An Asymptotic Approximation Scheme for Multigraph Edge Coloring

PETER SANDERS
University of Karlsruhe
and
DAVID STEURER
Max Planck Institute for Computer Science

The edge coloring problem asks for assigning colors from a minimum number of colors to edges of a graph such that no two edges with the same color are incident to the same node. We give polynomial time algorithms for approximate edge coloring of multigraphs, i.e., parallel edges are allowed. The best previous algorithms achieve a fixed constant approximation factor plus a small additive offset. One of our algorithms achieves solution quality $\mathrm{opt} + \sqrt{9\mathrm{opt}/2}$ and has execution time polynomial in the number of nodes and the *logarithm* of the maximum edge multiplicity.

Categories and Subject Descriptors: G.2.2 [**Mathematics of Computings**]: Discrete Mathematics—*graph algorithms*; F.2.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity

General Terms: Algorithms, Theory

Additional Key Words and Phrases: edge coloring, multigraphs, chromatic index, data migration

1. INTRODUCTION

One of the most fundamental coloring problems asks for assigning colors to edges of a (multi)graph such that no two edges with the same color meet at a node. The number of different colors is to be minimized. For example, if edges represent data packets then an edge coloring with q colors specifies a schedule for exchanging the packets directly and without node contention.

The minimal number of colors needed to color the edges of a graph $G = (V, E)$ is the *chromatic index* $\chi'(G)$, indicated by χ' for short. There are two obvious lower bounds:

$$\chi' \geq \Delta := \max_{v \in V} \mathrm{degree}(v) \tag{1}$$

$$\chi' \geq \Gamma := \max_{H \subseteq V, |H| \geq 2} \frac{|E(H)|}{\lfloor |H|/2 \rfloor} \tag{2}$$

where $E(H)$ denotes the set of edges of the subgraph induced by the vertex set H.

Partially supported by DFG grant SA 933/1-2.

For *bipartite* multigraphs we actually have $\chi' = \Delta$ and optimal colorings can be found very quickly [Cole et al. 2000]. For *simple* graphs, Vizing's algorithm [Vizing 1964] gives a coloring with $\Delta + 1$ colors in time $\mathcal{O}(|E|(|V| + \Delta))$. The same quality can also be achieved in time $\mathcal{O}\big(\min\{\Delta m \log n, m\sqrt{n \log n}\}\big)$ [Gabow et al. 1985].

It is NP-hard to decide whether $\chi' = \Delta$ [Holyer 1981]. Vizing's algorithm can be generalized to color multigraphs with $\Delta + \mu$ colors, where μ is the maximum multiplicity of an edge.

An algorithm by Shannon can color any multigraph with $3/2 \cdot \Delta$ colors [Shannon 1949]. There is a 4/3-approximation algorithm for multigraphs [Hochbaum et al. 1986] but any better constant factor approximation is NP-hard to obtain [Holyer 1981]. However, if we allow a small additive error, much better approximation factors can be obtained. In a sequence of results, approximation guarantees of $7\chi'/6 + 2/3$, $9\chi'/8 + 3/4$ [Hochbaum et al. 1986], and $11\chi'/10 + 4/5$ [Nishizeki and Kashiwagi 1990] have been obtained. All these algorithms have the same basic structure and it can be expected that any approximation of the form $(1 + 1/2k)\chi' + 1 - 1/k$ can be achieved. However, the actual algorithms became more and more complex with a large number of case distinctions that can only be managed using careful exploitation of symmetric cases. After eight more years, the most recent improvement in this direction only affected the additive constant improving it from $1 - 1/k$ to $1 - 3/2k$ [Caprara and Rizzi 1998]. To break out of this road block, we relax the requirement on the additive offset and in exchange obtain better approximation factors. To understand the basic idea behind this approach it is instructive to first have a look at the previous algorithms.

The basic operations are to *color* an edge, to *uncolor* it, or to *shift* it, i.e., on a path with edges alternatingly colored a and b, swap the colors a and b. The edges are colored sequentially in arbitrary order. To color an edge e, constant size subgraphs O containing e are investigated where the edges in O come from some small set of colors. Using an exhaustive case distinction, three basic outcomes are possible: (1) e can be colored using a small number of operations originating in O. (2) O forms a witness that the number of colors can be increased without getting too far away from the optimum. In that case e is colored with the new color. (3) O is enlarged by taking additional colors and nodes into account; now an exhaustive case distinction for the larger graph is necessary. This process eventually has to terminate since for sufficiently large subgraphs, case (1) or (2) has to be applicable. However, the approximation guarantee is determined by the size of the graph for which a complete case distinction is feasible.

Our algorithm uses a similar basic approach but avoids massive case distinctions by investing a small number of additional colors that make it possible to impose an additional structure on O so that the algorithm can handle arbitrarily large subgraphs O. Our algorithm is also more flexible in a number of other ways. Rather than insisting on coloring an arbitrary edge, it picks an uncolored edge e that is parallel to another uncolored edge. The coloring is then "balanced" by coloring e and possibly uncoloring another edge which is *not* parallel to another uncolored edge. Eventually this process will terminate with a graph without parallel uncolored edges. An additional coloring mechanism ensures that subgraphs induced by connected components of uncolored edges must eventually be small.

The remaining uncolored edges can then be colored using Vizing's algorithm. In Section 2 we give a summary of our approach and then a detailed derivation. We obtain two closely related algorithms. For any constant $\epsilon > 0$, the first algorithm computes a coloring with $\max\{\lfloor(1+\epsilon)\Delta\rfloor + 1/\epsilon, \chi' + 3/\epsilon\} \leq (1+\epsilon)\chi' + 3/\epsilon$ colors in time $\mathcal{O}(|E|(\Delta + |V|))$. Note that this is the same asymptotic execution time as for the constant factor approximation algorithms like [Nishizeki and Kashiwagi 1990]. The second algorithm has higher yet still polynomial execution time and gets rid of the parameter ϵ. It uses at most $\widetilde{\chi}' + \sqrt{9\widetilde{\chi}'/2}$ colors.

The above algorithms as well as all previous algorithms for general multigraph edge coloring have execution time polynomial in $|E|$ but are only pseudopolynomial in the number of bits needed to describe a multigraph since edge multiplicities can be encoded as binary numbers. This problem can be fixed by appropriately rounding edge multiplicities but this costs additional colors. In Section 3, we develop a more elegant solution that achieves the same approximation guarantees as the pseudopolynomial algorithms. We exploit that a graph with even edge multiplicities can be colored by coloring a graph with halved edge multiplicities and then using each color twice.

Section 4 summarizes the paper and mentions some open problems.

Related Work

The *fractional edge coloring* problem asks to find a set of matchings $\mathcal{M}$ and weights $w\colon \mathcal{M} \to \mathbb{R}_{\geq 0}$ such that $\sum_{M\in\mathcal{M}} w(M)$ is minimized subject to

$$\forall e \in E : \sum_{\{M\in\mathcal{M}:e\in M\}} w(M) \geq 1.$$

The *fractional chromatic index* $\widetilde{\chi}'$ denotes the total weight of the optimal solution. It is known that $\widetilde{\chi}' = \max(\Delta, \Gamma)$ and it is conjectured that $\widetilde{\chi}' \leq \chi' \leq \widetilde{\chi}' + 1$ [Goldberg 1973; Seymour 1979].

The fractional chromatic index can be found in time polynomial in $|E|$ [Padberg and Wolsey 1984; Feige et al. 2002]. Kahn showed that $\chi' \leq \widetilde{\chi}' + o(\widetilde{\chi}')$ using the probabilistic method [Kahn 1996]. Recently, Plantholt has sharpened this result to $\chi' \leq \widetilde{\chi}' + \mathcal{O}(\log\min(n, \widetilde{\chi}'))$ also using a nonconstructive approach [Plantholt 2003]. It looks like an interesting open problem to develop this approach into a polynomial time algorithm.

Plantholt has also developed a polynomial time algorithm that yields a coloring with at most $\widetilde{\chi}' + \mathcal{O}(\sqrt{n\log n})$ colors [Plantholt 1999]. Note that this may yield a better approximation than our algorithms for graphs with $\Delta = \Omega(n\log n)$. For $\widetilde{\chi}' < 450$ [Nishizeki and Kashiwagi 1990; Caprara and Rizzi 1998] remains the algorithm with the best performance guarantee.

2. A PSEUDOPOLYNOMIAL ALGORITHM

Since the details of our algorithm are fairly technical, we give an outline together with an overview of the technical sections first. In this overview, we do not quantify what adjectives like "small", "sufficiently many",... mean since the appropriate thresholds can only be derived when all the technical ingredients are assembled.

The algorithm works with a collection of matchings $E_1, \ldots, E_q \subset E$ with $q \geq \Delta$

which represent a partial coloring. The maximum color q is increased when it can be proven that q is closer to χ' than required for the claimed approximation guarantee. Let G_0 denote the subgraph induced by the uncolored edges of the input graph G. Color c is *missing* at node v if none of its incident edges is colored c.

Our algorithm first produces a partial coloring such that G_0 is simple and has small connected components. Then it calls Vizing's algorithm to color G_0 using fresh colors. Since the maximum degree of a simple graph with small components is small, this last step will only need a small number of additional colors.

It is easy to ensure that the connected components of G_0 are small: Section 2.2 explains how to color an edge when two nodes in the same component of G_0 have a common missing color. Hence, when this routine is no longer applicable, nodes in a component of G_0 have disjoint missing colors. If there are sufficiently many missing colors at each node, this disjointness property limits the size of components of G_0.

The difficult part of the algorithm is to make G_0 a simple graph. Progress towards this goal is measured using the potential function Φ that is defined as the total number of uncolored edges plus the number of *bad* edges where bad edges are uncolored edges that are parallel to other uncolored edges. Note that Φ can be reduced by coloring an edge, or by coloring a bad edge and uncoloring a *lean* edge where an edge e is lean if e itself and all edges parallel to it are colored.

In order to facilitate this *balancing* operation, we define the concept of an *edge orbit* O constructed around a bad edge e in Section 2.3. Edge orbits are subgraphs with properties that allow us to color the edge e in exchange for uncoloring any other edge in O. In particular, if O contains a lean edge, we can reduce Φ.

When an orbit O lacks a lean edge, we can try to grow it using the techniques described in Section 2.4. An orbit is grown using a "fresh" color c that has not been used previously to grow O and two nodes u, v with the property that either u and v miss c, or u and v have c-colored edges leaving O.

The additional structure imposed by only growing the orbit using fresh colors is the main reason why our algorithms are much simpler than the previous ones. In particular, although growing the orbit requires complex recoloring operations affecting the entire graph, the basic properties of the orbits are invariant under these transformations.

Finally, when an orbit O can neither be grown nor contains a lean edge, we show that it witnesses that G is hard to color—it either contains a very high degree node or it has a high ratio of edges to nodes. In that case, the number of colors q can be increased without going too far away from the lower bounds (1) and (2).

Section 2.5 puts all the pieces together and analyzes two algorithm variants. The faster variant follows the classical framework of an asymptotic approximation scheme. It starts with $(1+\epsilon)\Delta$ colors and terminates using at most $\max\{(1+\epsilon)\Delta + 1/\epsilon, \widetilde{\chi}' + 3/\epsilon\} \leq (1+\epsilon)\widetilde{\chi}' + 3/\epsilon$ colors. For constant ϵ, its running time is $\mathcal{O}(|E|(|V| + \Delta))$ which is asymptotically as good as the best previous algorithms [Nishizeki and Kashiwagi 1990; Caprara and Rizzi 1998] but gives a better approximation guarantee except for very small values of $\widetilde{\chi}'$. The second variant is slower but has a better approximation guarantee. This algorithm takes time $\mathcal{O}(|E|\sqrt{\Delta}(|V| + \Delta))$ and needs at most $\widetilde{\chi}' + \sqrt{9\widetilde{\chi}'/2}$ colors.

2.1 Preliminaries

When referring to graphs we allow parallel edges, unless explicitly stated otherwise. In order to avoid dealing with multi-sets of edges, we view edges not as unordered pairs of nodes but as abstract entities. Then the incidence relation is defined by an implicitly given function ι mapping edges to two element subsets of V. An edge e is *incident* to a node u if $u \in \iota(e)$.

Let $G = (V, E)$ be a (multi)graph. A collection $\mathcal{C} = \{E_i\}_{i \leq q}$ of $q \geq \Delta$ pairwise disjoint edge sets is called a *(partial) q-coloring* of G if each $E_i \subseteq E$ is a matching, i.e., no pair of distinct edges has a common endpoint. For a *color* $c \in \{1, \ldots, q\}$, the set $E_c \in \mathcal{C}$ is called the *color class* of c. An edge e has color c if $e \in E_c$; it is *uncolored* if $e \in E_0 := E \setminus \bigcup \mathcal{C}$. The graph induced by the uncolored edges in $\mathcal{C}$ is denoted $G_0 = (V, E_0)$. We say a color c is *missing* at a node $u \in V$ in coloring $\mathcal{C}$ if u is not covered by $E_c \in \mathcal{C}$, i.e., no edge in E_c is incident to u. By $M(u; \mathcal{C})$, or shortly $M(u)$, we denote the set of colors missing at u in $\mathcal{C}$. As indicated above, we assume that in all considered colorings at least Δ colors are available, so that every node that is incident to an uncolored edge misses at least one color. We also assume $\Delta \geq 3$, for otherwise the edge coloring problem is trivial.

We say that an edge e *leaves* a subgraph $H \subseteq G$ if exactly one of the endpoints of e is contained in $V(H)$. Similarly, a path $P \subseteq G$ *leaves* a subgraph H if $\emptyset \neq V(P) \cap V(H) \neq V(P)$. If an edge or a path does not leave a subgraph H then it is *running within* H. We denote by $H - u$ the subgraph of H obtained by removing node u and all edges incident to u.

For a node u and two colors c and d, let $\texttt{Apath}(u, c, d; \mathcal{C}) \subseteq G$, or shortly $\texttt{Apath}(u, c, d)$, be the connected component of $(V, E_c \cup E_d)$ containing u. Observe that $\texttt{Apath}(u, c, d)$ is a path or a simple cycle since each node has degree at most 2. If $c \in M(u)$ then $\texttt{Apath}(u, c, d)$ is the *c, d-alternating path starting at u*. One of our basic recoloring techniques, namely the shift operation, consists of swapping the colors of such a maximal alternating path $P = \texttt{Apath}(u, c, d; \mathcal{C})$. Let $\mathcal{C}'$ be the coloring obtained by shifting P. Since P is a maximal alternating path in $\mathcal{C}$, all color classes of $\mathcal{C}'$ are still matchings. Also observe that P remains a maximal c, d-alternating path in the shifted coloring, i.e., $\texttt{Apath}(u, c, d; \mathcal{C}) = \texttt{Apath}(u, c, d; \mathcal{C}')$.

For our algorithms, we represent a coloring $\mathcal{C}$ as an array of matchings, indexed by colors. Each matching is stored as an array indexed by nodes such that the entry for a node u contains the node that is matched with u or the entry empty if u is not covered by the matching. Note that with this representation, given u, c, and d, shifting a path $\texttt{Apath}(u, c, d)$ takes time proportional to the length of the path. In addition to the array of matchings, we also store the graph G_0 that is induced by the uncolored edges.

We can assume that in the course of our coloring algorithms we consider only q-colorings with $q < 3\Delta/2$, for otherwise we could directly compute a q-coloring with no uncolored edges [Shannon 1949]. Hence we can, for example, compute the missing colors of a node in time $\mathcal{O}(\Delta)$.

Potential Functions. For distinct nodes u and v, let $uv := \iota^{-1}(\{u, v\})$ denote the set of edges incident to both u and v. For an edge e, let $[e] := \iota^{-1}(\iota(e))$ be the set of edges *parallel* to e. According to the number of uncolored edges parallel to an edge e, we either say that e is *lean*, *even*, or *fat* in a coloring $\mathcal{C}$. Specifically, we partition E into three sets, namely

—$E^{(<)} := \{e \in E : |[e] \cap E_0| = 0\}$, the set of lean edges,
—$E^{(=)} := \{e \in E : |[e] \cap E_0| = 1\}$, the set of even edges,
—$E^{(>)} := \{e \in E : |[e] \cap E_0| > 1\}$, the set of fat edges,

where $E_0 := E \setminus \bigcup \mathcal{C}$ denotes, as before, the set of uncolored edges in $\mathcal{C}$. If an edge is both uncolored and fat then it is called *bad*. Let $E_0^{(>)} := E^{(>)} \cap E_0$ denote the set of bad edges. Now the *potential* of coloring $\mathcal{C}$ is defined as $\Phi(\mathcal{C}) := |E_0| + |E_0^{(>)}|$.

The polynomial-time algorithm devised in Section 3 uses a more general notion of potential. For $M \in \mathbb{N}$, an edge e is called *M-lean* if $|[e] \cap E_0| < M$; it is *M-even* if $|[e] \cap E_0| = M$. In all other cases, e is called *M-fat*. Let $E^{(>M)}$ denote the set of M-fat edges. Then $E_0^{(>M)} := E^{(>M)} \cap E_0$ is the set of *M-bad* edges and $\Phi^{(M)}(\mathcal{C}) := |E_0| + |E^{(>M)} \cap E_0|$ is the *M-potential* of $\mathcal{C}$. Note that $\Phi(\mathcal{C}) = \Phi^{(1)}(\mathcal{C})$.

In Sections 2.2–2.4 we provide recoloring operations for reducing the potential of a coloring. Basically the same operations can be employed to reduce the M-potential of a coloring. The only difference is that we would need to use the generalized notions of lean, even and fat edges instead of the original ones. Also the algorithms presented in Section 2.5 extend in a straight-forward manner to the general case.

2.2 Color Orbits

In the following we show that the potential of a coloring can be reduced if a connected component of G_0 contains two nodes with a common missing color. The main ingredient of our proof is the observation that the shift operation allows to "move" a missing color along an uncolored edge.

The next lemma provides the technical details. The additional conditions (2) and (3) shall enable us to iterate the lemma along paths of uncolored edges.

Lemma 2.1 Missing Color Move. *Given a q-coloring $\mathcal{C}$, an uncolored edge $e \in uv$, and a color $c \in M(u)$, we can compute a q-coloring $\mathcal{C}'$ that either is of lower potential than $\mathcal{C}$ or satisfies the following conditions:*

(1) color c is missing at node v in $\mathcal{C}'$,

(2) all nodes besides u and v miss the same colors in $\mathcal{C}'$ as in $\mathcal{C}$, i.e., $M(x; \mathcal{C}') = M(x; \mathcal{C})$ for each $x \in V \setminus \{u, v\}$,

(3) $\mathcal{C}$ and $\mathcal{C}'$ have the same uncolored edges.

Proof. Let d be a color missing at node v. We may assume $c \neq d$, for otherwise the lemma holds trivially with $\mathcal{C}' := \mathcal{C}$. Now let $\mathcal{C}'$ be the coloring obtained by shifting the c, d-alternating path P starting at node u.

First, we consider the case that P ends at node v. Then $c \in M(v; \mathcal{C}')$ since P was shifted. Hence condition (1) is fulfilled in $\mathcal{C}'$. The additional conditions (2) and (3) are also satisfied by $\mathcal{C}'$.

So we may assume that P does not contain node v. Then color d is still missing at v in $\mathcal{C}'$. On the other hand, d is missing at u in $\mathcal{C}'$ since P was shifted. Therefore we can paint edge e with color d, thereby decreasing the potential, for e was uncolored in $\mathcal{C}$. □

We introduce the notion of a color orbit. A color orbit shall describe a range in which missing colors can be moved by shift operations, specifically by Lemma 2.1.

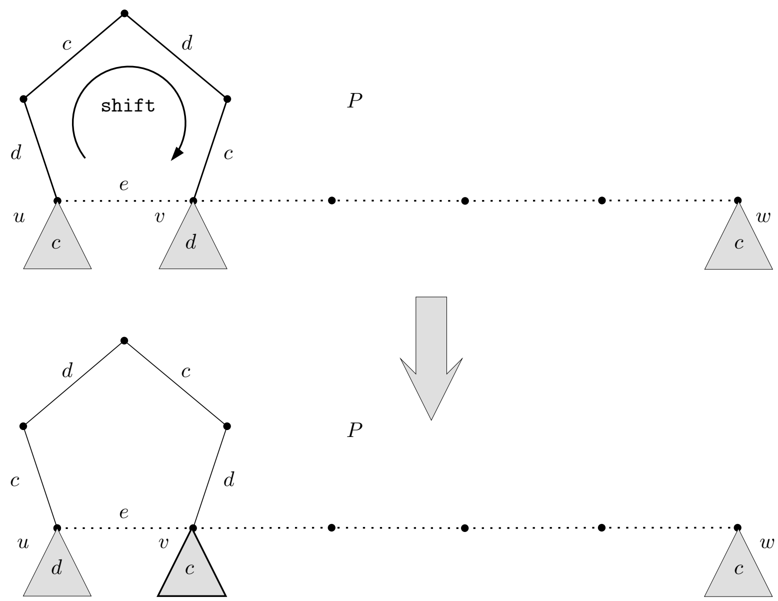

Fig. 1. Applying Lemma 2.1 to an uncolored edge, as in Proposition 2.3, decreases the distance between nodes with common missing color c. *Triangles* at nodes denote missing colors, and *dashed edges* denote uncolored edges.

Definition 2.2. A *color orbit* in a coloring $\mathcal{C}$ is a node set $U \subseteq V(G)$ that is connected by uncolored edges running within U.

A color orbit U is *weak* if some color is missing at two distinct nodes of U. Otherwise, U is called *strong*.

If two nodes in a color orbit have a common missing color, say c, then Lemma 2.1 can be used to move missing color c along a path of uncolored edges until c is missing at both endpoints of an uncolored edge, which then can be painted with color c.

Proposition 2.3. *If a q-coloring $\mathcal{C}$ contains a weak color orbit then we can compute a q-coloring of lower potential than $\mathcal{C}$. Moreover, the uncolored edges of the computed coloring are also uncolored in $\mathcal{C}$.*

Proof. Let U be a weak color orbit with some color, say c, missing at two distinct nodes u and w of U. As U is connected by uncolored edges, a simple path $P \subseteq G_0$ connects u and w. The proof is by induction on the number of edges in P.

If P contains only one edge, then u and w are connected by a single uncolored edge. As c is missing at both u and w, we can paint this edge with color c, thereby decreasing the potential.

So we may assume $|E(P)| > 1$. Let e be the edge in P incident to u and v be the node next to u in P, so that $e \in uv$ (cf. Fig. 1). By our assumption, v is distinct from w. Applying Lemma 2.1 to edge e and color c either decreases the potential directly, in this case we are done, or yields a coloring $\mathcal{C}'$ fulfilling conditions (1)–(3) of Lemma 2.1 (cf. Fig 1. Note that $\Phi(\mathcal{C}) = \Phi(\mathcal{C}')$, for $\mathcal{C}$ and $\mathcal{C}'$ have the same uncolored edges, by condition (3).

Now consider the path $\bar{P} := P - u$ between v and w. Since $\mathcal{C}$ and $\mathcal{C}'$ have the same uncolored edges, $\bar{P}$ is uncolored in $\mathcal{C}'$. Furthermore, conditions (1) and (2) of Lemma 2.1 ensure that c is a common missing color of v and w in $\mathcal{C}'$. Hence we can use the induction hypothesis to obtain a coloring of lower potential. □

By inspection of the proof above, we derive a time bound for the algorithm arising from Proposition 2.3.

Corollary 2.4. *Given a q-coloring $\mathcal{C}$, a weak color orbit U in $\mathcal{C}$, and an upper bound W on the maximum size of a strong color orbit contained in U, it takes time $\mathcal{O}(W \cdot (\Delta + |V|))$ to compute a q-coloring of lower potential.*

Proof. Note that we may assume $W \leq |U| - 1$. We traverse the subgraph of G_0 induced by U until the set of discovered nodes $\bar{U}$ has size $W + 1$. The node set $\bar{U}$ is a weak color orbit. The computation takes time $\mathcal{O}(W \cdot \Delta)$ since at most that many edges are touched during the traversal. Also in this time bound we can compute the missing colors of each node in $\bar{U}$, and find two nodes with a common missing color.

As in the proof of Proposition 2.3, we iterate Lemma 2.1 along a path of uncolored edges, at most W times. In each iteration we look up one missing color of a node and shift a path of length at most $|V|$. Thus the total time for shifting the alternating paths is $\mathcal{O}(W \cdot |V|)$.

Hence we can reduce the potential of $\mathcal{C}$ in time $\mathcal{O}(W \cdot (\Delta + |V|))$. □

2.3 Edge Orbits

In this section, we identify another subgraph structure, called edge orbit, that may allow us to decrease the potential of a coloring, like (weak) color orbits. Instead of directly coloring an uncolored edge, we will reduce the potential by trading lean edges for bad edges, that is, we uncolor a lean edge and in return we can color a bad edge. Thereby we come closer to our goal of eliminating all bad edges.

Edge orbits will be constructed in a way that allows us to iterate the lemma below. The main idea of the lemma is as follows. Let e be an uncolored edge and P be a maximal alternating path ending at the endpoints of P. Suppose that P contains a lean edge f. Then we can recolor the edges of P such that the number of uncolored edges in $[f]$ increases by one and the number of uncolored edges in $[e]$ decreases by one. In case that e was a bad edge in the original coloring, this recoloring operation will decrease the potential. Otherwise, e will be a lean edge in the new coloring, meaning that the leanness of f "moved" to e.

Lemma 2.5 Lean Edge Move. *Given a q-coloring $\mathcal{C}$, an edge $e \in xy$, and two distinct colors $a \in M(x)$ and $b \in M(y)$ such that* $\texttt{Apath}(x, a, b)$ *contains a lean edge, we can compute a q-coloring $\mathcal{C}'$ that either is of lower potential than $\mathcal{C}$ or satisfies the following conditions:*

(1) edge e is lean in coloring $\mathcal{C}'$,

(2) all color classes besides that of colors a and b are the same in $\mathcal{C}$ and $\mathcal{C}'$,

(3) all edges that are bad in $\mathcal{C}$ are also bad in $\mathcal{C}'$,

(4) $\Phi(\mathcal{C}') = \Phi(\mathcal{C})$.

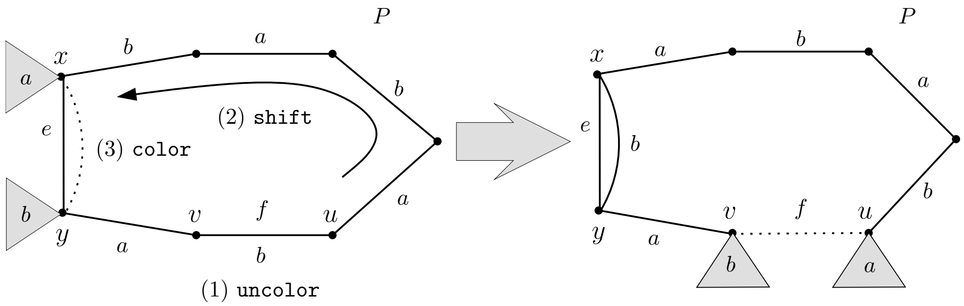

Fig. 2. Illustration of Lemma 2.5. First, we uncolor the edge f between u and v on path P. Then, the a, b-alternating path between u and x is shifted. Finally, we can color an uncolored edge between x and y. Notice that we can generally assume, at least for the figures, that $P = \texttt{Apath}(x, a, b)$ ends at node y, for otherwise we could directly reduce the potential, as in the proof of Lemma 2.1.

PROOF. Suppose f is a lean edge contained in the path $P := \texttt{Apath}(x, a, b)$. Let u, v be the endpoints of f and assume that in the path P node x is the closer to u than to v (cf. Fig. 2). We may assume that e is not lean, for otherwise the lemma is trivially true with $\mathcal{C}' := \mathcal{C}$. Now we proceed in three steps.

First, let $\mathcal{C}^{(1)}$ be the coloring obtained from $\mathcal{C}$ by uncoloring f. Since f was lean in $\mathcal{C}$, it is the only edge in $[f]$ that is uncolored in $\mathcal{C}^{(1)}$ and therefore f is not bad in $\mathcal{C}^{(1)}$. So the potential increased by one and we have $\Phi(\mathcal{C}^{(1)}) = \Phi(\mathcal{C}) + 1$.

Second, we shift the alternating path $Q := \texttt{Apath}(x, a, b; \mathcal{C}^{(1)})$ to get a coloring $\mathcal{C}^{(2)}$ with $\Phi(\mathcal{C}^{(2)}) = \Phi(\mathcal{C}^{(1)})$. Path $Q \subseteq P$ ends at node u, because e is uncolored in $\mathcal{C}^{(1)}$. Since $u \neq y$, color b is still missing at y in $\mathcal{C}^{(2)}$. And since the shifted path Q starts at x, color b is now missing at x in $\mathcal{C}^{(2)}$. Thus, we can paint one of the uncolored edges in $xy = [e]$, say g, with color b. Note that such an edge exists, because e is not lean in $\mathcal{C}$ by assumption.

Third, let $\mathcal{C}^{(3)}$ be the coloring resulting from painting g with color b. Since the number of uncolored edges decreased, $\mathcal{C}^{(3)}$ is of lower potential than $\mathcal{C}^{(2)}$. In case that $\Phi(\mathcal{C}^{(3)}) < \Phi(\mathcal{C}^{(2)}) - 1$, coloring $\mathcal{C}^{(3)}$ even has a lower potential than $\mathcal{C}$ and hence the lemma is true with $\mathcal{C}' := \mathcal{C}^{(3)}$. So we may assume $\Phi(\mathcal{C}^{(3)}) = \Phi(\mathcal{C}^{(2)}) - 1 = \Phi(\mathcal{C})$. Then g was not a bad edge in $\mathcal{C}^{(2)}$, i.e., it was the only uncolored edge in $\mathcal{C}^{(2)}$ parallel to e. Thus e is lean in $\mathcal{C}^{(3)}$, i.e., condition (1) holds. Furthermore $\mathcal{C}^{(3)}$ fulfills also conditions (2)–(4). So $\mathcal{C}' := \mathcal{C}^{(3)}$ is the desired coloring. □

An edge orbit O, as defined below, is a not necessarily induced subgraph of G that consists of two parallel uncolored edges and a collection of alternating paths. The pair of parallel uncolored edges witnesses that these two edges are bad. The alternating paths in O are arranged in a way that allows us, by means of Lemma 2.5, to "move" the leanness of an arbitrary edge in O towards the bad edges of O. In order to ensure that an application of Lemma 2.5 affects only one of the alternating paths in O, we simply demand that no two alternating paths in O have a color in common. In case that an edge orbit contains a lean edge, iterating Lemma 2.5 allows us to eliminate a bad edge in exchange for uncoloring the lean edge. In this way, the potential of the coloring can be decreased.

Definition 2.6. The set of *edge orbits* for a coloring $\mathcal{C}$ is inductively defined:

(1) A subgraph consisting of two parallel uncolored edges is a (trivial) edge orbit.

(2) Let $e \in xy$ be an edge contained in some edge orbit $O \subseteq G$ and for distinct colors $a \in M(x)$ and $b \in M(y)$, let $P := \texttt{Apath}(x, a, b)$ be the a, b-alternating path starting at x.
Then the subgraph $O \cup P$ is also an edge orbit if
—no edge of color a or b is already contained in O and
—the path P leaves the edge orbit O (but may enter O again).

(3) Nothing else is an edge orbit.

An edge orbit is *weak* if it contains a lean edge. Otherwise, it is called *strong*.

We say a color c is *free* for an edge orbit O if no edge of color c is contained in O. Note that we may extend the orbit O by a path P only if the edges of P are painted with free colors. The two uncolored edges contained in the edge orbit are called its *seed*.

A trivial edge orbit consists of two nodes and all colors are free for it. Attaching an alternating path to an edge orbit introduces at least one new node to the edge orbit and reduces the number of free colors by at most two. Therefore any edge orbit O in a q-coloring has at least $q - 2|V(O)| + 4$ free colors.

Also note that an edge orbit remains an edge orbit if only free color classes or uncolored edges besides the seed are changed.

As for weak color orbits, we can decrease the potential of a coloring in presence of a weak edge orbit. See Figure 3 for a small example.

Proposition 2.7. *If a q-coloring $\mathcal{C}$ contains a weak edge orbit, then we can compute a q-coloring of lower potential than $\mathcal{C}$.*

Proof. Let O be a weak edge orbit in coloring $\mathcal{C}$. Note that O is not a trivial edge orbit, for otherwise it would just consist of two uncolored parallel edges, contradicting the weakness of O. We proceed by induction on the number of nodes in the orbit.

Since O is non-trivial, the orbit $O = \bar{O} \cup P$ consists of a smaller edge orbit $\bar{O}$ and an a, b-alternating path P. Since O is weak, it contains a lean edge f. We may assume that f is contained in P but not in $\bar{O}$, for otherwise $\bar{O}$ would be weak and the induction hypothesis could be applied to $\bar{O}$. By definition, $\bar{O}$ contains an edge $e \in xy$ such that $a \in M(x)$, $b \in M(y)$, and $P = \texttt{Apath}(x, a, b)$. Applying Lemma 2.5 to path P either directly decreases Φ or yields a coloring $\mathcal{C}'$ fulfilling the conditions (1)–(4) of the lemma. We may assume the latter case. The induction hypothesis is now applicable to $\bar{O}$ in coloring $\mathcal{C}'$. First, $\bar{O}$ is still an edge orbit in $\mathcal{C}'$, since the seed was not changed by condition (3), and because, by condition (2), no colors besides a and b were changed in $\mathcal{C}'$. Note that a, b were free for $\bar{O}$ in $\mathcal{C}$, for otherwise $\bar{O} \cup P$ was not an edge orbit in $\mathcal{C}$. Furthermore, $\bar{O}$ is weak in $\mathcal{C}'$ as it contains edge e, which is lean in $\mathcal{C}'$ by condition (1). So we can apply the induction hypothesis to the weak edge orbit $\bar{O}$ in $\mathcal{C}'$ in order to obtain a coloring $\mathcal{C}''$ with $\Phi(\mathcal{C}'') < \Phi(\mathcal{C}')$. By condition (4) we also have $\Phi(\mathcal{C}'') < \Phi(\mathcal{C})$, proving the proposition. □

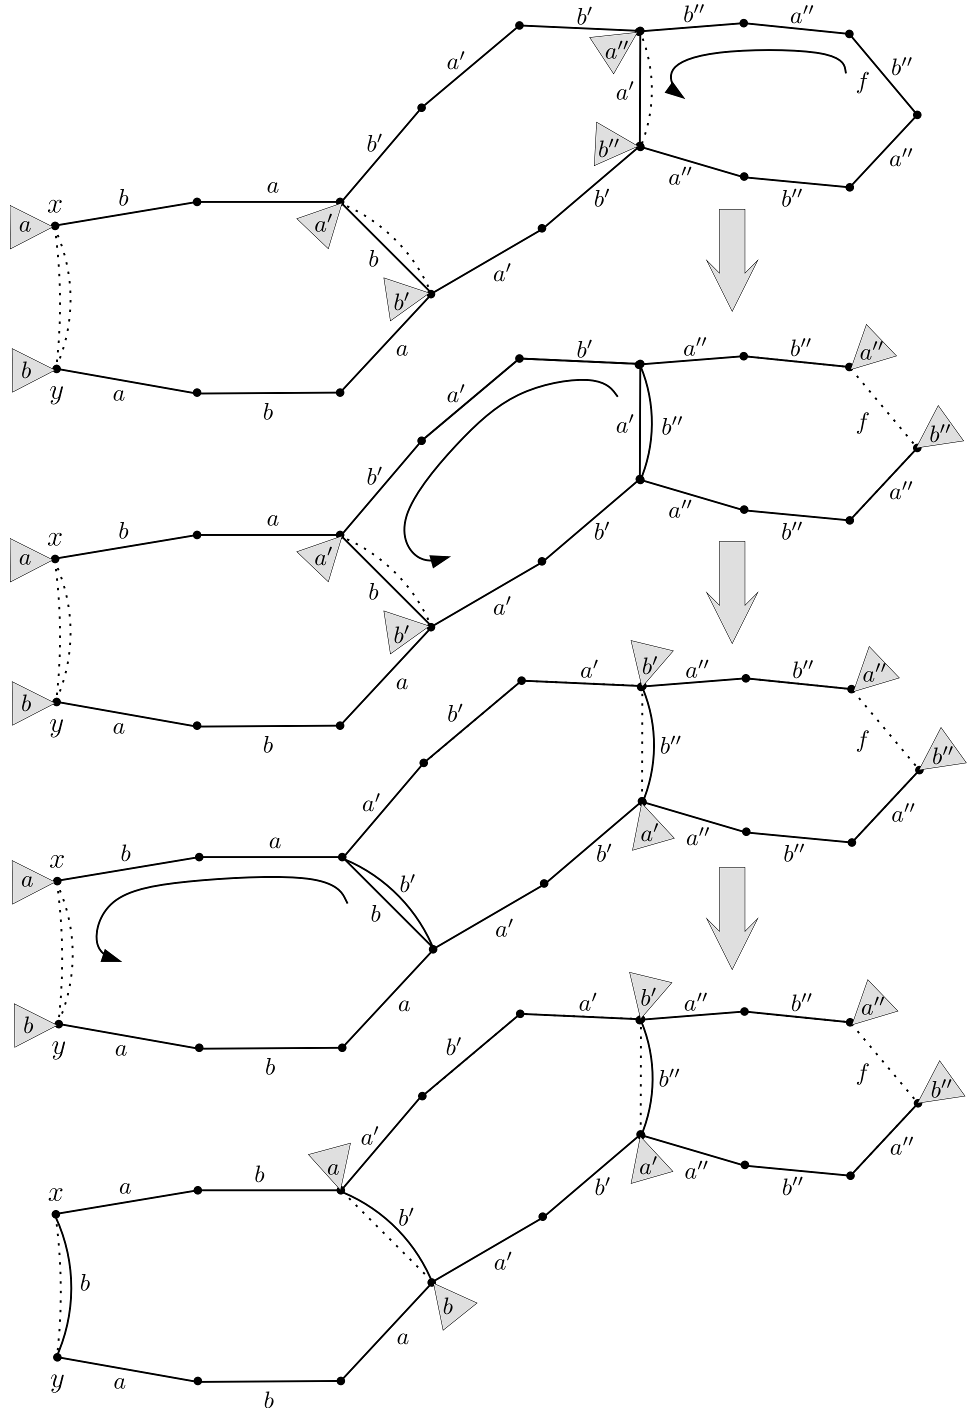

Fig. 3. Illustration of Proposition 2.7. Three applications of Lemma 2.5 allow us to color a bad edge between nodes x and y in exchange for uncoloring the lean edge f.

Corollary 2.8. *Given a q-coloring $\mathcal{C}$ and a weak edge orbit O in $\mathcal{C}$, it takes time $\mathcal{O}(|E(O)|)$ to compute a q-coloring of lower potential.*

Proof. We can find a lean edge of O in time proportional to the number of edges in O. In order to decrease the potential we shift each alternating path of O at most once. The total length of all these paths is at most $|E(O)|$. We assume that O is appropriately stored as a collection of (colored) paths where each path is either linked to an edge of another path or to the seed of O. Then it takes time $\mathcal{O}(|E(O)|)$ to alter the coloring as in the proof of Proposition 2.7. □

Observation 2.9. *The nodes of a strong edge orbit form a color orbit.*

Proof. The nodes of an edge orbit O are connected by the edges $E(O)$. For a strong edge orbit all edges $E(O)$ are non-lean, i.e., parallel to at least one uncolored edge. Thus $V(O)$ is connected by uncolored edges, so that $V(O)$ is a color orbit. □

Only in case that an orbit is both a strong edge orbit and a strong color orbit, neither Proposition 2.3 nor 2.7 yields a reduction in potential.

Definition 2.10. A strong edge orbit is a *hard orbit* if its node set forms a strong color orbit.

2.4 Growing Orbits

A color c is *full* in an edge orbit O if all but at most one node in O are covered by edges of color c that have both endpoints in O, or equivalently if $|E_c \cap E(V(O))| \geq \lfloor |V(O)|/2 \rfloor$ where $E(V(O))$ is the set of edges running within $V(O)$. Note that edge orbits are not necessarily induced subgraphs, so that $E(V(O))$ may differ from $E(O)$. Also note that the color class of a full color has, as the name suggests, a maximum number of edges in $E(V(O))$, since no color class can share more than $\lfloor |V(O)|/2 \rfloor$ edges with $E(V(O))$.

Definition 2.11 *Witnesses.* A hard orbit is a Δ-*witness* if all missing colors of some node are non-free; it is a Γ-*witness* if all free colors of the orbit are full.

The intuition of these witnesses is the following. Assume that almost all colors are free for O. In case of a Δ-witness, we found a node with degree almost as large as the number of available colors. And in case of a Γ-witness, a node subset was found in which almost all color classes are near-perfect matchings. Thus, these witnesses indicate that the number of available colors is close to either Δ or Γ.

In the proposition below, we observe that in absence of witnesses any hard orbit can be grown until it becomes weak.

Proposition 2.12. *Given a hard orbit in a q-coloring $\mathcal{C}$, we can compute a q-coloring of the same potential as $\mathcal{C}$ that either contains a witness or a larger edge orbit.*

For proving Proposition 2.12 we assume the following lemma.

Lemma 2.13. *Suppose a q-coloring $\mathcal{C}$ contains a hard orbit O with a free color c. In either of the following cases we can compute a q-coloring of the same potential as $\mathcal{C}$ that contains a Δ-witness or a larger orbit.*

(1) A node $u \in V(O)$ misses color c and an edge e of color c leaves O.

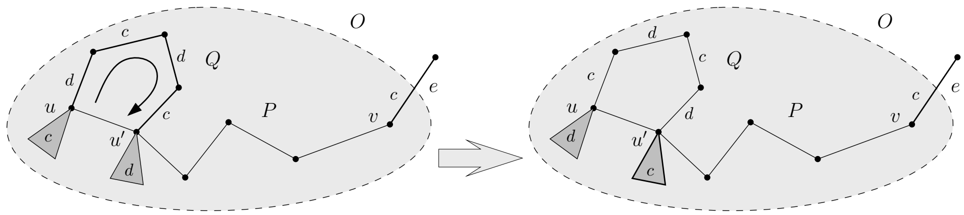

Fig. 4. Illustration of Lemma 2.13, Case (1). Shifting the c, d-alternating path Q reduces the distance between a node with missing color c and a node with leaving edge of color c.

(2) *Two distinct edges e and f of color c leave O.*

Recall that an edge is said to *leave* a subgraph if exactly one of its *endpoints* is contained in the subgraph.

Proof of Proposition 2.12. Let O be a hard orbit in $\mathcal{C}$. We may assume that O has at least one free color, say c, that is not full in O, for otherwise orbit O would be a Γ-witness in $\mathcal{C}$.

Since c is not full in O, two distinct nodes of O, say u and v, are not covered by edges in $E_c \cap E(V(O))$. So these nodes either miss color c or they are incident to edges of color c that leave the orbit O. As O is a hard orbit and henceforth a strong color orbit, color c cannot be missing at both u and v. Thus, one of the two cases of Lemma 2.13 applies. □

The basic idea of the construction below is similar to the proof of Lemma 2.3.

Proof of Lemma 2.13, Case (1). Let c be free color for the hard orbit O in $\mathcal{C}$. Assume that c is missing at a node u of O and that a c-colored edge e leaves O. Suppose v is the endpoint of e in O.

The nodes u and v are connected by a path $P \subseteq O$ since O is a connected graph. We may assume that each node in P misses at least one free color, for otherwise O is a Δ-witness. Suppose the node u' next to u in P misses a free color d. Then let Q be the alternating path $\texttt{Apath}(u', d, c)$. Observe that $O \cup Q$ would be a larger edge orbit if Q left O. So we may assume that the path Q runs within $V(O)$. As $V(O)$ is a strong color orbit, u and u' are the only nodes in O that miss color c or d. Thus, u and u' are the endpoints of Q.

Now, the proof is by induction on the number of edges in P. If P contains only one edge, we have $u' = v$ which implies that e is the first edge of Q. But this is not possible because e leaves O and Q was assumed to run within $V(O)$.

For $|E(P)| > 1$, shifting Q yields a new q-coloring $\mathcal{C}'$ of the same potential as G such that c is missing at u' in $\mathcal{C}'$. Since c and d were free colors, O is still a hard orbit. As Q runs within $V(O)$, e is not contained in Q and therefore it is still painted with color c in $\mathcal{C}'$. So we can apply the induction hypothesis to u' and e, since u' and v are connected in O by the path $P - u$, which contains $|E(P)| - 1$ edges. □

Proof of Lemma 2.13, Case (2). Let O be a hard orbit with free color c such that two distinct edges $e, f \in E_c$ leave O. Suppose $u \in V(O)$ is the endpoint e in O. We may assume that u misses a free color d, for otherwise O would be a Δ-witness.

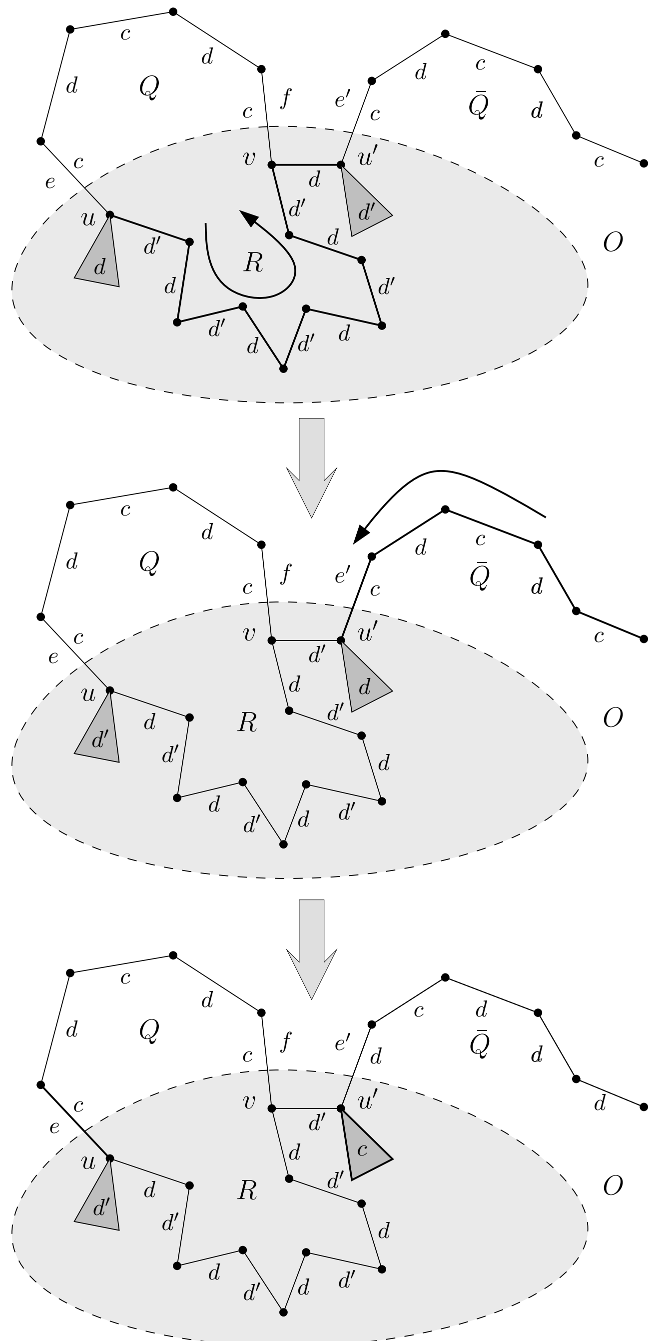

Fig. 5. Illustration of Lemma 2.13, Case (2). In the first step, we reduce case (2b) to case (2a) by shifting the d, d'-alternating path R. The second step reduces the case (2a) to case (1) of Lemma 2.13 by shifting the c, d-alternating path $\bar{Q}$. We can apply Lemma 2.13 to the situation in the third picture since node u' has missing color c and the c-colored edge e leaves $V(O)$.

Now consider the alternating path $Q := \texttt{Apath}(u, d, c)$. Notice that $e \in E(Q)$. If Q ends in O, either both its endpoints miss color d which is impossible since O is hard, or the endpoint other than u misses color c, and so case (1) of Lemma 2.13 could be applied to this node and edge e. Thus we may assume that Q has an endpoint outside of O. We distinguish now two cases.

Case (2a) $f \notin E(Q)$: Then shifting Q yields a coloring $\mathcal{C}'$ so that node $c \in M(u; \mathcal{C}')$ and leaving edge f is c-colored in $\mathcal{C}'$. Therefore case (1) of Lemma 2.13 applies to u and f in $\mathcal{C}'$.

Case (2b) $f \in E(Q)$: This case can be reduced to case (2a) as follows (cf. Fig. 5). Since Q has an endpoint outside of O, there exists a subpath $\bar{Q} \subseteq Q$ such that $\bar{Q}$ and O share only one node, say u'. Observe that $\bar{Q} \neq Q$ and hence $u' \neq u$ since Q enters O at least twice, for $f \in E(Q)$. Let e' be the edge in $\bar{Q}$ that leaves O. Note that e' is incident to u'. As before, we may assume that u' misses a free color, say d'. Then let $R := \texttt{Apath}(u', d', d)$ be the d', d-alternating path starting at u'. If R left O then at least one edge of color d or d' would leave O. Thus case (1) of Lemma 2.13 could be applied to this edge and either u or u'. So we can assume that R runs completely within the nodes of O. Hence, shifting R changes the colors of neither e nor $\bar{Q}$. Let $\mathcal{C}'$ denote the coloring obtained by shifting R. Since u' misses color d in $\mathcal{C}'$, the edge $e' \in E(\bar{Q})$ cannot have color d and hence has color c. Furthermore, e is not contained in $\texttt{Apath}(u', d, c; \mathcal{C}') = \bar{Q}$ since $V(O) \cap V(\bar{Q}) = \{u'\}$. Thus the case (2a) applies to the leaving c-colored edges e, e' and the c, d-alternating path $\bar{Q}$ with $e \notin E(\bar{Q})$. □

COROLLARY 2.14. *Let O be a hard orbit in q-coloring $\mathcal{C}$. Suppose that we are given a list of the colors that are free and non-full for O, and that for each node in O, we have a list of the missing colors that are free for O. Then it takes time $\mathcal{O}\big(|V(O)|^2 + |V(O')|\big)$ to compute a q-coloring $\mathcal{C}'$ with $\Phi(\mathcal{C}) = \Phi(\mathcal{C}')$ and an edge orbit O' in $\mathcal{C}'$ such that O' either is a witness in $\mathcal{C}'$ or it is larger than O'.*

PROOF. If the list of free and non-full colors was empty, we could report that $O' := O$ is a Γ-witness. Otherwise, O is not a Γ-witness in $\mathcal{C}$ and then, as in the proof of Proposition 2.12, we can apply Lemma 2.13. Suppose O contains W nodes. By a constant number of shift operations we can reduce case (2) or Lemma 2.13 to case (1). For case (1), we perform at most W shift operations. For each shift operation we might have to look up one missing color of a node that is free for O. With the given data structures this look-up takes constant time. Moreover, we have to swap the colors of at most W edges. So the shift operations take time $\mathcal{O}\big(W^2\big)$. If at some point a color look-up failed, i.e., all missing colors of some node were not free for O, then we would report that $O' := O$ is a Δ-witness. Otherwise, as in the proof of Lemma 2.13, we can attach an alternating path to O and obtain a larger edge orbit O' in q-coloring $\mathcal{C}'$. This last step takes time $\mathcal{O}(|V(O')|)$. □

2.5 Algorithms

In this section, we combine the tools developed in the previous sections and design algorithms for producing colorings with no bad edges and no weak color orbits.

The absence of weak color orbits will ensure that the graph G_0 induced by the uncolored edges has no large connected components, say, no component of size exceeding W nodes.

The absence of bad edges witnesses that G_0 is a simple graph. So we can use Vizing's algorithm to color G_0 using at most $\Delta(G_0) + 1 \leq W$ colors.

Thus, we can turn any partial coloring without weak color orbits and bad edges into a coloring where all edges are colored using at most W additional colors.

Let $W(q)$ be the maximum size of a strong color orbit in any q-coloring of G. Note that for any q-coloring without weak color orbit, all components of the graph G_0 induced by the uncolored edges are strong color orbits and thus have size at most $W(q)$. So $W(q)$ satisfies the condition on W stated above. In the following, we use W as shorthand for $W(q)$.

Our algorithm first applies the following proposition in order to eliminate all bad edges.

PROPOSITION 2.15. *Given a q-coloring $\mathcal{C}$, we can compute a q'-coloring $\mathcal{C}'$ with $q' \geq q$ in time proportional to $|E_0| \cdot W \cdot (\Delta + |V| + W^2)$ such that*

(1) $\mathcal{C}'$ contains no edge orbit and hence no bad edges,

(2) either $q' = q$ or there is a $(q'-1)$-coloring $\widetilde{\mathcal{C}}$ containing a witness.

Then it uses the next proposition for eliminating all weak color orbits. Note that this step will not introduce new bad edges.

PROPOSITION 2.16. *Given a q-coloring $\mathcal{C}$, we can compute in time at most proportional to $|E_0| \cdot W \cdot (\Delta + V)$ a q-coloring not containing a weak color orbit. Moreover, the uncolored edges of the computed coloring are also uncolored in $\mathcal{C}$.*

After applying the above propositions we have a coloring with no bad edges and no weak color orbits. Furthermore, the number of colors was only increased in the face of a witness.

For simplicity, we first show how to eliminate all weak color orbits.

PROOF OF PROPOSITION 2.16. We apply the following algorithm to coloring $\mathcal{C}$, maintaining a set $S \subseteq V$ of nodes that could be contained in a weak color orbit.

—Initialize S to the set of nodes incident to an edge in E_0.
—While $S \neq \emptyset$,
  —traverse G_0, starting at some node of S, until the set U of discovered nodes forms a weak color orbit or a connected component of G_0;
  —*Case 1:* if U is a weak color orbit,
    then decrease the potential by applying Proposition 2.3 to U;
  —*Case 2:* if U is a connected component of G_0 but not a weak color orbit,
    then remove U from S.

First, we analyze the cost of the iterations in which Case 2 applied and the potential of the coloring was not reduced. In these iterations each node incident to an edge in E_0 is discovered at most once. Therefore each edge of E_0 is touched at most twice. In order to test that the discovered nodes do not form a weak color orbit we need to compute the missing colors of the nodes. Since we consider at most $|E_0|$ nodes, computing the missing colors takes time $\mathcal{O}(|E_0| \cdot \Delta)$. Thus at most that much time is spent for the iterations in Case 2.

Now we bound the time for the iterations in Case 1. Since the potential of $\mathcal{C}$ is at most $2|E_0|$, we perform at most that many iterations in this case. In each iteration,

U has size at most $W+1$ because every color orbit with more than W nodes is weak and we stop traversing G_0 as soon as the discovered nodes form a weak color orbit. So we process at most $W \cdot \Delta$ edges in one iteration. By Corollary 2.4, it takes $\mathcal{O}(W \cdot (\Delta + |V|))$ time to decrease the potential, given a weak color orbit. Thus the total cost of the iterations in Case 1 is at most proportional to $|E_0| \cdot W \cdot (\Delta + |V|)$. □

Now we show how to eliminate all bad edges.

Proof of Proposition 2.15. We assume that $\mathcal{C}$ contains at least one bad edge, for otherwise the proposition is true with $\mathcal{C}' := \mathcal{C}$. Since the potential of $\mathcal{C}$ is at most $2|E_0|$, it is sufficient to show that we can decrease the potential in time $\mathcal{O}\big(W(\Delta + |V| + W^2)\big)$ if the coloring contains a (trivial) edge orbit. Throughout the algorithm we will maintain a list of bad edges, so that we can find trivial edge orbits in constant time. Given an edge orbit O, the following procedure computes in time $\mathcal{O}\big(W(\Delta + |V| + W^2)\big)$ a coloring of lower potential than the current coloring. In the beginning, we let O be one of the trivial edge orbits of the coloring.

—*Case 1:* If edge orbit O is weak,
then apply Proposition 2.7 to decrease the potential;
—*Case 2:* if the nodes of O form a weak color orbit,
then apply Proposition 2.3 to decrease the potential;
—*Case 3:* if O is a hard orbit then apply Proposition 2.12;
—*Case 3.1:* if Proposition 2.12 yields a larger orbit $O \cup P$,
then update $O := O \cup P$ and repeat;
—*Case 3.2:* if Proposition 2.12 yields a witness in some q-coloring $\widetilde{\mathcal{C}}$,
then increment $q := q + 1$, introducing a new color $c := q + 1$,
and decrease the potential of $\widetilde{\mathcal{C}}$ by painting an edge in the seed of O with c.

By Observation 2.9 the case distinction is complete. Note that in Case 3 the size of O never exceeds W since the nodes of O would form a weak color orbit if O was a strong edge orbit and $|V(O)| > W$. Also note that $|E(O)| \leq W \cdot \Delta + |V|$ is an invariant of the algorithm because in Case 3 the edge orbit has at most $W \cdot \Delta$ edges and attaching the path P might add no more than $|V|$ edges. So the total time for testing whether O forms a weak edge orbit is $\mathcal{O}(W \cdot \Delta + |V|)$.

Similarly, the total time for testing whether the node set of O forms a weak color orbit is $\mathcal{O}(W \cdot \Delta)$. For every node that gets inserted to O, we test whether one of its $\mathcal{O}(\Delta)$ missing colors is already missing at some node in O. After processing at most $W + 1$ nodes we are guaranteed to find a pair of nodes with a common missing color.

By Corollaries 2.4 and 2.8, the total time for Case 1 and Case 2 is $\mathcal{O}(W \cdot (\Delta + |V|))$.

Case 3 is repeated at most W times. The cost for Case 3.2 is constant. By Corollary 2.14, the cost for Case 3.1 is $\mathcal{O}\big(W^2 + |V|\big)$, assuming that the data structures required by the corollary are maintained. The total time for maintaining a list of missing free colors for each node is $\mathcal{O}(W \cdot \Delta)$ since there are at most Δ insertions and deletions per node. In order to keep track of the non-full color classes we just need to count for each color class the number of edges it shares with $E(V(O))$. The total time for maintaining these counters is $\mathcal{O}(W \cdot \Delta)$.

Thus the total time for the procedure above is $\mathcal{O}\big(W \cdot (\Delta + |V| + W^2)\big)$. □

Note that for a coloring with no edge orbits, all edges are either lean or even and hence the graph induced by the uncolored edges is simple.

For the special case of constant size strong color orbits, the complexity of our algorithms match the complexity of previous algorithms with weaker approximation guarantees. In the next section, we shall see that strong color orbits indeed have size at most a constant whenever q is by a constant factor larger than Δ.

COROLLARY 2.17. *Under the assumption that $V \leq |E_0|$ and that the maximum size of a strong color orbit is constant, the time complexity of the If the maximum size of a strong color orbit is bounded by a constant, then the algorithms in Proposition 2.15 and 2.16 run in time $\mathcal{O}(|E_0|(|V| + \Delta))$.*

2.6 Analysis

We relate now properties of our orbit structures to the known lower bounds of χ', namely the maximum degree Δ and $\Gamma = \max_{H \subset V, |H| \geq 2} |E(H)| / \lfloor |H|/2 \rfloor$. Recall from the introduction that the fractional chromatic index $\widetilde{\chi}' = \max\{\Delta, \Gamma\}$.

LEMMA 2.18. *Let $\mathcal{C}$ be a q-coloring, and let U be a strong color orbit in G. Then*

$$|U| \leq \frac{q+2}{q - \Delta + 2}.$$

PROOF. First we find a lower bound for the total number of colors missing at the nodes of U, i.e., for $\sum_{u \in U} |M(u)|$. Obviously, every node in U misses at least $q - \Delta$ colors. Since U is connected by uncolored edges of $\mathcal{C}$, at least $|U| - 1$ uncolored edges are incident to nodes of U, and therefore at least $2|U| - 2$ additional colors are missing at nodes of U. Thus we have

$$|U|(q - \Delta) + 2|U| - 2 \leq \sum_{u \in U} |M(u)|.$$

On the other hand, since no two nodes of U share a missing color, we have $\sum_{u \in U} |M(u)| \leq q$. This yields the claimed inequality. □

The next lemma shows that the two kinds of witnesses which we defined for a coloring are nothing but a relaxation of our lower bounds.

LEMMA 2.19. *Let O be a hard orbit in a q-coloring with node set $U = V(O)$.*

(1) *If O is a Δ-witness then $q \leq \Delta + 2|U| - 6$.*
(2) *If O is a Γ-witness then $q \leq \Gamma + 2|U| - 6$.*

PROOF. As noted below of Definition 2.6, the orbit O has at least $q - 2|U| + 4$ free colors.

If O is a Δ-witness then some node in O, say u, misses no free color. At least $q - \Delta + 2$ colors are missing at u, since every node of O is incident to at least two uncolored edges. For the endpoints of the seed, this is trivially true. All other nodes of O are inner nodes of some simple path in O with each edge parallel to at least one uncolored edge, for O is a hard orbit. Note that in a hard orbit, each alternating path forms a cycle with the the edge it was attached to, for otherwise one endnode of the path would share a missing color with an endpoint of that edge, contradicting the assumption that $V(O)$ forms a strong color orbit. Since the missing colors of u and the free colors of O are disjoint, we have

$$q \geq |M(u)| + q - 2|U| + 4 \geq 2q - \Delta - 2|U| + 6,$$

which is equivalent to (1).

If O is a Γ-witness then all free colors of O are full. Note that $E(U)$ contains at most $|E(U)| - |U|$ colored edges because the strong color orbit U is connected by uncolored edges and includes a bad edge. Hence, the number of full colors is at most

$$\frac{|E(U)| - |U|}{\lfloor |U|/2 \rfloor} \leq \Gamma - \frac{|U|}{\lfloor |U|/2 \rfloor} \leq \Gamma - 2.$$

Therefore $q - 2|U| + 4 \leq \Gamma - 2$, which is equivalent to (2). □

Combining the lemmata above, we can show that if sufficiently many colors are available then a strong color orbit can have at most constant size and any witness indicates that the number of colors exceeds $\widetilde{\chi}'$ by at most a constant.

Lemma 2.20. *If $q \geq \lfloor (1+\epsilon)\Delta \rfloor - 1$ for some $\epsilon > 0$ then the following statements hold in any q-coloring $\mathcal{C}$.*

(*1*) *If U is a strong color orbit in $\mathcal{C}$ then $|U| \leq 1/\epsilon + 1$*

(*2*) *If there is a witness then $q \leq \widetilde{\chi}' + 2/\epsilon - 2$.*

Proof. By Lemma 2.18 and the assumption $q \geq \lfloor (1+\epsilon)\Delta \rfloor - 1$, we have

$$|U| \leq \frac{q+2}{q - \Delta + 2} = \frac{\Delta}{q - \Delta + 2} + 1 \leq \frac{\Delta}{\lfloor \epsilon\Delta \rfloor + 1} + 1 \leq 1/\epsilon + 1$$

The second part of the lemma follows from $|U| \leq 1/\epsilon + 1$ and Lemma 2.19. □

Theorem 2.21. *Given $\epsilon > 0$ and a multigraph $G = (V, E)$, we can compute in time at most proportional to $|E| \cdot (\Delta + |V| + 1/\epsilon^2)/\epsilon$ an edge coloring of G that uses at most* $\max\{\lfloor (1+\epsilon)\Delta \rfloor + 1/\epsilon, \widetilde{\chi}' + 3/\epsilon\}$ *colors.*

Proof. We start with a coloring consisting of $\lfloor (1+\epsilon)\Delta \rfloor - 1$ empty color classes. Applying Propositions 2.15 and 2.16 yields a coloring $\mathcal{C}$ without bad edges and weak color orbits. The number of colors has only been increased if there was some witness. By Lemma 2.20, the presence of a witness implies that the number of available colors is at most $\widetilde{\chi}' + 2/\epsilon - 2$. Hence at most $\max\{\lfloor (1+\epsilon)\Delta \rfloor - 1, \widetilde{\chi}' + 2/\epsilon - 1\}$ colors are used in coloring $\mathcal{C}$.

Since all color orbits in $\mathcal{C}$ are strong and hence have size at most $1/\epsilon + 1$ by Lemma 2.20, the connected components of G_0 have maximum size at most $1/\epsilon + 1$. As $\mathcal{C}$ contains no bad edges, G_0 is a simple graph and has maximum degree no more than $1/\epsilon$. Thus, Vizing's algorithms yields a coloring of G_0 using at most $1/\epsilon + 1$ additional colors. Combining the partial coloring $\mathcal{C}$ and the coloring of G_0, we obtain a coloring of G where all edges are colored using at most $\max\{\lfloor (1+\epsilon)\Delta \rfloor + 1/\epsilon, \widetilde{\chi}' + 3/\epsilon\}$ colors.

By Lemma 2.20, the maximum size of a strong color orbit is $\mathcal{O}(1/\epsilon)$ for any q-coloring with $q \geq \lfloor (1+\epsilon)\Delta \rfloor - 1$. Thus by Propositions 2.15 and 2.16, coloring $\mathcal{C}$ is computed in time $\mathcal{O}\big(|E| \cdot (\Delta + |V| + 1/\epsilon^2)/\epsilon\big)$. □

Observe that the number of colors used by the coloring in Theorem 2.21 is at most $(1+\epsilon)\widetilde{\chi}'+3/\epsilon$, and hence at most $(1+\epsilon)\chi'+3/\epsilon$. Thus the theorem shows that the edge coloring problem for multigraphs exhibits a (fully polynomial) *asymptotic approximation scheme*.

In order to obtain the best possible approximation guarantees, we would need to choose ϵ depending on $\widetilde{\chi}'$. Specifically, if we choose $\epsilon = 1/\sqrt{\widetilde{\chi}'/2}$, Theorem 2.21 yields a coloring of G that uses at most $\widetilde{\chi}'+\sqrt{9\widetilde{\chi}'/2}$ colors. It might appear that we need to know the exact value of $\widetilde{\chi}'$ in order to compute such a coloring. However, this is not the case. We will show that if we start with about $\Delta+\sqrt{\Delta}$ colors and proceed as in the proof of Theorem 2.21, the obtained coloring of G uses at most $\widetilde{\chi}' + \sqrt{9\widetilde{\chi}'/2}$ colors.

We need the following refinement of Lemma 2.20.

LEMMA 2.22. *If a q-coloring $\mathcal{C}$ contains a witness then $q \le \widetilde{\chi}' + \sqrt{2\widetilde{\chi}'} - 2$.*

PROOF. Let O be some witness in $\mathcal{C}$, and let U denote the node set of O. Then

$$\frac{q-\widetilde{\chi}'+6}{2} \overset{La.2.19}{\le} |U| \overset{La.2.18}{\le} \frac{q+2}{q-\widetilde{\chi}'+2}.$$

A straight-forward calculation shows for all $q \ge \widetilde{\chi}'$ satisfying the inequality above,

$$q \le \widetilde{\chi}' + \sqrt{2\widetilde{\chi}'+1} - 3.$$

So q does not exceed $\widetilde{\chi}' + \sqrt{2\widetilde{\chi}'} - 2$. □

THEOREM 2.23. *Given a multigraph $G = (V, E)$, we can compute an edge coloring of G in time $\mathcal{O}(|E|\sqrt{\Delta}(\Delta + |V|))$ using $\widetilde{\chi}' + \sqrt{9\widetilde{\chi}'/2}$ colors.*

PROOF. We start with a coloring consisting of $\lfloor \Delta + \sqrt{\Delta} \rfloor$ empty color classes. Applying Propositions 2.15 and 2.16 we obtain a coloring without bad edges and weak color orbits. The number of colors has only been increased when the coloring contained a witness. So by Lemma 2.22, the current coloring uses at most $\widetilde{\chi}' + \sqrt{2\widetilde{\chi}'} - 1$ colors.

Now we compute a coloring of G_0 using Vizing's algorithm and combine it with the partial coloring of G. Since the maximum size of a connected component of G_0 is at most $(q+2)/(q-\Delta+2)$ by Lemma 2.18, the number of colors used by the coloring of G is at most

$$q+\Delta(G_0)+1 \le q + \frac{q+2}{q-\Delta+2} = q+1+\frac{\Delta}{q-\Delta+2} \le \widetilde{\chi}' + \sqrt{2\widetilde{\chi}'} + \sqrt{\widetilde{\chi}'/2},$$

where the last inequality uses that $q+\Delta/(q-\Delta+2)$ is monotonically increasing for $q-\Delta+2 > \sqrt{\Delta}$. Note that $\sqrt{2}+1/\sqrt{2} = \sqrt{9/2}$, and hence the computed coloring of G uses at most $\widetilde{\chi}' + \sqrt{9\widetilde{\chi}'/2}$ colors.

The computation takes time $\mathcal{O}(|E|\sqrt{\Delta}(\Delta + |V|)$ since the maximum size of a strong color orbit is at most $\Delta/(q-\Delta+2)+1 = \mathcal{O}(\sqrt{\Delta})$ for $q \ge \Delta + \sqrt{\Delta} - 1$. □

3. A POLYNOMIAL ALGORITHM

The running times of the algorithms in the previous section are polynomial in the number of edges of the graph. However, if edge multiplicities are encoded as binary numbers, the graph could have a number of edges that is exponential in the input size. In this section, we devise an algorithm with approximation guarantee as in Theorem 2.23 and running time polynomial in the number of nodes and the *logarithm* of the maximum edge multiplicity. Thus, the algorithm runs in polynomial time even if the edge multiplicities of the input graph are encoded in binary.

As indicated in Section 2.1, we will use generalizations of the tools developed in Sections 2.2–2.5. Most importantly, the notion of a bad edge is relaxed, so that, for some positive integer M, an uncolored edge is *(M-)bad* only if it is parallel to at least M other uncolored edges. The other notions are adjusted accordingly. For example, the seed of an edge orbit consists of $M+1$ uncolored edges instead of just two uncolored edges. The *(M-)potential* $\Phi^{(M)}$ of a coloring is the number of uncolored edges plus the number of M-bad edges.

With these generalized notions, the algorithm of Proposition 2.15 will tolerate up to M parallel edges in G_0. In exchange, a witness indicates for large enough M that the number of available colors is indeed less than the lower bound $\widetilde{\chi}'$ of the chromatic index.

The following is a generalization of Lemma 2.19.

LEMMA 3.1. *Let O be a hard orbit in a q-coloring $\mathcal{C}$.*

(*1*) *If O is a Δ-witness then $q \leq \Delta + 2|V(O)| - 2M - 4$.*
(*2*) *If O is a Γ-witness then $q \leq \Gamma + 2|V(O)| - 2M - 4$.*

PROOF. Since a trivial edge orbit cannot be a witness, we may assume that O is non-trivial. Recall from the proof of Lemma 2.19 that in a hard orbit each alternating path forms a cycle with the edge it was attached to. So in the subgraph O, every node has two neighbors and at least $|V(O)|-1$ edges are colored. As noted in the proof of Lemma 2.19, at least $q - 2|V(O)| + 4$ colors are free for O.

Every node of O misses at least $q - \Delta + 2M$ colors since it is connected in O to at least two neighbors by non-lean edges. So it is incident to at least $2M$ uncolored edges and hence at most $\Delta - 2M$ colored edges.

If O is a Δ-witness, some node in O misses no free color. So

$$q - \Delta + 2M \ + \ q - 2|V(O)| + 4 \leq q,$$

which is equivalent to the inequality in case (1).

Furthermore, at least $M \cdot |V(O)|$ edges are uncolored in $E(V(O))$ and thus the number of full colors in O is at most

$$\frac{|E(V(O))| - M \cdot |V(O)|}{\lfloor |V(O)|/2 \rfloor} \leq \Gamma - \frac{M|V(O)|}{|V(O)|/2} = \Gamma - 2M.$$

If O is a Γ-witness, all free colors are full in O and hence $q - 2|V(O)| + 4 \leq \Gamma - 2M$, which shows the inequality in case (2). □

LEMMA 3.2. *For $M = |V|$, any coloring with a witness uses at most $\widetilde{\chi}' - 1$ colors.*

PROOF. For any hard orbit O, we have $|V(O)| \leq M$, so that the inequality follows from Lemma 3.1. □

For the polynomial-time algorithm we need to contract consecutive color classes that have the same combinatorial structure, that is, we represent a q-coloring $\mathcal{C} =$

$\{E_i\}_{i \leq q}$ by the graph of uncolored edges $G_0 = (V, E_0)$ and a collection of matchings $M_{(a_0,a_1]}, \ldots, M_{(a_{I-1},a_I]}$ of V such that $0 = a_0 \leq \ldots \leq a_I = q$ and for every interval $(a_i, a_{i+1}]$ and color $c \in (a_i, a_{i+1}]$, the edge set E_c has the combinatorial structure of M_J, i.e., $\iota(E_c) = M_J$.

In this representation, the time needed by the algorithms of Proposition 2.15 and Proposition 2.16 is independent of Δ but depends polynomially on $I, |V|$, and $|E_0|$. Hence the same holds for the algorithm of Theorem 2.23.

Coloring or uncoloring a single edge increments the number of intervals by at most one. Shifting an alternating path might increase the number of intervals by at most two, assuming without loss of generality that only color classes corresponding to interval endpoints are altered.

The next lemma shows that for any coloring, the number of uncolored edges can be reduced to a number polynomial in $|V|$ without increasing the number of colors beyond $\widetilde{\chi}'$. This reduction step takes time polynomial in the size of the representation of the input coloring.

LEMMA 3.3. *Given a coloring $\mathcal{C}$ that uses at most $\widetilde{\chi}'$ colors contracted to I intervals, we can compute a coloring $\mathcal{C}'$ in time* $\mathrm{poly}(|E_0|, |V|, I)$ *using at most $\widetilde{\chi}'$ colors contracted to $I + \mathrm{poly}(|E_0|, |V|)$ intervals such that at most $|V|^3$ edges are uncolored in $\mathcal{C}'$.*

PROOF. We apply Proposition 2.15 to coloring $\mathcal{C}$ for $M = |V|$. The obtained coloring $\mathcal{C}'$ has no (M-)bad edges. Hence it contains at most $M \cdot |V|^2 = |V|^3$ uncolored edges.

The number of colors has only been increased if there was a witness, i.e., if at most $\widetilde{\chi}' - 1$ colors were available by Lemma 3.2. Hence at most $\widetilde{\chi}'$ colors are used by coloring $\mathcal{C}'$.

The algorithm of Proposition 2.15 colors and uncolors at most $|E_0|$ edges, and it performs $\mathrm{poly}(|E_0|, |V|)$ shift operations. Thus the number of additional intervals is polynomial in $|E_0|$ and $|V|$. □

The *(multiplicity-weighted) adjacency matrix* of a multigraphs $G = (V, E)$ is the matrix $A_G = (|uv|)_{u,v \in V}$. For a function $f : \mathbb{N} \to \mathbb{N}$, we denote by $f(G)$ the graph with adjacency matrix $(f(|uv|))_{u,v \in V}$. Similarly, $G + G'$ denotes the graph with adjacency matrix $A_G + A'_G$.

The next proposition shows how to find a good partial coloring by scaling edge multiplicities.

PROPOSITION 3.4. *Given a multigraphs G with maximum edge multiplicity μ, we can compute in time* $\mathrm{poly}(|V|, \log \mu)$ *a $\widetilde{\chi}'$-coloring of G with at most $|V|^3$ uncolored edges and* $\mathrm{poly}(|V|, \log \mu)$ *color intervals.*

PROOF. We compute the desired coloring using a recursive algorithm. For $\mu = 0$ the graph contains no edges and the proposition is trivially true. Suppose $\mu > 0$. We partition the input graph into three parts, writing $G = \lfloor G/2 \rfloor + \lfloor G/2 \rfloor + (G \bmod 2)$. Note that $\widetilde{\chi}'(G) \geq 2\widetilde{\chi}'(\lfloor G/2 \rfloor)$.

Now we recursively compute a coloring $\bar{\mathcal{C}}$ of $\lfloor G/2 \rfloor$ as in the proposition, so that $\bar{\mathcal{C}}$ uses at most $\widetilde{\chi}'(\lfloor G/2 \rfloor)$ colors and has at most $|V|^3$ uncolored edges.

By simply doubling the endpoints of the intervals in coloring $\bar{\mathcal{C}}$, we obtain a coloring of the graph $2 \lfloor G/2 \rfloor$ using at most $2\widetilde{\chi}'(\lfloor G/2 \rfloor) \leq \widetilde{\chi}'(G)$ colors and having at most $2|V|^3$ uncolored edges. The number of intervals did not increase by this doubling. By including the edges of $(G \bmod 2)$ as uncolored edges, this coloring gives rise to a $\widetilde{\chi}'$-coloring of G with at most $2|V|^3 + |E(G \bmod 2)| \leq 2|V|^3 + |V|^2$ uncolored edges. The algorithm of Lemma 3.3 can reduce the number of uncolored edges to $|V|^3$ again. It increases the current number of intervals I by at most a polynomial in $|V|$ and it runs in time $\mathrm{poly}(|V|, I)$ since the number of currently uncolored edges is polynomial in $|V|$, namely at most $2|V|^3 + |V|^2$. The resulting coloring is as required by the proposition.

It remains to estimate the total running time of the recursive algorithm. The depth of recursion is $\mathcal{O}(\log \mu)$ since the maximum edge multiplicity of the instance we recurse on is at most half the maximum multiplicity of the parent instance. In each level of recursion the number of intervals increases only polynomially in $|V|$. Therefore the maximum number I of intervals is polynomial in $|V|$ and $\log \mu$. Thus only $\mathrm{poly}(|V|, I) = \mathrm{poly}(|V|, \log \mu)$ time is spend in each level of recursion and the total time is polynomial in $|V|$ and $\log \mu$. □

We can apply the same reasoning as in Theorem 2.23 to the coloring obtained by the proposition above. In this way, we obtain a coloring with the same guarantees as in Theorem 2.23.

Theorem 3.5. *Given a multigraphs G, an edge coloring of G can be computed in time* $\mathrm{poly}(|V|, \log \mu)$ *which uses at most* $\widetilde{\chi}' + \sqrt{9\widetilde{\chi}'/2}$ *colors.*

Proof. For $M = 1$, we apply Propositions 2.15 and 2.16 to the $\widetilde{\chi}'$-coloring obtained by Proposition 3.4. The resulting coloring $\mathcal{C}$ has no strong colors and no parallel uncolored edges. The running times of the algorithms of Propositions 2.15 and 2.16 depend polynomially on the number of uncolored edges and the number of color intervals. By Proposition 3.4 both numbers are polynomial in $|V|$ and $\log \mu$.

Now we combine coloring $\mathcal{C}$ with a $\Delta(G_0)+1$-coloring of G_0 obtained by Vizing's algorithm. As in the proof of Theorem 2.23, this edge coloring of G uses at most $\widetilde{\chi}' + \sqrt{9\widetilde{\chi}'/2}$ colors. □

There is another possibility to obtain a polynomial algorithm. Let $\mathcal{M}$ be the set of matchings of V and consider the linear problem to minimize $\sum_{M \in \mathcal{M}} w(M)$ subject to the constraints

$$\forall \{u, v\} \subseteq V. \sum_{M \in \mathcal{M}: uv \in M} w(M) \geq |uv|,$$

where $w(M)$ are non-negative variables. As mentioned in the introduction, the linear program has optimal value $\widetilde{\chi}'$. We can compute in polynomial time an optimal solution w^* that has at most $|V|^2$ non-zero entries. Rounding down each component of w^* yields a $\widetilde{\chi}'$-coloring which is contracted to at most $|V|^2$ matchings and has at most $|V|^3$ uncolored edges. We can proceed with this coloring as in the proof of Theorem 2.23 to obtain an edge coloring using at most $\widetilde{\chi}' + \sqrt{9\widetilde{\chi}'/2}$ colors.

4. CONCLUSION

Our edge coloring algorithms offer a way out of the combinatorial explosion in the number of necessary case distinctions for edge coloring algorithms along the lines of [Hochbaum et al. 1986; Nishizeki and Kashiwagi 1990]. Our algorithms give better approximation except for graphs with very small [Nishizeki and Kashiwagi 1990] or very large [Plantholt 2003] maximum degree.

If one wants to implement our algorithm to solve real world instances, it would be interesting to add further heuristics. For example, the algorithm of Theorem 2.23 could be refined in such a way that it starts with only Δ colors instead of $\Delta + \sqrt{\Delta}$, and then, before adding colors, it first tries to color edges by shifting alternating paths as in Lemma 2.1. It would then get optimal solutions at least for bipartite multigraphs. It might also be interesting to attempt to reduce the maximum degree of G_0 before switching to Vizing's algorithm, e.g., by using balancing operations similar to the ones we apply to bad edges.

There are also many opportunities for speeding up the algorithm. For example, after adding a fresh color, one can color many edges at once by finding a maximal matching in G_0.

ACKNOWLEDGMENTS

We would like to thank Arthur Czumaj, Jeff Kahn, Michael Plantholt, Martin Skutella, Gerhard Wöginger and anonymous referees for useful feedback and valuable pointers to related work.

REFERENCES

Caprara, A. and Rizzi, R. 1998. Improving a family of approximation algorithms to edge color multigraphs. *Information Processing Letters 68*, 11–15.

Cole, R., Ost, K., and Schirra, S. 2000. Edge-coloring bipartite multigraphs in $O(E \log D)$ time. *Combinatorica 21,* 1, 5–12.

Feige, U., Ofek, E., and Wieder, U. 2002. Approximating maximum edge coloring in multigraphs. In *APPROX*. LNCS, vol. 2462. Springer, 108–121.

Gabow, H. N., Nishizeki, T., Kariv, O., Leven, D., and Terada, O. 1985. Algorithms for edge-coloring graphs. Tech. Rep. TR-41/85, Department of Computer Science, Tel Aviv University.

Goldberg, M. K. 1973. On multigraphs of almost maximal chromatic class (in russian). *Diskret Analiz 23*, 3–7.

Hochbaum, D. S., Nishizeki, T., and Shmoys, D. B. 1986. A better than "best possible" algorithm to edge color multigraphs. *Journal of Algorithms 7*, 79–104.

Holyer, I. 1981. The NP-completeness of edge-coloring. *SIAM Journal on Computing 10,* 4, 718–720.

Kahn, J. 1996. Asymptotics of the chromatic index for multigraphs. *Journal of Combinatorial Theory B 68*, 233–254.

Nishizeki, T. and Kashiwagi, K. 1990. On the 1.1 edge-coloring of multigraphs. *SIAM Journal on Discrete Mathematics 3,* 3 (Aug.), 391–410.

Padberg, M. W. and Wolsey, L. A. 1984. Fractional covers for forests and matchings. *Math. Programming 29,* 1, 1–14.

Plantholt, M. 1999. A sublinear bound on the chromatic index of multigraphs. *Discrete Mathematics 202,* 1–3, 201–213.

Plantholt, M. 2003. A combined logarithmic bound on the chromatic index of multigraphs. Manuscript in preparation. The result was presented at the 34th Southeastern Int. Conf. on Combinatorics, Graph Theory, and Computing, Boca Raton.

Seymour, P. D. 1979. Some unsolved problems on one-factorization of graphs. In *Graph Theory and Related Topics*, J. A. Bondy and U. S. R. Murty, Eds. Academic Press, 367–368.

Shannon, C. E. 1949. A theorem on colouring lines of a network. *J. Math. Phys. 28,* 148–151.

Vizing, V. G. 1964. On an estimate of the chromatic class of a p-graph (in Russian). *Diskret. Analiz 3*, 23–30.

Received Month Year; revised Month Year; accepted Month Year