An Asymptotic Approximation Scheme for Multigraph Edge Coloring

Peter Sanders, David Steurer
MPI Informatik, Stuhlsatzenhausweg 85,
66123 Saarbrücken, Germany,
[dsteurer,sanders]@mpi-sb.mpg.de.

July 14, 2004

Abstract

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. Our algorithms achieve arbitrarily good approximation factors at the cost of slightly larger additive term. In particular, for any $\epsilon > 0$ we achieve a solution quality of $(1+\epsilon)\text{opt} + \mathcal{O}(1/\epsilon)$. The execution times of one algorithm are independent of ϵ and polynomial in the number of nodes and the *logarithm* of the maximum edge multiplicity.

1 Introduction

One of the most fundamental coloring problems asks for assigning colors to edges such that no two edges with the same color meet at a node. The most important objective function here is to minimize the number of colors used.

Let us consider a motivating example. Assume that nodes are nodes of a (fully connected) communication network and edges represent data packets of identical size to be exchanged between nodes. Further assume that packets are exchanged directly[1] between sender and receiver in one *round* and that in each round, a node can only be involved in one packet exchange. In this model it does not matter in what direction the data moves so that the graph is an undirected graph. An edge coloring of this communication graph that assigns colors $1, \ldots, q$ to the edges yields a q round communication schedule — just exchange the packets with color c in round c. We see that minimizing the number of colors q will minimize the time needed for exchanging all the packets. Since there may be several packets exchanged between two nodes, we have to deal with a multigraph.

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

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

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

where $E(H)$ denotes the set of edges of the subgraph induced by the vertex set H. For *bipartite* multigraphs we actually have $\chi' = \Delta$ and optimal colorings can be found very quickly [2]. For *simple* graphs, Vizing's algorithm [11] gives a coloring with $\Delta + 1$ colors in time $\mathcal{O}(|E|(|V| + \Delta))$ but it is NP-hard to decide whether $\chi' = \Delta$. Vizing's algorithm can be generalized to color multigraphs with $\Delta + \mu$ colors where μ is the maximum multiplicity of an edge.

There is a 4/3-approximation algorithm for multigraphs but any better approximation is NP-hard to obtain [6]. However, if we allow a small additive error, much better approximation factors can be obtained. In a sequence of results, approximation guarantees of $7\chi'/6 + 2/3$, $9\chi'/8 + 0.75$ [5], and $11\chi'/10 + 0.8$ [8] 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

[1] Refer to [9] for what happens if this direct route assumption is lifted.

improvement only affected the additive constant improving it from $1 - 1/k$ to $1 - 3/2k$ [1]. 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 *color*ing an edge, *uncolor*ing an edge, and *shift*ing, 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 that are defined by edges colored with a small number 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 additional 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. This way 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 a multiply uncolored edge e and "balances" it by coloring one of the parallel edges of e — possibly by uncoloring a completely colored edge. Eventually this process will terminate with a graph without multiply uncolored edges. An additional coloring mechanism makes sure that subgraphs induced by connected components of uncolored edges must eventually be small. The remaining edges can then be colored using Vizing's algorithm. In Section 3 we give a summary of our algorithm and then a detailed derivation.

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 algorithm. This algorithm exploits 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.

More Related Work

The *fractional edge coloring* problem asks to find a set of matchings $\mathcal{M}$ and weights $w(M)$ 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 $\widetilde{\chi}' \leq \chi' \leq \widetilde{\chi}' + 1$ [4, 10].

The fractional chromatic index can be found in time polynomial in $|E|$ using the ellipsoid method [3]. Kahn [7] showed that $\chi' \leq (1+\epsilon)\widetilde{\chi}' + o(\chi')$ using the probabilistic method. Although the proof is not constructive, it might be possible to evolve it into a polynomial time randomized algorithm based on the ellipsoid method and randomized rounding that yields similar bounds as our algorithm. However our algorithm is deterministic, purely combinatorial, comparatively simple and fast, and should yield better approximation guarantees. Indeed, the polynomial algorithm from Theorem 23 can be used to approximate the fractional chromatic index faster than the ellipsoid method: Replace every edge of G by a bundle of M parallel edges (M may even be exponential in $|V|$) and run our algorithm on this graph. In time polynomial in $|V|$ and $\log M$ one gets a coloring that can be used to obtain a fractional coloring of G that is within a factor $(1 + \mathcal{O}(M^{-1/2}))$ from optimal.

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 intentionally do not quantify what adjectives like "small", "sufficiently many", ... mean since the right thresholds can only be derived when all the technical ingredients are assembled.

The algorithm massages a partial coloring of the edges $\tau : E \rightharpoonup \{1, \ldots, q\}$ where q is initially Δ or perhaps some larger value. 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. A color c is *missing* at a node v if none of its incident edges is colored c.

Our algorithm produces a partial coloring where G_0 is simple and has small connected components and then 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 consume few additional colors.

Making sure that the connected components of G_0 are small turns out to be easy. 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 free 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 simple. 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 not simple in G_0. 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 in Section 2.3 that has a bad edge e as its nucleus. Edge orbits are subgraphs with properties that allow us to color one edge in 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. We show that this is possible whenever (1) there is a color c available that has not been used before to grow the orbit. (2) There are at least two nodes in O that either miss c or are incident to a c edge 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 transformation.

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 simpler and faster variant follows the classical framework of an asymptotic approximation scheme. It starts with $(1+\epsilon)\Delta$ colors and terminates using at most $\max((1+\epsilon)\Delta + 1/\epsilon, \chi' + 3/\epsilon)$ colors. For constant ϵ, its running time is $\mathcal{O}(|E|(V+\Delta))$ which is asymptotically as good as the best previous algorithms [8, 1] but gives a better approximation guarantee except for very small values of χ'. The second variant is slower but more adaptive to the input — it only increases the number of colors when there is a necessity to do so. This algorithms needs at most $(1+\sqrt{4.5/\chi'})\chi'$ colors.

2.1 Notation

Since we always refer to multigraphs, we consider edges as abstract entities and not as two element sets or pairs of nodes. The incidence relation is realized by an implicitly given function ι mapping edges to two element subsets of V. An edge e is incident to a node u, if $u \in \iota(e)$. $G = (V, E, \tau)$ is a *partial edge coloring* or *coloring* with *partial color function* $\tau : E \rightharpoonup \{1, \ldots, q\}$. An edge e has *color* c, if $\tau(e) = c$. Only *proper* colorings are considered, i.e., colored edges incident to the same node must have different colors.

We consider a subgraph $H \subseteq G$ to be uncolored, i.e., we can write $H \subseteq G$ and $H \subseteq G'$ even if the edges of H are colored differently in the colorings G and G'. A subgraph P *leaves* another subgraph H, if $V(P) \nsubseteq V(H)$. Let H be a subgraph of G and u a node, then $H - u$ denotes the subgraph obtained by removing u and all edges incident to u from H. Similarly $H \backslash O$ denotes the subgraph obtained by removing all nodes of O and all edges incident to these nodes from H.

For the following definitions consider some arbitrary but fixed coloring G. Then $E_c := \tau^{-1}(c)$ is the set of edges of color c and $E_0 := E \setminus \tau^{-1}(\{1, \ldots, q\})$ is the set of uncolored edges. The graph $G_c := (V, E_c)$

(a matching) is the *color class* of color c and $G_0 := (V, E_0)$ is the graph of uncolored edges. If a node u is not incident to an edge of color c, then c is called *missing* at u and $M(u)$ is the set of all colors missing at a node $u \in V$. We always assume, that at least Δ colors are available in G, implying that every node incident to an uncolored edge has at least one missing color.

Let u be a node of a proper coloring G and c and d two colors, then $\mathtt{Apath}(u, c, d)$ denotes the unique maximal path $P \subseteq G$ containing u and solely consisting of edges of color c or d. If $c \in M(u)$, then we say that $\mathtt{Apath}(u, c, d)$ is the c, d-alternating path starting at u. One of our basic recoloring technique, namely the `shift` operation, consists of swapping the colors of such an maximal alternating path. Due to the maximality of these alternating paths a `shift` operation always preserves the properness of a coloring.

Let $uv := \iota^{-1}(\{u, v\})$ be the set of edges incident to both u and v and for each $e \in E$ let $[e] := \iota^{-1}(\iota(e))$ denotes the set of all edges parallel to e. We partition the edges E of G into three sets, namely

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

We define the set of *bad* edges as $E_0^{(>)} := E^{(>)} \cap E_0$. Now the potential $\Phi(G)$ of a coloring G is $\Phi(G) := |E_0| + |E_0^{(>)}|$. Observe that $\Phi(G) \leq 2|E_0|$.

The lemmata and propositions in the following three sections essentially represent functions mapping a coloring $G = (V, E, \tau)$ to a new coloring $G' = (V, E, \tau')$. For convenience we define every symbol $\circ$, that we defined above for G, also with respect to G' and write $\circ'$.

For the time complexity analysis we define $\mathrm{poly}(f_1, \ldots, f_n) := \bigcup_{k \in \mathbb{N}} \mathcal{O}\big(\prod_{i=1}^{n} f_i^k\big)$.

2.2 Coloring Edges in Large Components of G_0

The following lemma is just a more abstract view of the shift operation. With this operation we can move a missing color along an uncolored edge. By a simple iteration of the operation we can color edges in large components of G_0 and therefore can shrink components of G_0 to a certain extent.

With conditions 1b and 1c we assure that an iteration of the operation is possible.

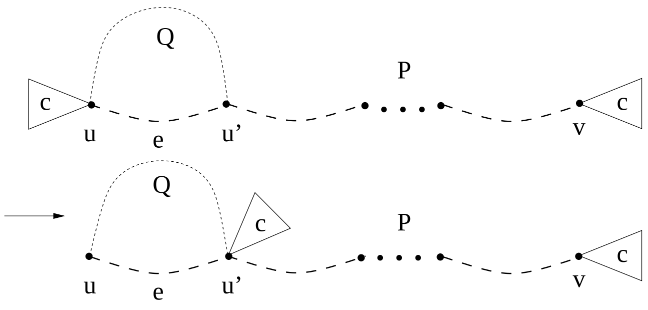

Figure 1: Proposition 3

Lemma 1 (Missing Color Move). *Let $e \in uv \cap E_0$ be an uncolored edge between u and v in G and $c \in M(u)$ a missing color of u. Then we can either decrease the potential Φ by assigning a color to e or compute a coloring G' such that*

a) $c \in M'(v)$, i.e., missing color c moved to v in G',

b) $\forall x \in V \setminus \{u, v\} : M'(x) = M(x)$, i.e., the missing colors of all other nodes were not changed and

c) $G'_0 = G_0$, i.e., the uncolored edges were also not changed.

Proof. Let $d \in M(v)$ be some missing color of y and $Q := \mathtt{Apath}(u, c, d)$ the c, d-alternating path starting at u and ending at ϑ. Now compute G' by shifting Q. If $\vartheta \neq v$, then color d became missing at u and is still missing at v hence we can assign color d to edge e. Otherwise, $\vartheta = v$ and color c became missing at v due to the shifting of Q. Hence G' fulfills condition 1a. Furthermore G' obviously fulfills conditions 1b and 1c. $\square$

Definition 2 (Color Orbit & Weakness). A *color orbit* $O \subseteq G$ is a node induced subgraph of G such that all nodes $V(O)$ are connected by uncolored edges.

A color orbit O is called *weak*, if there are nodes u and v in O that have a common missing color $c \in M(u) \cap M(v)$. Otherwise the color orbit O is called *strong*.

In the next proposition we observe, that given a weak color orbit O we can move the common missing color c by iterating Lemma 1 along a path of uncolored edges until an uncolored edge can be colored.

Proposition 3. *If there is a weak color orbit O in G, then we can decrease the potential Φ by coloring some uncolored edge of O.*

Proof. By definition of a weak color orbit, two nodes u and v in O have a common missing color c and a path $P \subseteq G_0$ joins u and v. Now the proof is by induction over the number of edges in P.

$|E(P)| = 1$: In this case, P consists of a single uncolored edges $e \in uv$. Since u and v are assumed to have a common missing color c, we can assign color c to edge e.

$|E(P)| > 1$: In that case, P contains an uncolored edge e incident to u and some other node $u' \neq v$. We compute G' by applying Lemma 1 on e. If e got colored in G', then the potential was decreased and our proposition is true. Otherwise, c became missing at u' and is still missing at v in G' by 1a respectively 1b. Furthermore, by 1c the uncolored edges were not changed and therefore $\bar{P} := P - u$ is a path of uncolored edges joining u' and v in G'. As $\bar{P}$ is strictly smaller than P, we can use the induction hypothesis to color some edge of $\bar{P}$. □

2.3 Edge Orbits

Again, the following lemma is just a more abstract view of the shift operation. It enables us to move the leanness of an edge along an alternating path. Together with our concept of an *edge orbit* , the operation is used to eliminate bad edges. Conditions 4b-4d are needed to maintain invariants of the edge orbit structure.

Lemma 4 (Lean Edge Move). *Let $e \in xy$ be some edge in G and $P := \texttt{Apath}(x, a, b)$ an alternating path for some colors $a \in M(x)$ and $b \in M(y)$ such that P contains a lean edge $f \in E(P) \cap E^{(<)}$.*

Then we can either decrease the potential Φ or compute a coloring G' such that

a) $e \in E^{(<)\prime}$, i.e., leanness of f moved to e in G',

b) $\forall c \notin \{a, b\} : G'_c = G_c$, i.e., no color class besides that of a or b was changed in G',

c) $E^{(>)\prime} \supseteq E^{(>)}$, i.e., all fat edges in G are fat in G'

d) $\Phi(G') = \Phi(G)$, i.e., the potential was not changed.

Proof. Suppose the lean edge f of P is incident to the nodes u and v and node u appears before v in P. We may assume that e is not lean, otherwise the proposition is trivially true with $G' := G$. Let $G^{(1)}$ be the coloring obtained by uncoloring f. Since f was lean in G, it is not fat in $G^{(1)}$. Thus we have $\Phi(G^{(1)}) = \Phi(G) + 1$. Let $P^{(1)}$ be the a, b-alternating path starting at x in $G^{(1)}$. Observe that $P^{(1)}$ ends at u, since either a or b is missing at u. Now $G^{(2)}$ is obtained by shifting $P^{(1)}$. Clearly $\Phi(G^{(2)}) = \Phi(G^{(1)})$. In $G^{(2)}$ node u has missing color b. Since b is still missing at y and e is assumed to be not lean, there is an uncolored edge in $[e]$, that can be colored with color b. Let $G^{(3)}$ be this new coloring. If e was fat in G, then in $G^{(3)}$ the number of uncolored and bad edges decreased each by at least one. Hence $\Phi(G^{(3)}) \leq \Phi(G^{(2)}) - 2 < \Phi(G)$, i.e., the potential was decreased. And if e was even in G, then e is lean in $G^{(3)}$ and $G' := G^{(3)}$ fulfills conditions 4a-4d. □

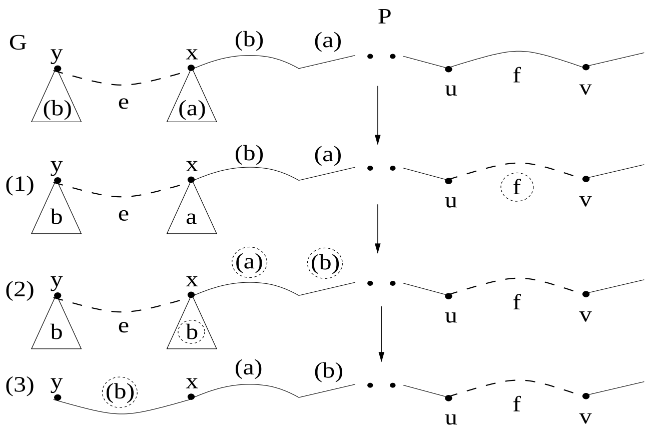

Figure 2: Lemma 4

For iterating Lemma 4 we introduce inductively defined subgraphs of G called *edge orbits.* In these subgraphs some edges are *marked.*

Definition 5 (Edge Orbit). The set of *edge orbits* in a coloring G is inductively defined as

a) For a bad edge $e \in xy$ the graph $O \subseteq G$ induced by x and y, in which all uncolored edges between x and y are marked, is an edge orbit.

b) Let $O \subseteq G$ be an edge orbit, x and y nodes in O and $a \in M(x)$ and $b \in M(y)$ colors, such that

  - an edge between x and y is marked in O,
  - no edge of color a or b is marked in O and
  - the a, b-alternating path $P := \texttt{Apath}(x, a, b)$ leaves O.

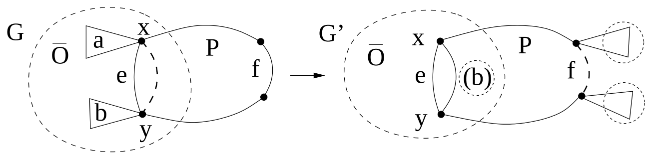

Figure 3: Proposition 7

If $\hat{O} \subseteq G$ is induced by $V(O) \cup V(P)$ and the edges of P and those marked in O are marked in $\hat{O}$, then $\hat{O}$ is an edge orbit. We write $\hat{O} = O + P$.

c) Nothing else is an edge orbit.

We say a *color* c is *marked* in an edge orbit O, if there are edges of color c, that are marked in O.

Note that the size of the orbit increases whenever a pair of colors gets marked. Thus an edge orbit O contains at least $i + 2$ nodes, if $2i$ colors are marked in O. In other words, the number of marked colors is at most $2|V(O)| - 4$.

Also note that an edge orbit is invariant under recoloring operations, that do not involve marked edges.

Definition 6 (Edge Orbit Weakness). An edge orbit $O \subseteq G$ is called *weak*, if an edge marked in O is lean. Otherwise, the edge orbit O is called *strong*.

In the next proposition we observe similarly to Proposition 3 that given a weak edge orbit O we can move the leanness of an edge towards the nucleus of O until some bad edge gets colored and the potential decreased.

Proposition 7. *If a coloring G contains a weak edge orbit O, then we can decrease the potential Φ.*

Proof. The proof is by induction over the size of the orbit.

$|V(O)| = 2$: If O contains only two nodes, then O is a trivial edge orbit induced by two nodes x and y such that the edges between x and y are fat. Thus O cannot be weak and the implication in the proposition is trivially true.

$|V(O)| > 2$: In this case $O = \bar{O} + P$ is induced by the nodes of a smaller orbit $\bar{O}$ and an alternating path P. Since O is weak, it contains a lean edge f. We may assume that f is not marked in $\bar{O}$ but in P, otherwise the induction hypothesis could be applied to $\bar{O}$. By definition of an edge orbit, there is an edge $e \in xy$ marked in $\bar{O}$ such that $P = \mathtt{Apath}(x, a, b)$ for some $a \in M(x), b \in M(y)$.

Let G' be the coloring obtained by applying Lemma 4 to edge e and path P. If the potential was decreased in G', then our proposition is true. So assume that the potential remained unchanged in G' and G' fulfills the conditions 4a-4d. Since colors a and b were not marked in $\bar{O}$, by conditions 4b and 4c, the marked edges of $\bar{O}$ were not changed in G' and therefore $\bar{O}$ is still an edge orbit in G'. By condition 4a, edge e of $\bar{O}$ became lean. Thus we can apply the induction hypothesis to compute a coloring of lower potential. □

2.4 Growing Orbits

We say a color c is *leaving* $O \subseteq G$ at node $u \in V(O)$, if there is an edge $e \in uu'$ of color c incident to u and a node $u' \notin V(O)$.

A color c is called *incomplete* in $O \subseteq G$, if there are two nodes such that no c edge in O is incident to either of them. Otherwise c is called *complete.*

Definition 8 (Hard Orbit & Witnesses). A subgraph $O \subseteq G$, that is a strong edge orbit and a strong color orbit, is called a *hard* orbit.

For a hard orbit O two types of witnesses are defined:

($\mathcal{A}$) : all missing colors of some node u in O are marked in O,

($\mathcal{B}$) : all incomplete colors of O are marked in O.

The intuition of these witnesses is the following. Assume that very few colors are marked in O. In case of an ($\mathcal{A}$) witness, we found a node where the number of incident edges is almost as large as the number of available colors. And in case of a ($\mathcal{B}$) witness, a subgraph was found, in which almost all color classes are maximal matchings. Thus these witnesses indicate, that it is 'almost' impossible to color an additional edge using only the available colors.

Proposition 9. *If O is a hard orbit, then we can either increase the size of the orbit or find an ($\mathcal{A}$) or ($\mathcal{B}$) witness.*

For proving Proposition 9 we assume the following lemma.

Lemma 10. *Suppose O is a hard orbit and color c is not marked in O.*

In either of the following cases we can increase the size of the orbit or find an ($\mathcal{A}$) witness.

a) color c is missing at a node u of O and leaving at a node v of O

b) color c is leaving at nodes u and v in O.

Proof (Proposition 9): We may assume, that there is an incomplete color c, otherwise O has a $(\mathcal{B})$ witness. Let $\{u,v\} \subseteq V(O)$ be the nodes with no incident c edges in O.

If c were missing at u and v, then O would not be a strong color orbit contradicting our hypothesis. Thus we can assume without loss of generality that c is leaving O at v. Hence, either Lemma 10a or 10b is applicable. □

Proof (Lemma 10a): Note that the nodes of an edge orbit are connected by the edges marked in the orbit. Thus there is a path P of edges marked in O joining u and v. We may assume that every node of O has at least one missing color not marked in O, otherwise we would have found an $(\mathcal{A})$ witness. The proof of the lemma is by induction over the number of edges in P.

$|E(P)| = 1$: There is an edge $e \in uv$ marked in O. Since color c is leaving O at v, the alternating path $Q := \texttt{Apath}(v,d,c)$ leaves O for any $d \in M(v)$. As mentioned before, we may assume that color d is not marked in O. Then $\hat{O} = O + Q$ is an edge orbit of G and strictly larger than O.

$|E(P)| > 1$: Let $e \in uu'$ be the first edge in P. Consider the alternating path $Q := \texttt{Apath}(u,c,d)$ for some unmarked color $d \in M(u')$. If Q leaves O, then $\hat{O} = O + Q$ is an edge orbit of G that is strictly larger than O. So suppose Q does not leave O and consider the coloring G' obtained by shifting Q. Note that u and u' are the only nodes of O that have missing colors c or d, since O is a hard orbit. Thus Q ends at node u' and the missing colors c and d of nodes u and u' were exchanged in G', in particular $c \in M'(u')$. Also note, that all edges marked in O or not contained in O remained unchanged in G'. Therefore O is still an edge orbit in G' and color c is still leaving at v. Now the induction hypothesis is applicable on the path $\bar{Q} := Q - u$. □

Lemma 10b which is proven in the appendix is similar to the `leave` routine in [5].

2.5 Algorithms

The following Proposition combines the tools introduced in the preceeding sections into an algorithm for producing a coloring without fat edges and where components of G_0 will turn out to be 'small'.

Proposition 11 (General Coloring Algorithm). *For a coloring G we can compute a coloring G^* such that every color orbit in G^* is strong, no edge in G^* is fat and during the computation of G^* the number q of colors used in G^* has only been increased if there was an $(\mathcal{A})$ or $(\mathcal{B})$ witness in some hard orbit $O \subseteq G'$ for some intermediate coloring G'.*

Proof. (By induction over the potential Φ of G.)

For $\Phi = 0$ the coloring G is complete and the proposition is trivially true. Obviously our proposition is correct, if there is no weak color orbit and no fat edge. But if there is a weak color orbit in G, then we can decrease the potential by Proposition 3 and the induction hypothesis becomes applicable.

Therefore suppose all color orbits are strong and e is a fat edge in G. Let O be the trivial edge orbit induced by $[e]$. By Proposition 9, we can increase the size of the orbit, until it is no longer hard or has a witness. In the case that the orbit is no longer a hard orbit, it either became a weak edge orbit or strong edge orbit and a weak color orbit, thus we can decrease the potential either by Proposition 7 or 3. In case of a witness we introduce a new color and can decrease the potential by assigning this color to some uncolored edge. In either case the induction hypothesis is applicable. □

Clearly the running time of the algorithm described above is in $\text{poly}(|E_0|, |V|, \Delta)$, if $|E_0|$ denotes the number of uncolored edges in G. The dependence on Δ stems from finding common missing colors and incomplete colors. For the special case of constant size strong color orbits it is worth having a closer look at the exact complexity of the algorithm since it turns out to match the complexity of previous algorithms with weaker approximation guarantee.

Proposition 12. *Under the assumption, that the size of a strong color orbit is always bounded by some constant, the time complexity of the algorithm in Proposition 11 is $\mathcal{O}(|E_0|(|V| + \Delta))$.*

Proof. Since $\Phi \leq 2|E_0|$, it suffices to show that the potential can be decreased in time $\mathcal{O}(|V| + \Delta)$. We use the collection $(G_c)_{c=1}^{q}$ of color classes and the graph G_0 of uncolored edges to represent the coloring G. Clearly we have $q \in \mathcal{O}(\Delta)$, therefore we can find missing colors and incomplete colors in $\mathcal{O}(\Delta)$. Assigning a color to an edge and uncoloring an edge can be done in constant time. Shifting an a,b-alternating path can be done in time proportional to the number of nodes in the path, since we only have to modify two matchings, G_a and G_b. Since we greedily eliminate weak color orbits in the algorithm, the maximum size of a weak color orbit considered in the algorithm is just one more than the maximum size of a strong color orbit.

We store a stack of fat edges, in order to be able to find an edge orbit in constant time. As long as it is hard, we can grow it by Proposition 9 in time $\mathcal{O}(|V|+\Delta)$, since we only have to perform a constant number of `shift` and 'color find' operations.

After a constant number of iterations of Proposition 9 there is a witness in the orbit or the orbit is no longer hard. In the first case we can reduce the potential in constant time. In the latter case we apply Proposition 3 or 7. In both propositions we only perform a constant number of `shift` and 'color find' operations.

If no more edges are fat, we compute the color orbits of each node. As soon as we found a weak color orbit, we use Proposition 3 to decrease the potential. Since all considered color orbits have constant size, time $\mathcal{O}(|V|+\Delta)$ is needed to decrease the potential.

Thus the total running time is $\mathcal{O}(|E_0|(|V|+\Delta)$. □

Now we relate properties of our orbit structures to the known lower bounds of χ'. This will finally enable us to design algorithms with guaranteed approximation ratios.

Lemma 13.
If O is a strong color orbit, then $|V(O)| \le \dfrac{q+2}{q-\Delta+2}$.

Proof. Since no two nodes share a missing color, we have $\sum_{u\in V(O)} |M(u)| \le q$. Obviously, every node in O has at least $q-\Delta$ missing colors. Since O is connected by uncolored edges, there are at least $|V(O)|-1$ uncolored edges in O and therefore at least $2(|V(O)|-1)$ additional missing colors. Thus we have, that the total number of missing colors is at least $|V(O)|(q-\Delta)+2(|V(O)|-1)$. □

Lemma 14. *Let O be a hard orbit.*

a) If there is an $(\mathcal{A})$ witness in O, then $q-\Delta+2 \le 2|V(O)|-4$.

b) If there is a $(\mathcal{B})$ witness in O, then $q-\Gamma+2 \le 2|V(O)|-4$.

Proof. As noted before the number of marked colors in an edge orbit O is at most $2|V(O)|-4$.

In a hard orbit every node is incident to at least two uncolored edges and at most $\Delta-2$ colored edges. Therefore every node in O has at least $q-\Delta+2$ missing colors. If there is an $(\mathcal{A})$ witness, then all missing colors of some node in O are marked in O, this implies $q-\Delta+2 \le 2|V(O)|-4$.

A hard orbit contains at least $|V(O)|$ uncolored edges and thus at most $|E(O)|-|V(O)|$ colored edges, i.e., there are at most $\frac{|E(O)|-|V(O)|}{\lfloor|V(O)|/2\rfloor} = \Gamma - \frac{|V(O)|}{\lfloor|V(O)|/2\rfloor} \le \Gamma-2$ complete colors and thus at least $q-\Gamma+2$ incomplete colors. And if O has a $(\mathcal{B})$ witness, then all incomplete colors of O are marked, implying $q-\Gamma+2 \le 2|V(O)|-4$. □

Lemma 15. *If $q \ge \lfloor(1+\epsilon)\Delta\rfloor - 1$ for some $\epsilon > 0$, then the following statements hold.*

a) If O is a strong color orbit, then $|V(O)| \le 1/\epsilon+1$

b) If there is a $(\mathcal{A})$ witness, then $q < \Delta + 2/\epsilon - 1$

c) If there is a $(\mathcal{B})$ witness, then $q < \Gamma + 2/\epsilon - 1$

Proof. By plugging $q \ge \lfloor(1+\epsilon)\Delta\rfloor - 1$ into the inequality of Lemma 13, we obtain

$$|V(O)| \overset{La.13}{\le} \frac{q+2}{q-\Delta+2} \le \frac{\lfloor(1+\epsilon)\Delta\rfloor+1}{\lfloor\epsilon\Delta\rfloor+1} \le \frac{\Delta}{\lfloor\epsilon\Delta\rfloor+1}+1 \le 1/\epsilon+1$$

If we plug this into the inequalities in Lemma 14, then we directly obtain the inequalities *b)* and *c)*. □

Theorem 16 (Algorithm I). *For every constant $\epsilon > 0$ there is an approximation algorithm for the multigraph edge coloring problem with time complexity $\mathcal{O}(|E|(|V|+\Delta))$ using at most $\max\{\lfloor(1+\epsilon)\Delta\rfloor + 1/\epsilon, \widetilde{\chi}' + 3/\epsilon\}$ colors.*

Proof. Start with $\lfloor(1+\epsilon)\Delta\rfloor - 1$ colors and apply Proposition 11 to obtain G'. The number of colors has only been increased, if there was some witness, i.e., if $q \overset{15b,15c}{<} \widetilde{\chi}'+2/\epsilon-1$ colors were available. Hence at most $\max\{\lfloor(1+\epsilon)\Delta\rfloor-1, \widetilde{\chi}'+2/\epsilon-1\}$ colors are used in G'.

No edge in G' is fat and by Lemma 15a all color orbits are strong and of size at most $1/\epsilon+1$. Using Vizing's algorithm and $1/\epsilon+1$ additional colors, we can now compute the desired complete coloring using a total of at most $\max\{\lfloor(1+\epsilon)\Delta\rfloor+1/\epsilon, \widetilde{\chi}'+3/\epsilon\}$ colors.

Since the size of a strong color orbit during the computation of G' was bounded by the constant $1/\epsilon+1$, the running time of the algorithm is by Proposition 12 $\mathcal{O}(|E|(|V|+\Delta)$. □

In the best case Algorithm I uses at least some $(1+\epsilon)\Delta$ colors. But from a practical point of view it may be worthwhile not to use that many colors in the beginning but to add colors in an adaptive manner.

Algorithm II relies on the following lemma.

Lemma 17. *If there is a witness in a hard color orbit O of G , then $q < \widetilde{\chi}' + \sqrt{2\widetilde{\chi}'} - 1$.*

Proof. Consider the following chain of inequalities

$$q - \widetilde{\chi}' + 2 \overset{La.14}{\leq} 2|V(O)| - 4 \overset{La.13}{\leq} 2\frac{q+2}{q-\widetilde{\chi}'+2} - 4.$$

For positive and integral q the solution is $q \leq \left\lfloor \widetilde{\chi}' + \sqrt{2\widetilde{\chi}'+1} \right\rfloor - 3$. Hence q fulfills the claimed inequality. □

The idea of Algorithm II is very simple. We start with Δ colors and simplify G_0 using Proposition 11. Thereafter we reduce the number of colors needed to color G_0 by iteratively adding new colors and applying Proposition 3. As soon as a stopping criterion is fulfilled, we stop adding colors and use Vizing's algorithm to compute a complete coloring. For the rather technical proof of the theorem refer to the appendix.

Theorem 18 (Algorithm II). *There is an approximation algorithm for the multigraph edge coloring problem with time complexity* $\mathrm{poly}(|V|, |\Delta|)$ *using at most* $\left(1+\sqrt{\frac{4.5}{\widetilde{\chi}'}}\right)\widetilde{\chi}'$ *colors.*

Note that the minimum of $(1+\epsilon)\widetilde{\chi}' + 3/\epsilon$ in ϵ is $\left(1+\sqrt{\frac{12}{\widetilde{\chi}'}}\right)\widetilde{\chi}'$, so the result of Theorem 18 is somewhat better than the naive approach.

3 A Polynomial Algorithm

In the following we will generalize the results of Section 2 to obtain a balancing algorithm that tolerates up to M uncolored parallel edges. This will be the main mechanism driving our polynomial algorithm.

For an arbitrary $M \in \mathbb{N}$ we partition the edges E of coloring G into three parts, namely

- the lean edges $E^{(<M)} := \{e \in E : |[e] \cap E_0| < M\}$,
- the even edges $E^{(=M)} := \{e \in E : |[e] \cap E_0| = M\}$,
- the fat edges $E^{(>M)} := \{e \in E : |[e] \cap E_0| > M\}$.

Now the potential $\Phi^{(M)}$ of coloring G is $\Phi^{(M)} := |E_0| + |E^{(>M)} \cap E_0|$. Note that all lemmata and propositions 1–12 are still true, if we just replace the old definitions of lean, even, and fat edges by these new ones. In the following we refer to orbits and witnesses with respect to the generalized definitions of lean, even and fat.

Now we refine the approximation lemmata of Section 2.5.

Lemma 19. *Let O be a hard orbit for some $M \in \mathbb{N}$*

a) If O has an $(\mathcal{A})$ witness, then $q - \Delta + 2M \leq 2|V(O)| - 4$.

b) If O has a $(\mathcal{B})$ witness, then $q - \Gamma + 2M \leq 2|V(O)| - 4$.

The proof of the lemma is very similar to the proof of Lemma 14. For completeness the proof of Lemma 19 can be found in the appendix.

Lemma 20. *If for $M := |V|$ a hard orbit O has some witness, then $q < \widetilde{\chi}'$.*

Proof. Trivially $|V(O)| \leq M$. And then the inequality follows from Lemma 19. □

In the following we contract consecutive colors with the same color class to color intervals, i.e. we represent a coloring G by a collection of matchings $(G_{I_k})_{k=1}^{I}$ where the $I_k = [a_k; b_k]$ are intervals of colors with the same color class and I is the number of these intervals. Then all 'color find' operations need time $\mathcal{O}(I)$. The `shift` and `color` operations need the same time than in the former representation, but may increase the number of intervals by at most a constant.

Lemma 21. *For any coloring G using at most $\widetilde{\chi}'$ colors contracted to I intervals we can compute a coloring G' 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.*

Proof. Let $M := |V|$. Then apply Proposition 11 on G to obtain a coloring G' with no fat edges. Therefore at most $M|V|^2 = |V|^3$ edges remain uncolored in G'.

The number of colors has only been increased, if there was a witness, i.e., $q \overset{La.19}{<} \widetilde{\chi}'$. Hence we still use at most $\widetilde{\chi}'$ colors.

The number of `shift` and `color` operations is polynomial in $|E_0|$ and $|V|$ and does not depend on I. Thus the number of intervals increased polynomially in $|E_0|$ and $|V|$.

Since the 'color find' operations can be done in $\mathcal{O}(I)$ we use total time of $\mathrm{poly}(|E_0|, |V|, I)$. □

We define the multiplicity-weighted adjacency matrix of the multigraph $G = (V, E)$ as $A = (|uv|)_{u,v\in V}$ For any function $f : \mathbb{N} \to \mathbb{N}$ and multigraph G the notation $f(G)$ means that f is applied on every entry of A. The notation $G + G'$ means the standard matrix addition of the multiplicity-weighted adjacency matrices A and A' of G respectively G'.

Proposition 22. *There is an algorithm with time complexity* $\mathrm{poly}(|V|, \log \mu)$ *that computes a coloring* G^* *of a multigraph* G *with maximum edge multiplicity* μ *such that at most* $|V|^3$ *edges of* G^* *are uncolored and at most* $\widetilde{\chi}'$ *colors are used in* G^*.

Proof. (By induction over μ of G.)

For $\mu = 0$ the graph contains no edges and our proposition is trivially true. Now suppose $\mu > 0$. We partition the input graph into three parts, i.e., $G = 2(\lfloor G/2 \rfloor) + (G \mod 2)$. Note that $\widetilde{\chi}'(G) \geq 2\widetilde{\chi}'(\lfloor G/2 \rfloor)$.

The algorithm recursively computes a coloring $\lfloor G/2 \rfloor^*$ of $\lfloor G/2 \rfloor$ that 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 $\lfloor G/2 \rfloor^*$, we obtain a coloring $2\lfloor G/2 \rfloor^*$ of $2\lfloor G/2 \rfloor$, that uses at most $2\widetilde{\chi}'(\lfloor G/2 \rfloor) \leq \widetilde{\chi}'(G)$ colors and has at most $2|V|^3$ uncolored edges. Obviously the number of intervals did not increase by this doubling.

Now we add the edges of the graph $(G \mod 2)$ to the coloring $2\lfloor G/2 \rfloor^*$ and obtain a coloring G' of G with at most $2|V|^3 + |V|^2$ uncolored edges.

The algorithm of Lemma 21 uses $\mathrm{poly}(|E_0'|, |V|, I) = \mathrm{poly}(|V|, I)$ time to color all but at most $|V|^3$ colors and increases the number of intervals polynomially in $|V|$. Let G^* be this new coloring.

Clearly the depth of recursion is $\mathcal{O}(\log \mu)$. In each recursive step the number of intervals increases 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 recursive step and therefore the total time is also $\mathrm{poly}(|V|, \log \mu)$. □

The running time of the following algorithm depends only logarithmically on μ and is therefore polynomial in the input size.

Theorem 23 (Polynomial Algorithm). *There is an approximation algorithm for the multigraph edge coloring problem with time complexity* $\mathrm{poly}(|V|, \log \mu)$ *using at most* $\left(1 + \sqrt{\frac{4.5}{\widetilde{\chi}'}}\right) \widetilde{\chi}'$ *colors.*

Proof. Use Proposition 22 and then apply Theorem 18. After application of Proposition 22 only $|V|^3$ uncolored edges remain and the number of intervals is polynomial in $|V|$ and $\log \mu$. In the coloring obtained by Proposition 22 at most $\widetilde{\chi}'$ colors are used. Therefore the algorithm in Theorem 18 runs in time polynomial in $|V|$ and $\log \mu$ to color the remaining uncolored edges and the number of used colors is then at most $\left(1 + \sqrt{\frac{4.5}{\widetilde{\chi}'}}\right) \widetilde{\chi}'$. □

Note that for $|V| \in \mathcal{O}(\log |\Delta|)$ the approximation ratio of this algorithm decreases exponentially in the size of the input.

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 [5, 8]. Our algorithms give better approximation except for graphs with very small maximum degree.

If one wants to implement our algorithm to solve real world instances, it would be interesting to add further heuristics. For example, Algorithm II from Section 2 could be refined such that before adding a fresh color, it first tries to color edges by swapping critical paths. 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., 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 by finding a maximal matching in G_0.

References

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

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

[3] U. Feige, E. Ofek, and U. Wieder. Approximating maximum edge coloring in multigraphs. In *APPROX*, volume 2462 of *LNCS*, pages 108–121. Springer, 2002.

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

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

[6] I. Holyer. The NP-completeness of edge-coloring. *SIAM Journal on Computing*, 10(4):718–720, 1981.

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

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

[9] P. Sanders and R. Solis-Oba. How helpers hasten h-relations. *Journal of Algorithms*, 41:86–98, 2001.

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

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

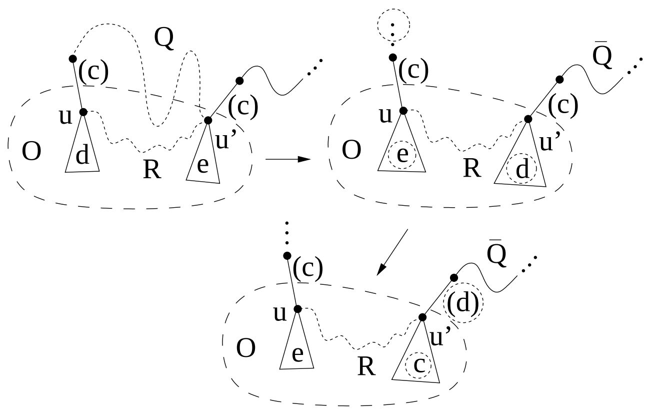

Figure 4: Lemma 10b

A Further Proofs

Proof (Lemma 10b): As in the proof of Lemma 10a we may assume that every node of O has at least one missing color not marked in O. Consider the alternating path $Q := \mathtt{Apath}(u, d, c)$ for some unmarked color $d \in M(u)$. We distinguish two cases.

1.) $V(O) \cap V(Q) = \{u\}$: If u is the only node of O contained in Q, we can shift Q to obtain a new coloring, such that color c is missing at u and still leaving at v, and apply Lemma 10a on this new coloring.

2.) $V(O) \cap V(Q) = \{u, ..., u'\}$: Let u' be the last node in Q that is still in O. Consider the alternating path $R := \mathtt{Apath}(u, d, e)$ for some unmarked color $e \in M(u')$. If R leaves O, then either d or e is leaving O and therefore Lemma 10a is applicable either on color d or e. So assume R does not leave O. Since O is a hard orbit, u and u' are the only nodes of O, that have missing colors d or e. Therefore R ends at u'. Now consider the coloring G' obtained by shifting R. In G' only edges contained in O and not marked in O were changed. Therefore O is an hard orbit in G' and the subpath $\bar{Q}$ of $Q \setminus (O - u')$ beginning at u' remained unchanged. Since $d \in M'(u')$, the alternating path $\mathtt{Apath}'(u', d, c)$ equals $\bar{Q}$ and therefore $\{u'\} = V(\bar{Q}) \cap V(O')$ and we are back to the first case. □

Proof (of Theorem 18): Start with Δ colors. Then compute a partial edge coloring of the input multigraph G by Proposition 11. Now every color orbit is strong an no edge is fat. Furthermore the number of colors has only been increased if there was some witness, i.e., if $q \overset{La.17}{<} (1 + \frac{\sqrt{2}}{\sqrt{\tilde{\chi}'}})\tilde{\chi}' - 1$.

Now iteratively add new colors and apply Proposition 3 until $q \geq \Delta + U$, where U is the number of colors Vizing's algorithm would use to color the current G_0.

Now we compute a complete coloring of G using Vizing's algorithm and U additional colors.

If $q \leq \Delta + \sqrt{\Delta}$, then

$$\begin{aligned} q + U &= q - \Delta + \Delta + U \\ &\overset{\Delta+U\leq q}{\leq} q - \Delta + q \overset{q\leq\Delta+\sqrt{\Delta}}{\leq} \Delta + 2\sqrt{\Delta} \end{aligned}$$

Otherwise,

$$q + U \leq q + \frac{q+2}{q-\Delta+2} \overset{q\leq\widetilde{\chi}'+\sqrt{2\widetilde{\chi}'}-1}{\leq} \widetilde{\chi}' + \sqrt{4.5\widetilde{\chi}'}$$

where the last inequality uses that $q + \frac{q+2}{q-\Delta+2}$ is monotonically increasing for $q > \Delta + \sqrt{\Delta}$. □

Proof (of Lemma 19): In a hard orbit O every node u has at least $q - \Delta + 2M$ missing colors, since u is connected to at least two neighbors by even or fat edges and thus is incident to at least $2M$ uncolored and at most $\Delta - 2M$ colored edges.

Furthermore, at least $M|V(O)|$ edges are uncolored in O and thus at most $\frac{|E(O)|-M|V(O)|}{\lfloor |V(O)|/2 \rfloor} \leq \Gamma - \frac{M|V(O)|}{|V(O)|/2} = \Gamma - 2M$ colors do not leave O, i.e., at least $q - \Gamma + 2M$ colors are leaving O. As noted before, at most $2|V(O)| - 4$ colors are marked in O.

By definition of the witnesses, we know that if there is an $(\mathcal{A})$ witness, then for some u all missing colors are marked in O, implying $q - \Delta + 2M \leq 2|V(O)| - 4$. If there is a $(\mathcal{B})$ witness, then all leaving colors are marked in O, so that $q - \Gamma + 2M \leq 2|V(O)| - 4$. □