Approximation Algorithms for Bounded Color Matchings via Convex Decompositions*

Georgios Stamoulis[1,2]

[1] LAMSADE, CNRS UMR 7243, Universitè Paris-Dauphine, France
[2] Universitá della Svizzera Italiana (USI), Lugano, Switzerland
stamoulis.georgios@dauphine.fr

Abstract. We study the following generalization of the maximum matching problem in general graphs: Given a simple non-directed graph $G = (V, E)$ and a partition of the edges into k classes (i.e. $E = E_1 \cup \cdots \cup E_k$), we would like to compute a matching M on G of maximum cardinality or profit, such that $|M \cap E_j| \leq w_j$ for every class E_j. Such problems were first studied in the context of network design in [17]. We study the problem from a linear programming point of view: We provide a polynomial time $\frac{1}{2}$-approximation algorithm for the weighted case, matching the integrality gap of the natural LP formulation of the problem. For this, we use and adapt the technique of *approximate convex decompositions* [19] together with a different analysis and a polyhedral characterization of the natural linear program to derive our result. This improves over the existing $\frac{1}{2}$, but with additive violation of the color bounds, approximation algorithm [14].

1 Introduction

In modern optical fiber network systems, we encode the information as an electromagnetic signal and we transfer it through the optical fiber as a beam of light in a specified frequency. Typically, at most one beam of light is allowed to travel through the fiber at any given time. In WDM[1] optical networks we allow *multiplexing* of a number of different light beams to travel simultaneously through the optical fiber as follows: We partition the electromagnetic spectrum into a number of k non-overlapping intervals. For each interval f_i we have an upper bound on how many different beams of light that have frequencies within this interval can travel at the same time through the optical carrier. This constraint is imposed since, otherwise, we would have quantum phenomena as *interference* of the light beams of a given interval. Naturally, if we allow a large number of beams of light with frequencies within a small interval to travel through the optical fiber, we can expect with very high probability two or more beams to be interfered. Our goal in an optical network is to establish communication between an as large as

* Part of this work was done while the author was a PhD student at IDSIA.
[1] WDM stands for Wavelength-Division Multiplexing

E. Csuhaj-Varjú et al. (Eds.): MFCS 2014, Part II, LNCS 8635, pp. 625–636, 2014.

possible number of pairs that want to communicate in their own frequency, such that in a given interval of frequencies f_i we allow no more than w_i connections to be established. This naturally reduces to the following problem:

Bounded Color Matching: We are given a (simple, non-directed) graph $G = (V, E)$. The edge set is partitioned into k sets $E_1 \cup E_2 \cup \cdots \cup E_k$ i.e. every edge e has color c_j if $e \in E_j$ and a profit $p_e \in \mathbb{Q}^+$. We are asked to find a maximum (weighted) matching M such that in M there are no more that w_j edges of color c_j, where $w_j \in \mathbb{Z}^+$ i.e. a matching M such that $|M \cap E_j| \leq w_j, \forall j \in [k]$.

In the following, we denote as $\mathcal{C}$ the collection of all the color classes. In other words, $\mathcal{C} = \{c_j\}_{j \in [k]}$. Moreover, for a given edge $e \in E(G)$, we denote by $c^{-1}(e)$ its color i.e. $c^{-1}(e) = c_j \Leftrightarrow e \in c_j$.

Bounded Color Matching is a budgeted version of the classical matching problem: For a given instance G, let $\mathcal{F}$ be the set of all feasible solutions. Associated with every feasible solution $M \in \mathcal{F}$ we are given a set of ℓ linear *cost* functions $\{\alpha_i\}_{i \in [\ell]}$ and a linear profit function π such that $\pi, \alpha_i : \mathcal{F} :\rightarrow \mathbb{Q}^+$ and for every cost function α_i a *budget* $\beta_i \in \mathbb{Q}^+$. The goal is to find $M \in \mathcal{F} : \alpha_i(F) \leq \beta_i,\ \forall i \in [\ell]$ that also maximizes $\pi(M)$. Budgeted versions of the maximum matching problem have been recently studied intensively. When G is bipartite there is a PTAS for the case where $\ell = 1$ [2] and the case where $\ell = \mathcal{O}(1)$ [11]. For general graphs there is a PTAS for the 2-budgeted maximum matching problem [12] and a bicriteria PTAS for $\ell = \mathcal{O}(1)$ [7] (where the returned solution might violate the budgets by a factor of $(1 + \epsilon)$). This approach works also for *unbounded* number of budgets albeit a *logarithmic* overflow of the budgets.

Bounded Color Matching (BCM) not only is **NP**-hard even in bipartite graphs where $w_j = 1,\ \forall c_j \in \mathcal{C}$ [10] but also **APX**-hardness can be deduced even in 2-regular bipartite graphs from [15]. In [13] the BCM was considered from a bi-criteria point of view: given a parameter $\lambda \in [0, 1]$ there is an $(\frac{2}{3+\lambda})$ approximation algorithm for BCM which might violate the budgets w_j by a factor of at most $(\frac{2}{1+\lambda})$.

To the best of our knowledge, the first case where matching problems with cardinality (disjoint) budgets were considered, was in [17] where the authors defined and studied the *blue-red* Matching problem: compute a maximum (cardinality) matching that has at most w blue and at most w red edges, in a blue-red colored (multi)-graph. A $\frac{3}{4}$ polynomial time combinatorial approximation algorithm and an $\mathbf{RNC^2}$ algorithm (that computes the maximum matching that respects both budget bounds with high probability) were presented. This was motivated by network design problems, in particular they showed how *blue-red* Matching can be used in approximately solving the Directed Maximum Routing and Wavelength Assignment problem (DirMRWA) [16] in *rings* which is a fundamental network topology, see [17] (also [4] for alternative and slightly better approximation algorithms and [1] for combinatorial algorithms). Here, approximately solving means that an (asymptotic) α-approximation algorithm for blue-red results in an (asymptotic) $\frac{\alpha+1}{\alpha+2}$-approximation algorithm for DirMRWA in rings.

We note that the exact complexity of the blue-red matching problem is not known: it is only known that blue-red matching is at least as hard as the Exact

Matching problem [18] whose complexity is open for more than 30 years. A polynomial time algorithm for the blue-red matching problem will imply that Exact Matching is polynomial time solvable. On the other hand, blue-red matching is probably not **NP**-hard since it admits an $\mathbf{RNC^2}$ algorithm. We note that this algorithm can be extended to a constant number of color classes with arbitrary bounds w_j. Using the results of [20] one can deduce an "almost" optimal algorithm for blue-red matching, i.e. an algorithm that returns a matching of maximum cardinality that violates the two color bounds by at most one edge. This is the best possible, unless of course blue-red matching (and, consequently, exact matching) is in **P**.

If we formulate BCM as a linear program, the polyhedron $\mathcal{M}_c$ containing all feasible matchings M for the BCM is

$$\mathcal{M}_c = \left\{ \boldsymbol{y} \in \{0,1\}^{|E|} : \quad \boldsymbol{y} \in \mathcal{M} \quad \bigwedge \quad \sum_{e \in E_j} y_e \leq w_j, \quad \forall j \in [k] \right\} \tag{1}$$

where $\mathcal{M}$ is the usual matching polyhedron: $\mathcal{M} = \{x \in \{0,1\}^{|E|} : \quad \sum_{e \in \delta(v)} x_e \leq 1, \ \forall v \in V\}$. We would like to find $\boldsymbol{y} \in \{0,1\}^{|E|}$ such that $\boldsymbol{y} = \max_{x \in \mathcal{M}_c}\{\boldsymbol{p}^T \boldsymbol{x} = \sum_{e \in E} p_e x_e\}$, $\boldsymbol{p} \in \mathbb{Q}^{|E|}_{\geq 0}$. As usual, we relax the integrality constraints $\boldsymbol{y} \in \{0,1\}^{|E|}$ to $\boldsymbol{y} \in [0,1]^{|E|}$ and we solve the corresponding linear relaxation efficiently to obtain a *fractional* $|E|$-dimensinal vector $\boldsymbol{y}$. It is not hard to show (see later section) that the *integrality gap* of $\mathcal{M}_c$ is essentially 2 and this is true even if we add the blossom inequalities i.e. if instead of $\mathcal{M}$ as defined here, we use the well known Edmond's LP [8].

Our Contribution: In this work we study the BCM problem with *unbounded* number of budgets: we provide a deterministic $\frac{1}{2}$ approximation algorithm based on the concept of *approximate convex decompositions* from [19] together with a different analysis and an extra step based on polyhedral properties of extreme point solutions of $\mathcal{M}_c$. This might be helpful also in the context of k-uniform b-matching problem. This result improves over the $\frac{1}{2}$ but with an additive violation of the color bounds w_j from [13] and matches the integrality gap of 2 of the natural linear formulation of the problem (captured by 1) which implies that a $\frac{1}{2}$ approximation algorithm is the best we can hope using this natural linear relaxation.

We note that the BCM problem can be easily seen as a special case of the 3-hypergraph β-matching problem [19] or 3-set packing. Using the existing LP-based results for these problems which guarantee a $\frac{1}{k-1+\frac{1}{k}}$ for the k-hypergraph β-matching, we can only guarantee a $\frac{3}{7}$-approximation algorithm [19]. We show that by taking advantage of the special structure of the problem we can do better than this. For 3-set packing there there exists a $\frac{1}{2} - \epsilon$ approximation algorithm for the *weighted* case [3] and a recent $\frac{3}{4}$ for the non-weighted case [9]. But these tell us nothing about the strengths (and limitations) of linear programming techniques for the problem, which is our main motivation. See also [6] for a similar study on the effect of linear programming techniques on k-dimensional matching problems.

2 A $\frac{1}{2}$ Approximation Based on Approximate Convex Decompositions

In this section we will provide a polynomial time $\frac{1}{2}$-approximation algorithm for BCM based on the technique of approximate convex decompositions from [19]. Given x^*, an optimal solution of the LP for the BCM, the main idea of the algorithm is to construct a collection of *feasible* (for the BCM problem) matchings $\mu_1, \mu_2, \ldots, \mu_\rho$ for some ρ, such that an approximate version of $p^T x$ can be written as a *convex combination* of these matchings μ_i. Recall the famous Carathéodory's theorem:

Theorem 1 (Carathéodory[5]). *Let $P = \{x \mid Ax \leq b\}$ where $x \in \mathbb{R}^n, A \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^m$ be a polyhedron in $\mathbb{R}^n$. Assume a point $z \in P$ and assume that this point satisfies r of the m inequalities with equality. Then z can be written as a convex combination of at most $n - r + 1$ vertices of the polyhedron P.*

As an immediate corollary, we have that any point z belonging in (the convex hull of) a bounded convex polyhedron $P \subseteq \mathbb{R}^n$ can be written as a convex combination of at most $n + 1$ vertices of P.

Let x^* be an optimal (fractional) solution of the relaxation of $\mathcal{M}_c$. Ideally we would like to use Carathéodory's theorem to write x^* as a *convex* combination of feasible integral extreme point solutions (vertices) of $\mathcal{M}_c$. But unfortunately this is not always the case, meaning that x^* might not belong in the convex hull of all feasible integral vertices of $\mathcal{M}_c$. Instead of that, we will settle for an *approximate* convex combination of x^* by vertices of $\mathcal{M}_c$. An approximate convex combination of x^* is a convex combination of ρ extreme points μ_i of $\mathcal{M}_c$ satisfying the following:

$$\alpha p^T x^* = \sum_{i\in[\rho]} \lambda_i \mu_i, \quad \alpha \in (0, 1], \quad \sum_{i\in[\rho]} \lambda_i = 1, \quad \mu_i \in \mathcal{M}_c \tag{2}$$

The fact that we insist for convex combination directly implies that $\lambda_i \geq 0, \forall i$. An important feature of the above convex combination is that it gives us immediately an α approximation algorithm. Indeed, since all points μ_i, for $i \in [\rho]$, are feasible and they constitute a *convex* combination of $p^T x^*$ then, by a standard averaging argument, at least one of the μ_is will have profit at least $\alpha \cdot p^T x^*$, and so we have the following:

Lemma 1. *Given an optimal fractional solution x^* of the relaxation of $\mathcal{M}_c$ and assume that x^* can be written as in (2) then we can retrieve a feasible integral solution for $\mathcal{M}_c$ with total profit at least α times $p^T x^*$.*

We will inductively construct an α approximate convex combination of $\bar{x}^*$ where $\bar{x}^*$ is x^* without a specific edge (i.e. after removing an edge e in $\mathsf{supp}(x^*)$ where we define for any vector $y \in \mathbb{R}^n$ $\mathsf{supp}(y) = \{j \in [n] : y_j \neq 0\}$). Then, if we can add e in a α fraction of the points $\bar{\mu}_i$ constituting the α approximate convex combination of $\bar{x}^*$, we will have a convex combination with the desired properties. In other words, we require that $\sum_{j\in\kappa}\lambda_j \geq \alpha x^*_e$, where $\kappa = \{\mu_i : \mu_i \cup e \in \mathcal{M}_c\}$ is the set of all the solutions (that constitute the approximate version of $\bar{x}^*$) that can facilitate e preserving feasibility.

The inductive process of obtaining (constructing) μ_is is roughly the following: at the "bottom" of the induction process we start with the trivial empty solution ($\mu_1 = \emptyset$). At each next step we try to pack any of the edges e_j into exactly an α fraction of the current set of the matchings. This is the step where we may create new solutions in order to maintain the invariant "pack any edge into exactly α fraction". Usually this process in its basic form greedily packs e_j and, moreover, is oblivious to any ordering of the edges. In our case however, this cannot lead to any meaningful approximation guarantee. We will show how we can carefully select an edge at the current step i such that edge e_i to be packed can fit into an $\frac{1}{2}$ fraction of the current set of solutions $\{\mu_j\}$. In other words, we will define an ordering of the edge set $(e_1,\ldots,e_m)$ $(m = |E(G)|)$ such that at step $i \in [m]$, assuming that we have an α-approximate convex decomposition for $(e_1,\ldots e_{i-1})$ (characterized by a set of feasible matchings $\{\mu_j\}$), then the current edge e_i can be successfully inserted into at least an α fraction of the μ_js. We will show that in our case we can in fact select $\alpha = \frac{1}{2}$, thus giving us the desired approximation guarantee.

Now we will define the ordering on the edges $\{e_i\}_{i\in[m]}$. Assume that we are in some inductive step where the remaining edges are $R_i = \{e_1, e_2, \ldots, e_i\}$ for some i, i.e. we have removed $m - i$ edges to iteratively obtain an α-approximate decomposition for the remaining solution. How do we choose which edge to select from R_i? Intuitively, the larger the fractional value of x_e is, the larger the fraction of matchings μ_i that can be added is. This is because high fractional value of x_e implies low fractional values of the other "blocking" components. And low value of these components means "few" matchings μ_j that are actually blocking e. So, a good starting strategy is to select at each step edges with high fractional value. Unfortunately, we cannot always guarantee that such edges are present in x^*. But there is hope: let x' be x^* restricted to the set of edges R_i. Now, it is not hard to see (we omit the easy proof) that x' is an extreme point solution for the *reduced* instance where we set $w_j := w_j - \sum_{e\in E_j\setminus R_i} x^*_e$, $\forall C_j$ and $\beta_v := \sum_{e\in\delta_v\setminus R_i} x^*_e$, $\forall v \in V$. Such extreme point solutions have very nice properties. To see that, we use the following slight generalization of a result due to [13] where we consider a version where for each vertex $v \in V$ we have a bound $\beta_v \leq 1$ and where w_j are no longer integers such that, when we say "tight" vertex (with respect to x^*) we mean a vertex v such that $\sum_{e\in\delta(v)} x_e = \beta_v$. The same for tight color classes. Initially, all w_j are integers and $\beta_v = 1$, $\forall v$. Define the residual graph with respect to a solution vector x to be the graph where we include an edge e only if $x_e > 0$.

Lemma 2. *Take any basic feasible solution x of $\mathcal{M}_c$ (where we no longer require the degree bounds on vertices and the bounds on color classes be integer anymore) such that $0 < x_e < 1$, $\forall e$ (i.e. remove all integer variables reducing the bounds appropriately if necessary). Then one the following must be true:*

1. *either there is a tight color class* $c_j \in Q$ *such that* $|E_j| \leq \lceil w_j \rceil + 1$ *in the residual graph,*
2. *or there is a tight vertex* $v \in F$ *such that the degree of* v *in the residual graph* $\leq \lceil \beta_v \rceil + 1$.

The proof follows similar lines as in [13] and we omit it from the current version. In our lemma, F, Q are the linearly independent set of rows of the LP $(\mathcal{M}_c)$ that characterize the basic solution x. This lemma will give us what we need in order to successfully select the right edge from R_i and prove that we can successfully pack it into $\frac{1}{2}$ of the matchings for the inductively obtained approximate convex decomposition for the rest of the edges i.e. for R_{i-1}:

1. Given x^* select e according to Lemma 2:
   (a) if $\exists v \in F$: $|\mathsf{supp}(x^* \cap \delta(v))| \leq 2$ then select as e the edge $\in \delta(v) : x_e \geq \frac{\beta_v}{2}$.
   (b) else $\exists c_j \in Q : |\mathsf{supp}(x^*) \cap E_j| \leq \lceil w_j \rceil + 1$. Select in this iteration an edge $e \in c_j$ with $x^*_e \geq \frac{w_j}{\lceil w_j \rceil + 1}$.
2. Zero out the coordinate of x^* corresponding to e and let x' the resulting vector.
3. Iteratively obtain an approximate convex decomposition for x'.
4. Add e to the convex decomposition of x' obtained in the previous step.

Using Lemma 2 we will show now how, in each iteration of our algorithm, we can successfully pack an edge $e = \{u, v\} \in \mathsf{supp}(x^*)$ into a large number (i.e. half) of the solutions that constitute the approximate convex decomposition of the residual solution vector x'. In order to insert e into a large number of the solutions μ_i that constitute the approximate convex decomposition of the residual solution vector x', we need to see in what fraction of the μ_i's the edge e *cannot* be added. These are all μ_i's such that

(**1**): $\exists e' \in \mu_i$: $u \vee v \in e'$, or
(**2**): $|\mu_i \cap E_j| = w_j$, where $j \in [k]$ is the color of edge e.

The first condition says that e *cannot* be added to those matchings μ_i (which constitute that approximate convex decomposition of x') that have edges incident to either of the endpoints of e. The second condition says that, additionally, e cannot be added to all μ_is that are "full" of color c_j. All these matchings are blocking the insertion of e, meaning that for such a μ_i, $\mu_i \cup e$ is not feasible anymore for either of the previous two reasons. In order to guarantee that e can be added to exactly α fraction of the matchings we may need to double a current solution μ_i and break its multiplier λ_i appropriately (see appendix).

We will distinguish between two cases (one for each case of the algorithm, i.e. step 1.(a) or 1.(b)) and we will prove the result inductively. The base case of the induction inside the algorithm is the trivial case of the empty graph. Focus, w.l.o.g., at the first execution of the algorithm and assume that we have an α-approximate convex decomposition of x' (x^* without edge e) i.e. $\alpha p^T x' = \sum_i \lambda_i \mu_i$ where μ_i's $\in \mathcal{M}_c$ and $\mathbf{1}^T \boldsymbol{\mu} = 1$. Moreover, by Carathéodory's Theorem, this collection of matchings is *sparse* (at most $|\mathsf{supp}(x')| + 1$ matchings μ_i). Let u, v the endpoints of e selected in the first step of the algorithm and assume that $c^{-1}(e) = j \in [k]$. Firstly, we will handle the case where the edge e has been selected according to the rule 1.(b) of the algorithm, which is slightly easier.

According to the algorithm (using Lemma 2), we know that there must exist a tight color class $c_j \in \mathcal{C}$ with the property that

$$|\mathsf{supp}(x^*) \cap E_j| \leq \lceil w_j \rceil + 1 \Rightarrow \exists e \in \mathsf{supp}(x^*) \cap c_j : \; x_e^* \geq \frac{w_j}{\lceil w_j \rceil + 1}$$

Lemma 3. *If we select to pack an edge e of color c_j according to rule 1.(b) of the algorithm, then the fraction of the solutions μ_i that e can be added is at least $\frac{1}{2}$ i.e. $\sum_{j \in \kappa_e} \lambda_j \geq \frac{1}{2}$.*

Proof. We will show that at least $\frac{1}{2}$ fraction of the solutions μ_i can facilitate such an e. Assume that we can identify a color class $c_j \in Q$ in the *reduced instance* induced by the current set of edges such that for these reduced color bounds we have that $|\mathsf{supp}(x^*) \cap E_j| \leq \lceil w_j \rceil + 1$. Let $e = \{u, v\}$ be an edge from $\mathsf{supp}(x^*) \cap E_j$. Let $\xi \in \mathsf{supp}(x^*)$ be the index of e. Observe that in the reduced solution x' (defined as x^* without e), there are exactly $\lceil w_j \rceil$ edges of color c_j (summing up to $w_j - x_e^*$) and, moreover, this number $\lceil w_j \rceil$ is *at most* the initial (integer) color bound for E_j.

All the solutions μ_i that constitute the α-approximate decomposition for the residual solution x' that block the insertion of e (i.e. the solutions μ_i such that $\mu_i \cup \{e\}$ is not feasible anymore) are these μ_is that have edges adjacent to either of u, v and those that have $\lceil w_j \rceil$ edges of color j and so we have that in an α-approximate decomposition, the fraction of μ_i's that block e is

$$\alpha(1 - x_e) + \alpha(1 - x_e) + \alpha(\frac{w_j - x_e}{\lceil w_j \rceil}) = A$$

and so if we would like to pack edge e to an α fraction of the μ_i's that constitute the α approximate decomposition of x', then we must require that $1 - A \geq \alpha x_e^*$. From this we conclude that the fraction of the solutions that e can be added is

$$1 - \Big(\alpha(1 - x_e) + \alpha(1 - x_e) + \alpha(\frac{w_j - x_e}{\lceil w_j \rceil})\Big)$$

from which we get that $\alpha \leq \frac{1}{2 + \frac{w_j}{\lceil w_j \rceil} - x_e(1 + \frac{1}{\lceil w_j \rceil})} = \frac{1}{\sigma}$. We will deliver an upper bound on σ. Using the bound on the variable x_e, we have that

$$\begin{aligned} \sigma &\leq 2 + \Big(\frac{w_j}{\lceil w_j \rceil} - \frac{w_j}{\lceil w_j \rceil + 1}(1 + \frac{1}{\lceil w_j \rceil})\Big) \\ &= 2 + \Big(\frac{w_j}{\lceil w_j \rceil} - \frac{w_j}{\lceil w_j \rceil + 1} - \frac{w_j}{\lceil w_j \rceil} \cdot \frac{1}{\lceil w_j \rceil + 1}\Big) \\ &= 2 + \Big(\frac{(\lceil w_j \rceil + 1)w_j - w_j \lceil w_j \rceil - w_j}{\lceil w_j \rceil(\lceil w_j \rceil + 1)}\Big) \;=\; 2 \end{aligned}$$

and so $1/\sigma \geq 1/2$ so we can select $\alpha = \frac{1}{2}$ to satisfy the bound on α delivered above i.e. $\alpha \leq \frac{1}{\sigma}$.

Now we move on to prove the second case: if we select an edge e (with endpoints $u, v \in V$ and color $j \in [k]$) according to rule 1.(a) of the algorithm, then this edge can be packed again into an α fraction of the solutions (matchings) that constitute an α-approximate decomposition of the subgraph induced by the current residual solution x' without the edge e. For this, we need to define the set of the solutions μ_i that block the insertion of e slightly more carefully.

Let B_e be the set of all solutions μ_i that constitute an α-approximate decomposition of x' (current solution without edge e) that block the insertion of e. In B_e, as before, we add all μ_i that have edges adjacent to either of the endpoints u, v of e. We need to describe which solutions μ_i are blocking solutions for e with respect to its color class c_j. The natural way is to consider the solutions that are "full" of color c_j i.e. have w_j edges of color c_j (condition (2)). Unfortunately, if we follow this rule, the result would be a slightly worse approximation guarantee than our goal i.e. we can guarantee that e can be packed into a $\frac{2}{5}$ fraction of the μ_i's resulting in a $\frac{2}{5}$ approximation guarantee.

Instead, we will define blocking solutions of edge e, with respect to color c_j, all the solutions μ_i such that $|\mu_i \cap E_j| = \lceil \sum_{e' \in E_j} x^*_{e'} \rceil$ (in the subgraph induced by the current solution vector x^*).

Lemma 4. *Let $x \in [0,1]^E$ be any fractional feasible solution of the natural LP for the BCM problem and let $\theta_j = \lceil \sum_{e \in E_j} x_e \rceil \leq w_j$, for every color class $c_j \in \mathcal{C}$. Assume that we have an α-approximate decomposition I for x, i.e. a collection of feasible matchings $\{\mu_i\}_{i \in I}$ together with their multipliers such that $\alpha \cdot x = \sum_i \lambda_i \mu_i$. For a color class c_j define $W(j) = \{\mu_i : |\mu_i \cap E_j| = \theta_j\}$. Then we have that*

$$\Lambda(W_j) \;=\; \sum_{\mu_i \in W(j)} \lambda_i \;\leq\; \alpha\Big(1 + \mathcal{F}\Big(\sum_{e' \in E_j} x_{e'}\Big)\Big)$$

where $\mathcal{F}(y) = y - \lceil y \rceil$ for $y > 0$ is the fractional part of y. Moreover, let $e = \{u, v\}$ be an edge with $c(e) = c_j, j \in [k]$ such that $x_e > 0$. Let $G[x'] := G[x] \setminus \{e\}$ and assume that we have an α approximate convex decomposition for $G[x']$ for some $\alpha \in (0,1)$. Define the set of blocking solutions (with respect to the approximate decomposition of x') for e due to color c_j as:

$$B_e(j) = \left\{ \mu_i : |\mu_i \cap E_j| = \lceil \sum_{e' \in E_j} x_{e'} \rceil \right\}$$

Then we have that

$$\sum_{i : \mu_i \in B_e(j)} \lambda_i \;\leq\; \alpha(1 - x_e)$$

Proof. We will prove the first claim with induction on x. The second claim will follow easily from the first one. For the base case assume that $x = 0^E$. Then the condition is automatically satisfied (take any empty matching with multiplier equal to 1). In this case $W_j = \emptyset$, for all color classes c_j.

Assume that the claim is true for all but the *first* edge of color c_j to be removed (first with respect to the inductive process of obtaining approximate decompositions), i.e., assume that it is true for the subgraph $G[x'] := G[x] \setminus \{e\}$. Let e_f be that edge. Define $X_{\bar{e}_f} = \sum_{e \in E_j \setminus \{e_f\}} x_e$ in the current subgraph. We will distinguish between two cases:

First case: $\lceil X_{\bar{e}_f} \rceil = \lceil X_{\bar{e}_f} + x_{e_f} \rceil$**:** In this case we have that

$$\begin{aligned}
\sum_{j \in W(j)} \lambda_j \;&\leq\; \sum_{\mu_\gamma \in \Gamma} \lambda_{\mu_\gamma} + \alpha x_{\bar{e}_f} \;\leq\; \alpha\Big(1 + \mathcal{F}\Big(X_{\bar{e}_f}\Big)\Big) + \alpha x_{\bar{e}_f} \;= \\
&= \alpha\Big(1 + \sum_{e \in E_j \setminus \{e_f\}} x_e - \lceil \sum_{e' \in E_j \setminus \{e_f\}} x_{e'} \rceil + x_{e_f}\Big) \\
&= \alpha\Big(1 + \sum_{e \in E_j} x_e - \lceil X_{\bar{e}_f} \rceil\Big) \\
&= \alpha\Big(1 + \sum_{e \in E_j} x_e - \lceil X_{\bar{e}_f} + x_{e_f} \rceil\Big) \;=\; \alpha\Big(1 + \sum_{e \in E_j} x_e - \lceil \sum_{e \in E_j} x_e \rceil\Big) \\
&= \alpha\Big(1 + \mathcal{F}\Big(\sum_{e' \in E_j} x_{e'}\Big)\Big)
\end{aligned}$$

where $\Gamma = \{\mu_\gamma : |\mu_i \cap E_j| = \lceil X_{\bar{e}_f} \rceil\}$ for the feasible matchings μ_γ that constitute the approximate decomposition of the current subgraph $G[x']$ (without the edge e_f). The first inequality is true because we have inserted an α fraction of e_f into the current approximate decomposition. The second inequality follows by the inductive hypothesis.

Second case: $\lceil X_{\bar{e}_f} \rceil \neq \lceil X_{\bar{e}_f} + x_{e_f} \rceil$**:** Observe that in this case we have that $\lceil X_{\bar{e}_f} \rceil = \lceil X_{\bar{e}_f} + x_{e_f} \rceil - 1$. Assume that we have an α-approximate decomposition for the subgraph induced by all the edges except e_f. Now, the set of matchings that block the insertion of edge e_f to the current approximate decomposition contains all the matchings μ_i such that $|\mu_i \cap E_j| = \lceil \sum_{e \in E_j} x_e \rceil$ (for the current set of edges of color c_j). But observe that since $\lceil X_{\bar{e}_f} \rceil \neq \lceil X_{\bar{e}_f} + x_{e_f} \rceil = \lceil \sum_{e \in E_j} x_e \rceil$, no matching μ_i from the current approximate decomposition has this property (that $|\mu_i \cap E_j| = \lceil \sum_{e \in E_j} x_e \rceil$). If either of ax_{e_f} or $\Lambda(W_j)$ (for the subgraph induced by all edges but e_f) is less than $\alpha(1 + \mathcal{F}(\sum_{e' \in E_j} x_{e'}))$, then we are done. Otherwise we might need to duplicate some solutions μ_i that constitute the approximate convex decomposition of x' to make sure that $\Lambda(W_j) = \alpha(1 + \mathcal{F}(\sum_{e' \in E_j} x_{e'}))$.

We need to prove the second claim i.e., $\sum_{i:\mu_i \in B_e(j)} \lambda_i \leq \alpha(1 - x_e)$. In the first case ($\lceil X_{\bar{e}_f} \rceil = \lceil X_{\bar{e}_f} + x_{e_f} \rceil$), observe that $B_e(j)$ does not change after the removal of e_f, i.e., in the subgraph induced by the remaining edges, the matchings μ_i with the property $|\mu_i \cap E_j| = \lceil \sum_{e' \in E_j} x_{e'} \rceil$ are the same in both cases. So we have that

$$\begin{aligned}\sum_{i:\mu_i \in B_e(j)} \lambda_i = \Lambda(W_j) &= \sum_{\mu_i \in W(j)} \lambda_i \leq \alpha\Big(1 + \mathcal{F}\Big(\sum_{e' \in E_j \setminus \{x_{e_f}\}} x_{e'}\Big)\Big) \\ &= \alpha\Big(1 + \sum_{e \in E_j \setminus \{e_f\}} x_{e'} - \lceil \sum_{e \in E_j \setminus \{e_f\}} x_{e'} \rceil\Big) \\ &= \alpha\Big(1 + \sum_{e \in E_j} x_{e'} - x_{e_f} - \lceil X_{e_f} + x_{e_f} \rceil\Big) \\ &= \alpha\Big(1 - x_{e_f} + \underbrace{\sum_{e \in E_j} x_{e'} - \lceil X_{e_f} + x_{e_f} \rceil}_{\leq 0}\Big) \\ &\leq \alpha(1 - x_{e_f})\end{aligned}$$

where the quantities $\Lambda(W_j)$ and $\sum_{\mu_i \in W(j)} \lambda_i$ are defined in the subgraph without e_f and with respect to the α-approximate decomposition defined by x'.

As for the second case, we already claimed that in an α-approximate decomposition I for $G[x'] = G[x] \setminus \{e_f\}$, no solution $\mu_i \in I$ can block the insertion of e_f because in $G[x']$ we have that $\lceil \bar{X}_{e_f} \rceil = \lceil X_{\bar{e}_f} + x_{e_f} \rceil - 1$ and by construction in I we do not have any μ_i such that $|\mu_i \cap E_j| = \lceil \bar{X}_{e_f} \rceil + 1$ to block the insertion of e_f. Now, as we already argued, if $\Lambda(W_j)$ (in $G[x']$ which to avoid confusion we denote as $\Lambda_{x'}(W_j)$is less than $\alpha(1 + \mathcal{F}(\sum_{e' \in E_j} x_{e'}))$ then $B_{e_f}(j) = \emptyset$ in this case and we are done. But in the case that $\Lambda(W_j)$ (in $G[x']$)is strictly more than $\alpha(1 + \mathcal{F}(\sum_{e' \in E_j} x_{e'}))$, then we should set $B_{e_f}(j) \subseteq \Lambda_{x'}(W_j)$ such that

$$\begin{aligned}\sum_{i \in B_{e_f}(j)} &= \Lambda_{x'}(W_j) - \alpha(1 + \mathcal{F}(\sum_{e' \in E_j} x_{e'})) \\ &\leq \alpha\Big(1 + \bar{X}_{e_f} - \lceil \bar{X}_{e_f} \rceil\Big) - \alpha\Big(1 + \bar{X}_{e_f} + x_{e_f} - \lceil \bar{X}_{e_f} + x_{e_f} \rceil\Big) \\ &= \alpha\Big(- x_{e_f} \underbrace{- \lceil \bar{X}_{e_f} \rceil + \lceil \bar{X}_{e_f} + x_{e_f} \rceil}_{=1}\Big) \\ &= \alpha\Big(1 - x_{e_f}\Big)\end{aligned}$$

Now, with the help of the previous claim, we will show that when we apply rule 1.(a) of our algorithm, we can add the selected edge e into an $\alpha = \frac{1}{2}$ of the matchings μ_i that constitute an α approximate convex decomposition I of the residual solution. Remember that the rule 1.(a) says that

$$\exists v \in F : |\mathsf{supp}(x^* \cap \delta(v))| \leq 2 \Rightarrow \exists e \in \delta(v) : x_e \geq \frac{\beta_v}{2}$$

As before, we want to calculate the fraction of the matchings μ_i that $e = \{u, v\}$ of color $c_j \in \mathcal{C}$ can be inserted preserving feasibility (i.e., μ_i is still a matching) and the above rule (that in the resulting matching μ_i after the addition of e we have that $|\mu_i \cap E_j| \leq \theta_j$). For this, we will calculate the fraction of μ_i that block the insertion of e: these are all the matchings μ_i that have edges adjacent to either u or v, and all the matchings μ_j that have θ_j edges of color c_j. In the residual solution vector x' (x without e) we have that (1) $\sum_{e' \in \delta(u)} x'_{e'} \leq 1 - x_e$, (2) the single edge e_2 adjacent to v has $x_{e_2} = \beta_v - x_e \leq \beta_v/2$ (remember that we have selected v such that the degree of v is equal to 2), and since we have an α-approximate convex decomposition of x', this means that the fraction of solutions that block the insertion of e (using also the previous claim) is at most

$$\alpha(1 - x_e) + \alpha(\beta_v - x_e) + \alpha(1 - x_e) = B$$

In clear analogy with the previous case (Lemma 3), since we want to insert e into an α fraction of the matchings in I, we want that $1 - B \geq \alpha x_e$ from which we conclude that the fraction of the matchings μ_i of I that e can be inserted is at least

$$1 - \alpha\Big(1 - x_e + \beta_v - x_e + 1 - x_e\Big) \geq \alpha x_e \Rightarrow \alpha \leq \frac{1}{1 + \beta_v - x_e + 1 - x_e}$$

and using the fact that $x_e \geq \frac{\beta_v}{2} \Rightarrow \beta_v - 2x_e \leq 0$ we conclude that $\alpha \leq \frac{1}{2}$ such that we can select $\alpha = \frac{1}{2}$ in this case as well. And so, in clear analogy with Lemma 3 we have proved the following:

Lemma 5. *If we select to insert an edge e of color c_j according to rule 1.(a) of the algorithm, then the fraction of the solutions μ_i of an α-approximate convex decomposition of the residual solution x' that e can be added is at least $\frac{1}{2}$.*

Theorem 2. *We can, in polynomial time, construct an $\frac{1}{2}$-approximate convex decomposition of x^*, resulting in a polynomial time $\frac{1}{2}$-approximation algorithm for BCM in general graphs.*

Acknowledgements. The author would like to thank Monaldo Mastrolilli for his support during the development of this work, Christos Nomikos and Panagiotis Cheilaris for discussions various issues in preliminary versions of this work.

References

1. Bampas, E., Pagourtzis, A., Potika, K.: An experimental study of maximum profit wavelength assignment in wdm rings. Networks 57(3), 285–293 (2011)
2. Berger, A., Bonifaci, V., Grandoni, F., Schäfer, G.: Budgeted matching and budgeted matroid intersection via the gasoline puzzle. In: Lodi, A., Panconesi, A., Rinaldi, G. (eds.) IPCO 2008. LNCS, vol. 5035, pp. 273–287. Springer, Heidelberg (2008)
3. Berman, P.: A $d/2$ approximation for maximum weight independent set in d-claw free graphs. In: Halldórsson, M.M. (ed.) SWAT 2000. LNCS, vol. 1851, pp. 214–219. Springer, Heidelberg (2000)
4. Caragiannis, I.: Wavelength management in wdm rings to maximize the number of connections. SIAM J. Discrete Math. 23(2), 959–978 (2009)

5. Carathéodory, C.: Über den variabilitätsbereich der fourierschen konstanten von positiven harmonischen funktionen. Rendiconti del Circolo Matematico di Palermo 32, 193–217 (1911)
6. Chan, Y.H., Lau, L.C.: On linear and semidefinite programming relaxations for hypergraph matching. Math. Program. 135(1-2), 123–148 (2012)
7. Chekuri, C., Vondrák, J., Zenklusen, R.: Multi-budgeted matchings and matroid intersection via dependent rounding. In: SODA, pp. 1080–1097 (2011)
8. Edmonds, J.: Maximum matching and a polyhedron with $0, 1$ vertices. J. of Res. the Nat. Bureau of Standards 69B, 125–130 (1965)
9. Fürer, M., Yu, H.: Approximate the k-set packing problem by local improvements. In: ISCO-3rd International Symbosium on Combinatorial Optimization, Lisboa, Portugal, March 5-7, Lisboa, Portugal, March 5-7 (2014)
10. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman (1979)
11. Grandoni, F., Ravi, R., Singh, M.: Iterative rounding for multi-objective optimization problems. In: Fiat, A., Sanders, P. (eds.) ESA 2009. LNCS, vol. 5757, pp. 95–106. Springer, Heidelberg (2009)
12. Grandoni, F., Zenklusen, R.: Approximation schemes for multi-budgeted independence systems. In: de Berg, M., Meyer, U. (eds.) ESA 2010, Part I. LNCS, vol. 6346, pp. 536–548. Springer, Heidelberg (2010)
13. Mastrolilli, M., Stamoulis, G.: Constrained matching problems in bipartite graphs. In: Mahjoub, A.R., Markakis, V., Milis, I., Paschos, V.T. (eds.) ISCO 2012. LNCS, vol. 7422, pp. 344–355. Springer, Heidelberg (2012)
14. Mastrolilli, M., Stamoulis, G.: Bi-criteria approximation algorithms for restricted matchings. Theoretical Computer Science 540-541, 115–132 (2014)
15. Monnot, J.: The labeled perfect matching in bipartite graphs. Inf. Process. Lett. 96(3), 81–88 (2005)
16. Nomikos, C., Pagourtzis, A., Zachos, S.: Minimizing request blocking in all-optical rings. In: IEEE INFOCOM (2003)
17. Nomikos, C., Pagourtzis, A., Zachos, S.: Randomized and approximation algorithms for blue-red matching. In: Kučera, L., Kučera, A. (eds.) MFCS 2007. LNCS, vol. 4708, pp. 715–725. Springer, Heidelberg (2007)
18. Papadimitriou, C.H., Yannakakis, M.: The complexity of restricted spanning tree problems. J. ACM 29(2), 285–309 (1982)
19. Parekh, O.: Iterative packing for demand and hypergraph matching. In: Günlük, O., Woeginger, G.J. (eds.) IPCO 2011. LNCS, vol. 6655, pp. 349–361. Springer, Heidelberg (2011)
20. Yuster, R.: Almost exact matchings. Algorithmica 63(1-2), 39–50 (2012)