f(b_i/(3 \cdot 2^{i-1})) \\ &\leq \sum_{i=0}^{\ell} 3 \cdot 2^{i-1} \cdot f(a_i/2^{i-1}). \end{aligned}$$

Then, a comparison with Inequality (2.3) gives a ratio of 6.

Sub-additive cost functions: Recall that a cost function f is *sub-additive* if the cost of a whole is never more than sum of its parts; i.e., breaking an independent set into pieces will not be cheaper w.r.t. f. When f is separable, sub-additivity implies that

$$f(x) = \min_{x' \leq x} f(x') \cdot \left\lceil \frac{x}{x'} \right\rceil.$$

We say that a function f is *ρ-approximated* by another function $\hat{f}$ if, $f(x) \leq \hat{f}(x) \leq \rho \cdot f(x)$ holds for all x in the domain of f.

Lemma 2.3. *Let f' be a sub-additive separable cost function defined on the positive reals. Then, there is a concave separable function $\hat{f}$ that 2-approximates f'.*

Proof. Let $\hat{f}$ be the concave hull of f', i.e. the tightest function upper bounding f' which forms a concave polyhedron with the x-axis. We claim that

$$\hat{f}(x) \leq 2 \cdot f'(x), \quad \forall x \in \mathbf{R}^+ \tag{2.4}$$

To see this, let x_m be any point in $\mathbf{R}^+$ with $\hat{f}(x_m) > f'(x_m)$ and x_l and x_r be the hull points on either side of x_m, i.e., the points such that $x_l < x_m < x_r$ and $\hat{f}(x_m)$ is on the straight line from $\hat{f}(x_l) = f'(x_l)$ to $\hat{f}(x_r) = f'(x_r)$. The slope of that line is $(f'(x_r) - f'(x_l))/(x_r - x_l)$. Thus, using that $x_m - x_l > x_r - x_l$, $f'(x_l) \leq f'(x_m) \leq f'(x_r)$ and $x_m < x_r$, we have that

$$\begin{aligned} \hat{f}(x_m) &= f'(x_l) + \frac{f'(x_r) - f'(x_l)}{x_r - x_l}(x_m - x_l) \\ &\leq f'(x_l) + \frac{f'(x_r) - f'(x_l)}{x_r} x_m \\ &\leq f'(x_m) + \frac{f'(x_r) - f'(x_m)}{x_r} \cdot x_m. \end{aligned}$$

Observe that for $a \leq b \leq c$, it holds that $\lceil c/a \rceil \leq \lceil c/b \rceil \cdot \lceil b/a \rceil$. By definition of f', we can verify that $f'(x_r) \leq f'(x_m) \cdot \lceil \frac{x_r}{x_m} \rceil$ as follows:

$$\begin{aligned} f'(x_r) &= \min_{x_0 \leq x_r} f(x_0) \cdot \left\lceil \frac{x_r}{x_0} \right\rceil \leq \min_{x_0 \leq x_m} f(x_0) \cdot \left\lceil \frac{x_r}{x_0} \right\rceil \\ &= \min_{x_0 \leq x_m} f(x_0) \left\lceil \frac{x_m}{x_0} \right\rceil \frac{\lceil \frac{x_r}{x_0} \rceil}{\lceil \frac{x_m}{x_0} \rceil} \\ &\leq \min_{x_0 \leq x_m} f(x_0) \left\lceil \frac{x_m}{x_0} \right\rceil \left\lceil \frac{x_r}{x_m} \right\rceil \\ &= f'(x_m) \left\lceil \frac{x_r}{x_m} \right\rceil. \end{aligned}$$

Thus,

$$\hat{f}(x_m) \leq f'(x_m) \left[1 + \left(\left\lceil \frac{x_r}{x_m} \right\rceil - 1 \right) \frac{x_m}{x_r} \right] \leq 2 \cdot f'(x_m).$$

It is easy to see that Lemma 2.3 implies that a ρ-robust approximation with respect to concave functions yields a 2ρ-robust approximation with respect to sub-additive functions.

Theorem 2.4. *ACS finds a coloring of weighted interval graphs that simultaneously approximates any sub-additive cost function f within a factor of 8.*

Remark: We cannot expect to do much with arbitrary (separable) cost functions, in the case of non-uniform weights. Namely, it is easy to construct functions so that even if the graph is empty, it is NP-hard to get any reasonable approximation ratio. Consider for instance the function

$$f(x) = \begin{cases} \epsilon & \text{if } x = B \\ 1 & \text{otherwise.} \end{cases}$$

Then, by setting $B = w(G)/3$, it is equivalent to 3-Partition to determine if there is a coloring of cost 3ϵ, while otherwise any coloring has cost at least 1. Thus, setting $\epsilon = 2^{-p(n)}$, for some polynomial p, it is NP-hard to approximate cost coloring under f within $2^{p(n)}$ factor. For uniform weights, this problem does not arise, since there are only n possible weights.

3 Arbitrary Monotone Cost Functions

We can actually handle arbitrary monotone cost functions in the case of uniform weights.

We say that a coloring C' *refines* a coloring C if each class in C is a union of classes in C'.

Theorem 3.1. *Let G be a graph with uniform weights, and let f be any monotone cost function. Let C be a coloring found by ACS on G that is ρ-approximate with respect to any concave cost function. Then, there is a refinement C' of C that is 2ρ-approximate with respect to f.*

Proof. Let w be the uniform vertex weight. Given f, form the cost function f' defined by

$$f'(x) = \min_{w \le x_0 \le x} f(x_0) \cdot \lceil x/x_0 \rceil.$$

Namely, we form C' by refining each color class I of C into $\lceil w(I)/x_0 \rceil$ classes of size at most x_0 each, for the best possible $x_0 \le w(I)$. Observe that f' is dominated by f, or $f'(x) \le f(x)$, for all x.

By Lemma 2.3, there is a concave function $\hat{f}$ that 2-approximates f'. Let C_f^{opt} ($C_{\hat{f}}^{opt}$) be an optimal coloring with respect to f ($\hat{f}$), respectively. We can bound the cost of our refined coloring by

$$\begin{aligned} f(C') &\le f'(C) \le \hat{f}(C) \le \rho \cdot \hat{f}(C_{\hat{f}}^{opt}) \le \rho \cdot \hat{f}(C_f^{opt}) \\ &\le 2\rho \cdot f'(C_f^{opt}) \le 2\rho \cdot f(C_f^{opt}). \end{aligned}$$

Corollary 3.1. *Let f be a monotone cost function. In the case of uniform weights, there is an algorithm that finds colorings with respect to f that are 8-approximate on interval graphs and 12-approximate on perfect graphs.*

The refined colorings depend on the actual function f, rather than being robust in the sense of being oblivious of f as in the case of sub-additive functions. This is necessarily so, in that there are no colorings that are robust for all monotone cost functions, even in the uniform case. This holds for any class that includes simple threshold functions. Assuming huge super-threshold cost, no coloring can both approximate threshold q and threshold $p \cdot q$ within less than $\sqrt{p}$.

4 Cost Coloring General Graphs

We give upper and lower bounds for arbitrary cost coloring problems on general graphs.

There is a standard reduction to the MIS problem. The following lemma is a slight strengthening of an argument made numerous times before (see, e.g., [?]).

Lemma 4.1. *Suppose the maximum independent set (MIS) problem can be approximated within a factor of ρ on a hereditary class of graphs. Then, there is a $(\rho \log n, 1)$-approximation of k-MCS. Further, if $\rho = n^{\Omega(1)}$, then there is a $(O(\rho), 1)$-approximation.*

Given the known solvability range of MIS on general graphs [?], we have the following.

Corollary 4.1. *Let ρ_{IS} be the best possible approximation ratio of MIS on general graphs. Then, any concave separable cost coloring problem can be approximated within a factor of $O(\rho_{IS})$. In particular, we get a simultaneous $O(n(\log\log n)^2/\log^3 n)$ approximation.*

This matches the best approximation factor known for the ordinary graph coloring problem [?].

We next prove a strong hardness result on general graphs, that applies to individual cost functions.

Claim 4.2. *Let f be a concave cost function, possibly dependent on n, the number of vertices. Let $\epsilon > 0$. Let $\hat{x}$ be the positive value for which*

$$R = \frac{f(\hat{x} \cdot n^{1-\epsilon})}{f(\hat{x})}$$

is minimized. Then, it is NP-hard to obtain an $n^{1-\epsilon}/R$-approximate coloring with respect to f, even in the case of uniform weights.

Note that the concavity is not an important restriction, since any cost function for the uniform case can be made approximately concave, by our manipulations.

Proof. Let G be a graph input to the Graph Coloring problem. We assign each vertex the weight $\hat{x}/n^{\epsilon'}$, $\epsilon' = \epsilon/2$. Recall that any cost function becomes separable in the case of uniform weights. From the hardness result of Feige and Kilian [?] (and derandomized by Zuckerman [?]), it is hard to distinguish between two cases: when there is a coloring of G with $n^{\epsilon'}$ colors, and when there is no independent set with $n^{\epsilon'}$ vertices. In the former case, the cost of the coloring is by concavity at most $n^{\epsilon'} \cdot f(\hat{x}/n^{\epsilon'} \cdot n^{1-\epsilon'}) = n^{\epsilon'} \cdot f(\hat{x} \cdot n^{1-\epsilon}$. In the latter case, no coloring has color classes larger than $n^{\epsilon'}$, for a cost of at least $n^{1-\epsilon'} \cdot f(\hat{x})$. Hence, the ratio between the two is at least

$$\frac{n^{1-\epsilon'} \cdot f(\hat{x})}{n^{\epsilon'} \cdot f(\hat{x} \cdot n^{1-\epsilon})} \le \frac{n^{1-\epsilon}}{R}.$$

5 Hardness and Exact Algorithms

In this section, we give a number of hardness results on different classes of graphs that hold for any strictly concave function: strong NP-hardness for interval graphs, split graphs, and degree-3 planar graphs. On the positive side, we can give exact algorithms for proper interval graphs, and for graphs with independence number 2.

Limitations on robustness We first show a limitation for robust approximations. It holds for merely distinguishing two different "Rent-or-Buy" functions on uniformly weighted interval graphs. Note that this kind of result is not dependent on a complexity-theoretic assumption (like $P \neq NP$). It is plausible that polynomially computable colorings can only satisfy weaker robustness properties.

Theorem 5.1. *No coloring gives a c-approximation for every concave separable function, for $c < 4/3$.*

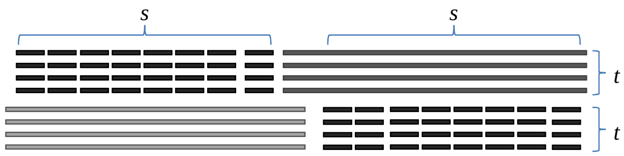

Figure 1: Construction showing the impossibility of simultaneous approximation of better than 4/3

Proof. Consider the interval graph formed by the intervals in Figure 1. There are $2t \cdot s$ small intervals, arranged in disjoint cliques of size t, and $2t$ mutually intersecting large intervals. Each vertex is of unit weight.

Consider a coloring C of this graph, that uses $q \cdot t$ colors, for some $q \geq 2$. Only the $(q-2)t$ colors that contain no large interval can contain $2s$ intervals; the others can only contain s small ones and one large one. Consider the two RBC problems, where the threshold W is 1 and s, respectively. For the former, the ordinary coloring problem, the approximation ratio of C for that measure is $q/2$. For the latter, the optimal solution is to use t colors for the small intervals and $2t$ for the large ones, for a cost of $t \cdot W + 2t$. The coloring C, however, must have at least $2t - (q-2)t = (4-q)t$ colors with s or more nodes, for a cost of at least $(4-q)tW$. The approximation ratio of C for this measure is then at least $(4-q)W/(W+2) = 4 - q - O(1/W)$. For s sufficiently large, this is $4 - q - o(1)$. The larger of $q/2$ and $4-q$ is $4/3$, attained when $q = 8/3$.

Remark: The above construction shows the limits of even *asymptotic* approximation.

We can also get a tighter bound of 1.5 on *absolute* approximation. Setting $t = 1$ in the proof above, we see that there is no coloring of this 2-colorable graph with less than 3 colors that has approximation ratio less than $2 - o(1)$ with respect to the other measure.

NP-hardness We observe that the strong NP-hardness result of Cardinal et al. [**?**] for Entropy Coloring has wider applicability. Rather than duplicate their proof here, we observe that the only property of entropy colorings used in the proof is that of strong concavity.

Proposition 5.1. *Let f be a strongly concave separable cost function. Then, it is strongly NP-hard to find an optimal coloring with respect to f on interval graphs.*

Compare the above result with the polynomial solvability of RBC and TTR on interval graphs [**?**]. It shows that strong concavity is necessary.

Finally, we show the NP-hardness of another subclass of chordal graphs.

Theorem 5.2. *Let f be a strongly concave separable cost function. Then, it is strongly NP-hard to find an optimal coloring with respect to f on split graphs.*

Proof. The reduction is the same as used by Yannakakis and Gavril [**?**] for the maximum k-subgraph problem.

Consider an instance $(X, \mathcal{S})$ of the set cover problem, where X is a base set and $\mathcal{S}$ is a collection of subsets of S. Additionally, each set in $\mathcal{S}$ contains exactly k elements. Let $n = |X|$.

We form a split graph G on the vertex set $V \cup U$, where the independent set V contains a node for each element of X, and the clique U contains a node for each set of $\mathcal{S}$. Additionally, there is an edge between $v \in V$ and $u \in U$ iff the element corresponding to v is *not* contained in the set in $\mathcal{S}$ corresponding to u. We assign weights $1/n$ to vertices in V and weight 1 to vertices in U.

An exact cover of $(X, \mathcal{S})$ with k sets corresponds to a coloring C of G, where $|V|/k$ color classes contain one node in U and k nodes in V each, for a weight of $1+k/n$. The remaining color classes contain one node of U each, of weight 1 each. The total number of colors used is $|U|$. No independent set in G has weight more than $1 + k/n$. Thus, no coloring of G has color classes of weight more than $1 + k/n$ or fewer than $|U|$ color classes. By strong concavity, any coloring of G other than C has cost more than C, with respect to f. Hence, determinining if there is a coloring of cost at most

$$\frac{n}{k} \cdot f(1 + k/n) + |\mathcal{S}| \cdot f(1)$$

is equivalent to determining if $(X, \mathcal{S})$ has an exact cover, which is NP-hard.

A concave function f is said to *strictly concave at* x, if $f(x) < \min_{x' < x} f(x') \cdot (x/x')$. A function f is *additive* if $f(x) = c \cdot x$, for some constant c. A concave cost function is said to be *non-additive* if, it is strictly concave at some point $x > 0$. A concave cost function is either additive, or non-additive.

Theorem 5.3. *Let f be a non-additive concave function. Then, it is NP-hard to compute optimal f-colorings on degree-3 planar graphs, even in the uniform case.*

Proof. We modify the NP-hardness reduction for sum coloring planar graphs of [**?**]. Given a graph G, we form a graph G' by replacing each edge uv with a gadget on six vertices: x, y, z, l, in addition to u and v, with the edges uy, xy, xl, xz, yz, zv. See Figure 2. Observe that G' is planar if G is, and of degree-3 if G is.

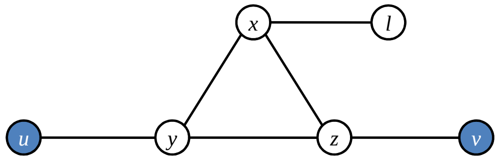

Figure 2: Edge gadget for the hardness on degree-3 planar graphs

We argue that there is an algorithm that reorganizes any coloring C of G' to give a coloring C' that satisfies the following properties, where color i refers to the i-th most frequent color of C: Each l is colored 1, each x is colored 2, each y and z pair is colored 1 and 3, and each of the original nodes are colored 1 and 2 such that those colored 1 form an independent set. In addition, we can argue that the coloring C' *dominates* that of C, in that the number of nodes colored $1, \ldots, i$ in C' is at least that in C, for each $i = 1, 2, 3$. Therefore, under any concave cost function, the cost of C' is at most that of C.

The number of nodes colored i in C' is: $|I_1| + 2|E|$ $(i = 1)$, $n - |I_1| + |E|$ $(i = 2)$, and $|E|$ $(i = 3)$, where I_1 is the independent set of nodes in G colored 1.

Given k, let the weight of each node be $w = \hat{x}/(k + 2|E|)$. Then, by the strict concavity at $\hat{x}$, there is a coloring of cost at most $f(\hat{x}) + f(w \cdot (n - k + |E|)) + f(w|E|)$ if and only if G contains an independent set of size k. By the NP-hardness of the independent set problem in cubic graphs, we obtain the claim.

Exact Algorithms We can show that all cost functions are easy on co-bipartite graphs, i.e., complements of bipartite graph, and more generally on graphs with the independence number $\alpha(G)$ equal to 2. Note that we make no requirement about separability or monotonicity. The result generalizes the same results shown for Probabilistic coloring [**?**] and Entropy coloring [**?**].

THEOREM 5.4. *Let f be any cost function, and G be a graph with $\alpha(G) = 2$. Then, a minimum cost coloring of G with respect to f can be computed in polynomial time.*

Proof. The collection of independent sets in G can also be viewed as a (non-simple) graph H, which is the complement of G along with all self-loops. We assign to each edge in H (and each self-loop) a weight that equals the cost of that independent set according to f. A minimum cost coloring of G under f now corresponds to a minimum cost edge cover of H, or a minimum weight collection of edges that cover all the vertices. It is not hard to see how to reduce an edge cover to a minimum weight matching problem, in the case of monotone costs. We describe here how to solve it even when costs are not monotone, i.e., when the cost of a doubleton set $\{u, v\}$ is greater than that of the singleton $\{u\}$.

We form a complete weighted graph H' containing $2n = 2|V(G)|$ vertices, $V(G) = \{v_1, \ldots, v_n\}$ and $u_1, u_2, \ldots, u_n$. The edge-weights are given as follows. The weight of (v_i, v_j) equals $f(\{v_i, v_j\})$ if v_i and v_j are independent in G, and ∞ otherwise. The weight of (v_i, u_j) equals $f(\{v_i\})$, for each i, j, while the weight of (u_i, u_j) is 0.

A maximal matching in H' induces an edge cover in H, by ignoring the dummy (u_i) vertices. Since H' is complete, every maximal matching is perfect. Similarly, a edge cover in H maps to a matchings in H'; in fact, each permutation of the dummy vertices gives a different matching, but of the same cost. Therefore, a minimum cost perfect matching in H' is equivalent to a minimum cost edge cover in H. The former is solvable in polynomial time [**?**].

The following result contrasts with the NP-hardness for general interval graphs. It holds also for the subclass of containment interval graphs. The proof idea is borrowed from [**?**].

THEOREM 5.5. *Let G be a proper interval graph with uniform weights. Then, there is a polynomial time algorithm to compute a coloring of G that simultaneously optimizes all concave cost functions.*

Proof. Compute a coloring $I_1, I_2, \ldots, I_\chi$ using the greedy exact coloring algorithm for interval graphs. One view of that algorithm is that it constructs independent sets one at a time as follows: Find the interval with the leftmost starting point, add it to the set, remove all conflicting intervals from the graph, and repeat until the graph is empty. Since the interval graph is proper, that means that the leftmost-starting point rule is also a leftmost-ending point rule, which yields a maximum independent set. This implies that for each k, the color classes $I_1, I_2, \ldots, I_k$ form a maximum k-colorable subgraph. Thus, if $J_1, J_2, \ldots, J_s$ is any other coloring, with $w(J_1) \geq w(J_2) \geq \cdots \geq w(J_s)$, we have that

$$\sum_{i=j} w(J_i) \geq \sum_{i=j} w(I_i), \quad j = 1, 2, \ldots.$$

By concavity and Lemma 2.2 we then have that

$$\sum_i f(w(J_i)) \geq \sum_i f(w(I_i)).$$

Hence, the greedy coloring is optimal with respect to f.

A simple example (path on four vertices, with the end vertices of larger weight) shows that simultaneous optimization does not extend to the case of non-uniform weights.

6 Conclusions

We conclude with some open questions.

- We have shown that for any weighted graph there is a 4-robust coloring, while it is not possible to go below 4/3. Can we close this gap?
- Can we extend the result of Section 3 to non-uniform weights?
- Does the hardness for interval graphs extend to the unweighted case?
- Can we generalize the class of cost functions approximable on interval and perfect graphs, e.g. allowing some form of non-separability, or sub-additivity over the color classes?