

On the cost of essentially fair clusterings

Ioana O. Bercea¹, Martin Groß², Samir Khuller¹, Aounon Kumar¹, Clemens Rösner³, Daniel R. Schmidt⁴, and Melanie Schmidt³

¹ Department of Computer Science, University of Maryland, College Park, USA

² RWTH Aachen, Germany

³ University of Bonn, Germany

⁴ University of Cologne, Germany

Abstract. Clustering is a fundamental tool in data mining. It partitions points into groups (clusters) and may be used to make decisions for each point based on its group. However, this process may harm protected (minority) classes if the clustering algorithm does not adequately represent them in desirable clusters – especially if the data is already biased.

At NIPS 2017, Chierichetti et al. [14] proposed a model for *fair clustering* requiring the representation in each cluster to (approximately) preserve the fraction of each protected class. Restricting to two protected classes, they developed both a 4-approximation for the fair k -center problem and a $\mathcal{O}(t)$ -approximation for the fair k -median problem, where t is a parameter for the fairness model. For multiple protected classes, the best known result is a 14-approximation for fair k -center [34].

We extend and improve the known results. Firstly, we give a 6-approximation for the fair k -center problem with multiple protected classes. Secondly, we propose a relaxed fairness notion under which we can give bicriteria constant-factor approximations for all of the classical clustering objectives k -center, k -supplier, k -median, k -means and facility location. All approximations are achieved in a single framework: It takes an existing unfair solution and a fair LP solution and combines them into an essentially fair clustering with a weakly supervised rounding scheme. In this way, a fair clustering can be established belatedly, in a situation where the centers are already fixed.

1 Introduction

Suppose we are to reorganize school assignments in a big city. Given a long list of children starting school next year and a short list of all available teachers, the goal is to assign the students-to-be to (public) schools such that the maximum distance to the school is small. The school capacity is given by the number of its teachers: For each teacher, s students can be admitted.

This challenge is in fact an instance of the capacitated (metric) k -center problem. A naïve solution may, however, result in some schools having an excess of boys while others might have a surplus of girls. We would prefer an assignment where the classes are more balanced. Thus a new challenge arises: Assign the children such that the ratio is (approximately) 1:1 between boys and girls, and minimize the maximum distance under this condition.⁵ This can be modeled by

⁵ Or, incorporating the capacities, ensure that the teacher:boys:girls ratio is $1:\frac{s}{2}:\frac{s}{2}$.

the following combinatorial optimization problem: Given a point set, half of the points are red, the other half is blue. Compute a clustering where each cluster has an equal number of red and blue points, and minimize the maximum radius.

In this form, our example is a special case of the *fair k-center* problem, as proposed by Chierichetti et al. [14] in the context of maintaining fairness in *unsupervised* machine learning tasks. Their model is based on the concept of *disparate impact* [33] (and the p%-rule). The input points are assumed to have a binary sensitive attribute modeled by two colors, and discrimination based on this attribute is to be avoided. Since preserving exact balance in each cluster may be very costly or even be impossible⁶, the idea is to ensure that at least $1/t$ of the points of each cluster are of the minority color, where t is a parameter. A cluster with this property is called *fair*, and the fairness constraint can now be added to any clustering problem, giving rise to fair k -center, fair k -median, etc. Chierichetti et al. develop a 4-approximation for fair k -center and a $(t + 1 + \sqrt{3} + \epsilon)$ -approximation for fair k -median.

The fair clustering model as proposed by Chierichetti et al. can also be used to incorporate other aspects into our school assignment example: For example, we might want to mitigate effects of gentrification or segregation. For these use cases, we need multiple colors. Then, in each cluster, the ratio between the number of points with one specific color and the total number of points shall be in some given range. If the allowed range is $[0.20, 0.25]$ for red points, we require that in each cluster, at least a fifth and at most a fourth of the points are red. This models well established notions of fairness (statistical parity, group fairness), which require that each cluster exhibits the same compositional makeup as the overall data with respect to a given attribute. One downside of this notion is that a malicious user could create an illusion of fairness by including proxy points: If we wanted to create an boy-heavy school in our above example, we could still achieve the desired parity by assigning only girls that are very unlikely to attend. Thus, instead of enforcing *equal representation* in the above sense, one could also ask for *equal opportunity* as proposed by Hardt et al. [20] for the case where we take binary decisions (i.e., $k = 2$) and have access to a labeled training set. This approach, however, raises the philosophical question if this equality of opportunity is a sufficient condition for the absence of discrimination. Rather than delving into this complex and much debated issue in this algorithmic paper, we refer to the excellent surveys by Romei and Ruggieri [33] and Žliobaitė et al. [36] that systematically discuss different forms of discrimination and how they can be detected. We assume that it is the intent of the user to achieve a truly fair solution.

Finding fair clusterings turns out to be an interesting challenge from the point of view of combinatorial optimization. As other clustering problems with side constraints, it loses the property that points can be assigned locally. But while many other constraint problems at least allow polynomial algorithms that assign

⁶ Imagine a point set with 49 red and 51 blue points: This can not at all be divided into true subsets with exact the same ratio.

points to given centers optimally, we show that even this restricted problem is NP-hard in the case of fair k -center.

Chierichetti et al. tackle fair clustering problems by a two-step procedure: First, they compute a micro clustering into so-called *fairlets*, which are groups of points that are fair and can not be split further into true subsets that are also fair. Secondly, representative points of the fairlets are clustered by an approximation algorithm for the unconstrained problem. Consider the special case of a point set with 1:1 ratio of red and blue points. Then a fairlet is a pair of one red and one blue point, and a good micro clustering can be found by computing a suitable bipartite matching between the two color classes.

The problem of computing good fairlets gets increasingly difficult when considering more general variants of the problem. For multiple colors and the special case of exact ratio preservation (i.e., for all colors, the allowed range for its ratio is one specific number), the fairlet computation problem can be reduced to a capacitated clustering problem. This is used in [34] to obtain a 14/15-approximation for fair k -center/ k -supplier with multiple colors and exact ratio preservation.

In the full version, we give an extensive overview of the existing results and further the fairlet approach in order to explore its applicability for different variants of fair clustering. Two major issues arise: Firstly, capacitated clustering is not solved for all clustering objectives; indeed, finding a constant-factor approximation for k -median is a long-standing open problem. Secondly, (even for k -center) it is unclear how fairlets even look like when we have multiple colors and want to allow ranges for the ratios. In this situation, subsets of very different size and composition may satisfy the desired ratio.

The main contribution of this paper is a very different approach. We start with a solution to the unconstrained problem. Since this can be any solution, we say our algorithm is a *black-box* approximation. Based on the given solution, we derive a fair clustering solution with the same centers. That is achieved by a technique that we call *weakly supervised LP rounding*: We solve an LP for the fair clustering problem and then combine it with the integral unfair solution by careful rounding. We use this method to prove the following statements.

Theorem 1. *There exists a 5/7-approximation for the fair k -center/ k -supplier problem with exact preservation of ratios.*

Theorem 2. *Given any set of centers S , there exists an assignment ϕ' : which is essentially fair, and which incurs a cost that is linear in the cost of S for the unconstrained problem and the cost of an optimal fair clustering of P , for all objectives k -center, k -supplier, k -median, k -means and facility location.*

Corollary 1. *There exists a 3/5/3.488/4.675/62.856-approximation for the fair k -center/ k -supplier/facility location/ k -median/ k -means problem.*

Here, *essentially fair* refers to our notion of bicriteria approximation: A cluster is *essentially fair* if by fractionally changing its assigned points by a total amount of *at most* 1, we can get all ratios into their allowed ranges. So this is a small additive fairness violation.

Our results have two advantages. Firstly, we get results for a wide range of clustering problems, and these results improve previous results. For example, we get a 5-approximation for the fair k -center problem with exact ratio preservation, where the best known guarantee was 14. All our bicriteria results work for multiple colors and approximate ratio preservation, a case for which no previous algorithm was known. As for the quality of the guarantees, compare the 4.675-approximation for essentially fair k -median clusterings with the best previously known $\Theta(t)$ -approximation, which is only applicable to the case of two colors. Notice that a similar result can *not* be achieved by using bicriteria approximation algorithms for capacitated clustering. The reduction from capacitated clustering only works when the capacities are not violated.

Secondly, the black-box approach has the advantage that fairness can be established belatedly, in a situation where the centers are already given. [17,37]. Consider our school example and notice that the location of the schools cannot be chosen. Our result says that if we are alright with essentially fair clusterings, we get a clustering which is not much more expensive than a fair clustering where the centers were chosen with the fairness constraint at hand.

Additional related work The unconstrained k -center/ k -supplier problem can be $2/3$ -approximated [18,21], and this is tight [22]. Facility location can be 1.488-approximated [29], and the best lower bound is 1.463 [19]. For k -median, the best upper bound is 2.765 [32,11], while the best hardness result lies below two [23]. A recent breakthrough gives a 6.357-approximation for k -means [4], but the newest hardness result is marginally above 1 [8,26].

The k -center problem allows for constant-factor approximations for many useful constraints such as capacity constraints [10,15,24], lower bounds on the size of each cluster [3,6] or allowing for outliers [12,16]. This is also true for facility location and capacities [2,7,9], uniform lower bounds [5,35], and outliers [12]. Much less is known for k -median and k -means. True constant-factor approximations so far exist only for the outlier constraint [13,25]. A major problem for obtaining constant factor approximations is that the natural LP has an unbounded integrality gap, which is also true for the LP with fairness constraints. Bicriteria approximations are known that either violate the capacity constraints [28,30,31] or the cardinality constraint [1].

A clustering problem where the points have a color was considered by Li, Yi and Zhang [27]. They provided a 2-approximation for a constraint called *diversity*, which allows at most one point per color in each cluster. Fair clustering as considered in this paper is studied in [14] and [34].

Preliminaries

Points and locations. We are given a set of n points P and a set of potential locations L . We allow L to be infinite (when $L = \mathbb{R}^d$). The task is to open a subset $S \subseteq L$ of the locations and to assign each point in P to an open location via a mapping $\phi : P \rightarrow S$. We refer to the set of all points assigned to a location

$c_i \in S$ by $P(i) := \phi^{-1}(i)$. The assignment incurs a cost governed by a semi-metric $d : (P \cup L) \times (P \cup L) \rightarrow \mathbb{R}_{\geq 0}$ that fulfills a β -relaxed triangle inequality

$$d(x, z) \leq \beta(d(x, y) + d(y, z)) \quad \text{for all } x, y, z \in P \cup L \quad (1)$$

for some $\beta \geq 1$. Additionally, we may have opening costs $f_i \geq 0$ for every potential location $c_i \in L$ or a maximum number of centers $k \in \mathbb{N}$.

Colors and fairness. We are also given a set of *colors* $Col := \{col_1, \dots, col_g\}$, and a coloring $col : P \rightarrow Col$ that assigns a color to each point $j \in P$. For any set of points $P' \subseteq P$ and any color $col_h \in Col$ we define $col_h(P') = \{j \in P' \mid col(j) = col_h\}$ to be the set of points colored with col_h in P' . We call $r_h(P') := \frac{|col_h(P')|}{|P'|}$ the *ratio* of col_h in P' . If an implicit assignment ϕ is clear from the context, we write $col_h(i)$ to denote the set of all points of a color $col_h \in Col$ assigned to an $i \in S$, i.e., $col_h(i) = col_h(P(i))$.

In our fair clustering problems, we want to preserve the ratios of colors found in P in our clusters. We distinguish two cases: *exact* preservation of ratios, and *relaxed* preservation of ratios. For the exact preservation of ratios, we ask that all the points assigned to a location $i \in S$ have exactly the same ratios as P , i.e.,

$$r_h(P(i)) = r_h(P) \quad \text{for all } col_h \in Col \text{ and } i \in S. \quad (2)$$

For the relaxed preservation of ratios, we have upper and lower bounds on the ratio of a color in each cluster. More precisely, we are given $p_1^h, q_1^h, p_2^h, p_2^h \in \mathbb{N}$ for every $c_h \in Col$, and the ratios have to fulfill

$$\ell_h := \frac{p_1^h}{q_1^h} \leq r_h(P(i)) \leq \frac{p_2^h}{q_2^h} =: u_h \quad \text{for all } col_h \in Col \text{ and } i \in S. \quad (3)$$

The exact case is a special case of the relaxed case where we set $\ell_h = u_h = r_h(P)$ for every color $col_h \in Col$.

A set of points $P' \subseteq P$ is *exactly fair* if P' has the same ratio for every color as P , i.e., for each $col_h \in Col$ we have $r_h(P') = r_h(P)$. We say that P' is ℓ, u -*fair* or just *fair* for some $\ell = (\ell_1 = p_1^1/q_1^1, \dots, \ell_g = p_1^g/q_1^g)$ and $u = (u_1 = p_2^1/q_2^1, \dots, u_g = p_2^g/q_2^g)$ if we have $r_h(P') \in [\ell_h, u_h]$ for every color $col_h \in Col$. *Essentially fair* clusterings are defined below (see Definition 1).

Objectives. We consider fair versions of several classical clustering problems. An instance is given by $I := (P, L, col, d, f, k, \ell, u)$, and our goal is to choose a solution (S, ϕ) according to one of the following objectives.

- **k -center** and **k -supplier**: minimize the maximum distance between a point and its assigned location: $\min \max_{j \in P} d(j, \phi(j))$. In these problems, we have $f \equiv 0$ and d is a metric. Furthermore, in k -center, $L = P$, whereas in k -supplier, $L \neq P$ is some finite set.
- **k -median**: minimize $\sum_{j \in P} d(j, \phi(j))$, d is a metric, $f \equiv 0$ and $L \subseteq P$.
- **k -means**: minimize $\sum_{j \in P} d(j, \phi(j))$, where $P \subseteq \mathbb{R}^m$ for some $m \in \mathbb{N}$, $L = \mathbb{R}^m$ and $d(x, y) = \|y - x\|^2$ is a semi-metric for $\beta = 2$ and $f \equiv 0$.
- **facility location**: minimize $\sum_{j \in P} d(j, \phi(j)) + \sum_{i \in S} f_i$, where $k = n$, d is a metric and L is a finite set.

The fair assignment problem. For all the objectives above, we call the subproblem of computing a cost-minimal fair assignment of points to given centers the *fair assignment problem*. We show the following theorem in the full version.

Theorem 3. *Finding an α -approximation for the fair assignment problem for k -center for $\alpha < 3$ is NP-hard.*

(I)LP formulations for fair clustering problems

Let $I = (P, L, col, d, f, k, \ell, u)$ be a problem instance for a fair clustering problem. We introduce a binary variable $y_i \in \{0, 1\}$ for all $i \in L$ that decides if i is opened, i.e. $y_i = 1 \Leftrightarrow i \in S$. Similarly, we introduce binary variables $x_{ij} \in \{0, 1\}$ for all $i \in L, j \in P$ with $x_{ij} = 1$ if j is assigned to i , i.e. $\phi(j) = i$. All ILP formulations have the inequalities (4) $\sum_{i \in L} x_{ij} = 1 \forall j \in P$ saying that every point j is assigned to a center, the inequalities (5) $x_{ij} \leq y_i \forall i \in L, j \in P$ ensuring that if we assign j to i , then i must be open, and the integrality constraints (6) $y_i, x_{ij} \in \{0, 1\} \forall i \in L, j \in P$. We may restrict the number of open centers to k with (7) $\sum_{i \in L} y_i \leq k$. For k -center and k -supplier, the objective is commonly encoded in the constraints of the problem, and the (I)LP has no objective function. The idea is to guess the optimum value τ . Since there is only a polynomial number of choices for τ , this is easily done. Given τ , we construct a *threshold graph* $G_\tau = (P \cup L, E_\tau)$ on the points and locations, where a connection between $i \in L$ and $j \in P$ is added iff i and j are close, i.e., $\{i, j\} \in E_\tau \Leftrightarrow d(i, j) \leq \tau$. Then, we ensure that points are not assigned to centers outside their range:

$$x_{ij} = 0 \quad \text{for all } i \in L, j \in P, \{i, j\} \notin E_\tau \quad (8)$$

For the remaining clustering problems, we pick the adequate objective function from the following three (let $d_{ij} := d(i, j)$):

$$\min \sum_{i \in L, j \in P} x_{ij} d_{ij} \quad (9) \quad \min \sum_{i \in L, j \in P} x_{ij} d_{ij}^2 \quad (10) \quad \min \sum_{i \in L, j \in P} x_{ij} d_{ij} + \sum_{i \in L} y_i f_i \quad (11)$$

We now have all necessary constraints and objectives. For k -center and k -supplier, we use inequalities (4)-(8), no objective, and define the optimum to be the smallest τ for which the ILP has a solution. We get k -median and k -means by combining inequalities (4)-(7) with (9) and (10), respectively, and we get facility location by combining (4)-(6) with the objective (11). LP relaxations arise from all ILP formulations by replacing (6) by $y_i, x_{ij} \in [0, 1]$ for all $i \in L, j \in P$. To create the fair variants of the ILP formulations, we add fairness constraints modeling the upper and lower bound on the balances.

$$\ell_h \sum_{j \in P} x_{ij} \leq \sum_{col(p_j)=col_h} x_{ij} \leq u_h \sum_{j \in P} x_{ij} \quad \text{for all } i \in L, h \in Col \quad (12)$$

Although very similar to the canonical clustering LPs, the resulting LPs become much harder to round even for k -center with two colors. We show the following in the full version.

Lemma 1. *There is a choice of non-trivial fairness intervals such that the integrality gap of the LP-relaxation of the canonical fair clustering ILP is $\Omega(n)$ for the fair k -center/ k -supplier/facility location problem. The integrality gap is $\Omega(n^2)$ for the fair k -means problem.*

Essential fairness. For a point set P' , $\text{mass}_h(P') = |\text{col}_h(P')|$ is the mass of color col_h in P' . For a possibly fractional LP solution (x, y) , we extend this notion to $\text{mass}_h(x, i) := \sum_{j \in \text{col}_h(P)} x_{ij}$. We denote the total mass assigned to i in (x, y) by $\text{mass}(x, i) = \sum_{j \in P} x_{ij}$.

With this notation, we can now formalize our notion of *essential fairness*.

Definition 1 (Essential fairness). *Let I be an instance of a fair clustering problem, let (x, y) be a fractional solution for I and $(x', y') := \vartheta(x, y)$ be an integral solution to I . We say that ϑ is essentially fair if $\forall i \in L$ with $y_i > 0$:*

$$\lfloor \text{mass}_h(x, i) \rfloor \leq \text{mass}_h(x', i) \leq \lceil \text{mass}_h(x, i) \rceil \quad \forall \text{col}_h \in \text{Col} \quad (13)$$

$$\text{and } \lfloor \text{mass}(x, i) \rfloor \leq \text{mass}(x', i) \leq \lceil \text{mass}(x, i) \rceil. \quad (14)$$

2 Essential fair clusterings via black-box approximation

For essentially fair clustering, we give a powerful framework that employs approximation algorithms for (unfair) clustering problems as a black-box and transforms their output into an essentially fair solution. In this framework, we start by computing an approximate solution for the standard variant of the clustering problem at hand. Next, we solve the LP for the fair variant of the clustering problem. Now we have an integral unfair solution, and a fractional fair solution. Our final and most important step is to combine these two solutions into an integral and essentially fair solution. It consists of two conceptual sub-steps: Firstly, we show that it is possible to find a fractional fair assignment to the centers of the integral solution that is sufficiently cheap. Secondly, we round the assignment. This last sub-step introduces the potential fairness violation of one point per color per cluster.

We show that this approach yields constant-factor approximations with fairness violation for all mentioned clustering objectives. The description will be neutral whenever the objective does not matter. Thus, descriptions like *the LP* mean the appropriate LP for the desired clustering problem. When the problem gets relevant, we will specifically discuss the distinctions. Notice that for all clustering problems defined in Section 1, P and L are finite except for k -means. However, for the k -means problem, we can assume that $L = P$ if we accept an additional factor of 2 in the approximation guarantee. Thus, we assume in the following that L and P are finite sets. Indeed, we even assume at least $L \subseteq P$ for all problems except k -supplier and facility location.

2.1 Step 1: Obtaining a fair solution with integral y

In the first step, we assume that we are given two solutions. Let (x^{LP}, y^{LP}) be an optimal solution to the LP. This solution has the property that the assignments

to all centers are fair, however, the centers may be fractionally open and the points may be fractionally assigned to several centers. Let c^{LP} be the objective value of this solution. For k -supplier and k -center, it is the smallest τ for which the LP is feasible, for the other objectives, it is the value of the LP. We denote the cost of the best *integral* solution to the LP by c^* . We know that $c^{LP} \leq c^*$.

Let (\bar{x}, \bar{y}) be any integral solution to the LP that may violate fairness, i.e., inequality (12), and let \bar{c} be the objective value of this solution. We think of (\bar{x}, \bar{y}) as being a solution of an α -approximation algorithm for the standard (unfair) clustering problem for some constant α . Since the unconstrained version can only have a lower optimum cost, we then have $\bar{c} \leq \alpha \cdot c^*$.

Our goal is now to combine (x^{LP}, y^{LP}) and (\bar{x}, \bar{y}) into a third solution, (\hat{x}, \hat{y}) , such that the cost of (\hat{x}, \hat{y}) is bounded by $O(c^{LP} + \bar{c}) \subseteq O(c^*)$. Furthermore, the entries of \hat{y} shall be integral. The entries of \hat{x} may still be fractional after step 1.

Let S be the set of centers that are open in (\bar{x}, \bar{y}) . For all $j \in P$, we use $\bar{\phi}(j)$ to denote the center in S closest to j , i.e., $\bar{\phi}(j) = \arg \min_{i \in S} d(j, i)$ (ties broken arbitrarily). Notice that the objective value of using S with assignment $\bar{\phi}$ for all points in P is at most \bar{c} , since assigning to the closest center is always optimal for the standard clustering problems without fairness constraint.

Depending on the objective, L is a subset of P or not, i.e., $\bar{\phi}$ is not necessarily defined for all locations in L . We then extend $\bar{\phi}$ in the following way. Let $i \in L \setminus P$ be any center, and let j^* be the closest point to it in P . Then we set $\bar{\phi}(i) := \bar{\phi}(j^*)$, i.e., i is assigned to the center in S which is closest to the point in P which is closest to i . Finally, let $\bar{C}(i) = \bar{\phi}^{-1}(i)$ be the set of all points and centers assigned to i by $\bar{\phi}$. We show the following lemma (proof is in the appendix).

Lemma 2. *Let (x^{LP}, y^{LP}) and (\bar{x}, \bar{y}) be two solutions to the LP, where (\bar{x}, \bar{y}) may violate inequality (12), but is integral. Then the solution defined by $\hat{y} := \bar{y}$ and*

$$\hat{x}_{ij} := \sum_{i' \in \bar{C}(i)} x_{i'j}^{LP} \quad \text{for all } i \in S, j \in P, \quad \hat{x}_{ij} := 0 \quad \text{for all } i \notin S, j \in P.$$

satisfies inequality (12), \hat{y} is integral, and the cost \hat{c} of (\hat{x}, \hat{y}) is bounded by $c^{LP} + \bar{c}$ for k -center, by $2 \cdot c^{LP} + \bar{c}$ for k -supplier, k -median, and facility location, and by $12 \cdot c^{LP} + 8 \cdot \bar{c}$ for k -means.

2.2 Step 2: Rounding the x -variables

For rounding the x -variables, we need to distinguish between two cases of objectives. Let $j \in P$ be a point that is fractionally assigned to some centers $L_j \subseteq L$.

First, we have objectives where we can shift mass from an assignment of j to $i' \in L_j$ to an assignment of j to $i'' \in L_j$ without modifying the objective. We say that such objectives are *reassignable* (in the sense that we can reassign j to centers in L_j without changing the cost). k -center and k -supplier have this property.

Second, we have objectives where the assignment cost is separable, i.e., where the distances influence the cost via a term of the form $\sum_{i \in L, j \in P} c_{ij} \cdot x_{ij}$ for some

$c_{ij} \in \mathbb{R}_{\geq 0}$. We call such objectives *separable*. Facility location, k -median and k -means fall into the this category.

Lemma 3. *Let (x, y) be an α -approximative fractional solution for a fair clustering problem with the property that all $y_i, i \in L$ are integral. Then we can obtain an α -approximative integral solution (x', y') with an additive fairness violation of at most one in time $O(\text{poly}(|S| + |P|))$, with $S := \{i \in L \mid y_i \geq 1\}$ being the set of locations that are opened in (x, y) .*

Proof. We create our rounded α -approximate integral solution (x', y') by min-cost flow computations. We begin by constructing a min-cost flow instance which depends on our starting solution (x, y) as well as on the objective of the problem we are studying.

We define a min-cost flow instance $(G = (V, A), c, b)$ (also see Figure 1) with unit capacities and costs c on the edges as well as balances b on the nodes. We begin by defining a graph $G^h = (V^h, A^h)$ for every color $h \in \text{Col}$ with

$$\begin{aligned} V^h &:= V_S^h \cup V_P^h, & V_S^h &:= \{v_i^h \mid i \in S\}, & V_P^h &:= \{v_j^h \mid j \in \text{col}_h(P)\}, \\ A^h &:= \{(v_j^h, v_i^h) \mid i \in S, j \in \text{col}_h(P) : x_{ij} > 0\}, \end{aligned}$$

as well as costs c^h by $c_a^h := c_{ij}$ for $a = (v_j^h, v_i^h) \in A^h, i \in S, j \in \text{col}_h(P)$ and balances b^h by $b_v^h := 1$ if $v \in V_P^h$ and $b_v^h := -\lfloor \text{mass}_h(x, i) \rfloor$ if $v = v_i^h \in V_S^h$. We use the graphs G_h to define $G = (V, A)$ by

$$\begin{aligned} V &:= \{t\} \cup V_S \cup \bigcup_{h \in \text{Col}} V^h, & V_S &:= \{v_i \mid i \in S\} \\ A &:= \bigcup_{h \in \text{Col}} A^h \cup \{(v_i^h, v_i) \mid i \in S, h \in \text{Col} : \text{mass}_h(x, i) - \lfloor \text{mass}_h(x, i) \rfloor > 0\} \\ &\quad \cup \{(v_i, t) \mid i \in S : \text{mass}(x, i) - \lfloor \text{mass}(x, i) \rfloor > 0\}, \end{aligned}$$

together with costs c of $c_a := c_a^h$ for $a \in A^h$ and 0 otherwise, and balances b of $b_v := b_v^h$ if $v \in V^h$ for some $h \in \text{Col}$, $b_v := -B_i$ if $v = v_i \in V_S$ and $b_t := -B$ with $B_i = \lfloor \text{mass}(x, i) \rfloor - \sum_{h \in \text{Col}} \lfloor \text{mass}_h(x, i) \rfloor$ and $B := |P| - \sum_{i \in S} \lfloor \text{mass}(x, i) \rfloor$.

Separable objectives – k -median and k -means. We observe that:

1. B and B_i are integers for all $i \in S$, and so are all capacities, costs and balances. Consequently, there are integral optimal solutions for the min-cost flow instance (G, c, b) ,
2. (x, y) induces a feasible solution for (G, c, b) , by defining a flow x in G as follows:

$$x_a := \begin{cases} x_{ij} & \text{if } a = (v_j^h, v_i^h) \in A^h, j \in P, i \in S, \\ \text{mass}_h(x, i) - \lfloor \text{mass}_h(x, i) \rfloor & \text{if } a = (v_i^h, v_i) \in A, h \in \text{Col}, i \in S, \\ \text{mass}(x, i) - \lfloor \text{mass}(x, i) \rfloor & \text{if } a = (v_i, t) \in A, i \in S. \end{cases}$$

Since (x, y) is a fractional solution, x satisfies capacity and non-negativity constraints because $x_{ij} \in [0, 1]$ for all $i \in L, j \in P$ and $\text{mass}_h(x, i) - \lfloor \text{mass}_h(x, i) \rfloor, \text{mass}(x, i) - \lfloor \text{mass}(x, i) \rfloor \in [0, 1]$ for all $i \in S$ and $\text{col}_h \in \text{Col}$ as well. We have flow conservation since the fractional solution needs to assign all points, and the flow of the edges (v_i^h, v_i) and (v_i, t) as well as the demand of v_i and t are chosen in such a way that we have flow conservation for all the other nodes as well.

3. Integral solutions x to the min-cost flow instance (G, c, b) induce an integral solution (\bar{x}, y) to the original clustering problem by setting $\bar{x}_{ij} := x_a$ for $a = (v_j^h, v_i^h) \in A^h$ if $j \in \text{col}_h(P), i \in S$. Since the flow x is integral, this gives us an integral assignment of all points to centers which have been opened, since y was already integral before this step.

This incurs the additive fairness violation of at most one, since every $i \in S$ is guaranteed by our balances to have at least $\lfloor \text{mass}_h(x, i) \rfloor$ points of color $h \in \text{Col}$ and at least $\lfloor \text{mass}(x, i) \rfloor$ points in total assigned to it. Since there is at most one outgoing arc of unit capacity (v_i^h, v_i) and (v_i, t) for an $i \in S$ if $\text{mass}_h(x, i) - \lfloor \text{mass}_h(x, i) \rfloor > 0$, we have at most $\lceil \text{mass}_h(x, i) \rceil$ points of color col_h and $\lceil \text{mass}(x, i) \rceil$ total points assigned to i .

Together, this yields that computing a min-cost flow \hat{x} for (G, c, b) followed by applying the third observation to \hat{x} yields a solution (\hat{x}, y) to the clustering with an additive fairness violation of at most one.

Since (x, y) was inducing the fractional solution x with $\text{cost}(x) = \text{cost}(x, y)$ to the min-cost flow instances, and $\text{cost}(x) \geq \text{cost}(\hat{x})$ by construction we have $\text{cost}(\hat{x}, y) \leq \text{cost}(x, y)$.

Reassignable objectives – k-center and k-supplier. In the case of reassignable objectives, we do not have to care about costs, as long as the reassignments happen to centers in L_j for all points $j \in P$. We essentially use the same strategy as before, but instead of a min cost flow problem we solve the transshipment problem $(G = (V, A), b)$ with unit capacities on the edges and balances b on the nodes. Notice that the three observations from the previous case apply here as well, and reassignability guarantees that the cost does not increase.

Lemmas 2 and 3 then lead directly to Theorem 2, or, in more detail, to:

Theorem 4. *Black-box approximation for fair clustering gives essentially fair solutions with a cost of $c^{LP} + \bar{c}$ for k-center, $2c^{LP} + \bar{c}$ for k-supplier, k-median and facility location, and $12c^{LP} + 8\bar{c}$ for k-means where c^{LP} is the cost of an optimal solution to the fair LP relaxation and \bar{c} is the cost of the given solution.*

We know that c^{LP} is not more expensive than an optimal solution to the fair clustering problem. If we use an α -approximation to obtain the unfair clustering solution, we have that \bar{c} is at most α times the cost of an optimal solution to the fair clustering problem. Currently, the best know approximation factors are 2 for k-center [18,21], 3 for k-supplier [21], 1.488 for facility location [29], 2.675 for k-median [11,32] and 6.357 for k-means [4], which yields Corollary 1.

References

1. Aardal, K., van den Berg, P.L., Gijswijt, D., Li, S.: Approximation algorithms for hard capacitated k -facility location problems. *European Journal of Operational Research* **242**, 358–368 (2015). <https://doi.org/10.1016/j.ejor.2014.10.011>
2. Aggarwal, A., Louis, A., Bansal, M., Garg, N., Gupta, N., Gupta, S., Jain, S.: A 3-approximation algorithm for the facility location problem with uniform capacities. *Mathematical Programming* **141**, 527–547 (2013). <https://doi.org/10.1007/s10107-012-0565-4>
3. Aggarwal, G., Panigrahy, R., Feder, T., Thomas, D., Kenthapadi, K., Khuller, S., Zhu, A.: Achieving anonymity via clustering. *ACM Transactions on Algorithms* **6**, 49:1–49:19 (2010). <https://doi.org/10.1145/1798596.1798602>
4. Ahmadian, S., Norouzi-Fard, A., Svensson, O., Ward, J.: Better guarantees for k -means and Euclidean k -median by primal-dual algorithms. In: Umans, C. (ed.) *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science (FOCS 2017)*. pp. 61–72. IEEE Computer Society (2017). <https://doi.org/10.1109/FOCS.2017.15>
5. Ahmadian, S., Swamy, C.: Improved approximation guarantees for lower-bounded facility location. In: Erlebach, T., Persiano, G. (eds.) *10th International Workshop on Approximation and Online Algorithms (WAOA 2012)*. *Lecture Notes in Computer Science (LNCS)*, vol. 7846, pp. 257–271. Springer Berlin Heidelberg (2012). https://doi.org/10.1007/978-3-642-38016-7_21
6. Ahmadian, S., Swamy, C.: Approximation algorithms for clustering problems with lower bounds and outliers. In: Chatzigiannakis, I., Mitzenmacher, M., Rabani, Y., Sangiorgi, D. (eds.) *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*. *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 55, pp. 69:1–69:15. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2016). <https://doi.org/10.4230/LIPIcs.ICALP.2016.69>
7. An, H., Singh, M., Svensson, O.: LP-based algorithms for capacitated facility location. *SIAM Journal on Computing* **46**, 272–306 (2017). <https://doi.org/10.1137/151002320>
8. Awasthi, P., Charikar, M., Krishnaswamy, R., Sinop, A.K.: The hardness of approximation of euclidean k -means. In: Arge, L., Pach, J. (eds.) *31st International Symposium on Computational Geometry (SoCG 2015)*. *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 34, pp. 754–767. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2015). <https://doi.org/10.4230/LIPIcs.SOCG.2015.754>
9. Bansal, M., Garg, N., Gupta, N.: A 5-approximation for capacitated facility location. In: Epstein, L., Ferragina, P. (eds.) *Algorithms – ESA 2012*. *Lecture Notes in Computer Science (LNCS)*, vol. 7501, pp. 133–144. Springer Berlin Heidelberg (2012). https://doi.org/10.1007/978-3-642-33090-2_13
10. Bar-Ilan, J., Kortsarz, G., Peleg, D.: How to allocate network centers. *Journal of Algorithms* **15**, 385–415 (1993). <https://doi.org/10.1006/jagm.1993.1047>
11. Byrka, J., Pensyl, T., Rybicki, B., Srinivasan, A., Trinh, K.: An improved approximation for k -median and positive correlation in budgeted optimization. *ACM Transactions on Algorithms* **13**, 23:1–23:31 (2017). <https://doi.org/10.1145/2981561>
12. Charikar, M., Khuller, S., Mount, D.M., Narasimhan, G.: Algorithms for facility location problems with outliers. In: Kosaraju, S.R. (ed.) *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms (SODA 2001)*. pp. 642–651 (2001), <http://dl.acm.org/citation.cfm?id=365411.365555>

13. Chen, K.: A constant factor approximation algorithm for k -median clustering with outliers. In: Teng, S.H. (ed.) Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008). pp. 826–835. SIAM (2008), <http://dl.acm.org/citation.cfm?id=1347082.1347173>
14. Chierichetti, F., Kumar, R., Lattanzi, S., Vassilvitskii, S.: Fair clustering through fairlets. In: Guyon, I., von Luxburg, U., Bengio, S., Wallach, H.M., Fergus, R., Vishwanathan, S.V.N., Garnett, R. (eds.) Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017 (NIPS 2017). pp. 5036–5044 (2017), <http://papers.nips.cc/paper/7088-fair-clustering-through-fairlets>
15. Cygan, M., Hajiaghayi, M., Khuller, S.: LP rounding for k -centers with non-uniform hard capacities. In: Guruswami, V. (ed.) 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2012). pp. 273–282. IEEE Computer Society (2012). <https://doi.org/10.1109/FOCS.2012.63>
16. Cygan, M., Kociumaka, T.: Constant factor approximation for capacitated k -center with outliers. In: Mayr, E.W., Portier, N. (eds.) 31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014). pp. 251–262. Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2014). <https://doi.org/10.4230/LIPIcs.STACS.2014.251>
17. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.S.: Fairness through awareness. In: Goldwasser, S. (ed.) Innovations in Theoretical Computer Science 2012 (ITCS 2012). pp. 214–226. ACM (2012). <https://doi.org/10.1145/2090236.2090255>
18. Gonzalez, T.F.: Clustering to minimize the maximum intercluster distance. Theoretical Computer Science **38**, 293–306 (1985). [https://doi.org/10.1016/0304-3975\(85\)90224-5](https://doi.org/10.1016/0304-3975(85)90224-5)
19. Guha, S., Khuller, S.: Greedy strikes back: Improved facility location algorithms. Journal of Algorithms **31**, 228–248 (1999). <https://doi.org/10.1006/jagm.1998.0993>
20. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. In: Advances in Neural Information Processing Systems 29 (NIPS 2016). pp. 3315–3323 (2016)
21. Hochbaum, D.S., Shmoys, D.B.: A unified approach to approximation algorithms for bottleneck problems. Journal of the ACM **33**(3), 533–550 (1986). <https://doi.org/10.1145/5925.5933>
22. Hsu, W., Nemhauser, G.L.: Easy and hard bottleneck location problems. Discrete Applied Mathematics **1**, 209–215 (1979). [https://doi.org/10.1016/0166-218X\(79\)90044-1](https://doi.org/10.1016/0166-218X(79)90044-1)
23. Jain, K., Mahdian, M., Saberi, A.: A new greedy approach for facility location problems. In: Reif, J.H. (ed.) Proceedings on 34th Annual ACM Symposium on Theory of Computing (STOC 2002). pp. 731–740. ACM (2002). <https://doi.org/10.1145/509907.510012>
24. Khuller, S., Sussmann, Y.J.: The capacitated k -center problem. SIAM Journal on Discrete Mathematics **13**, 403–418 (2000). <https://doi.org/10.1137/S0895480197329776>
25. Krishnaswamy, R., Li, S., Sandeep, S.: Constant approximation for k -median and k -means with outliers via iterative rounding. In: Diakonikolas, I., Kempe, D., Henzinger, M. (eds.) Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2018). pp. 646–659. ACM (2018). <https://doi.org/10.1145/3188745.3188882>

26. Lee, E., Schmidt, M., Wright, J.: Improved and simplified inapproximability for k -means. *Information Processing Letters* **120**, 40–43 (2017). <https://doi.org/10.1016/j.ipl.2016.11.009>
27. Li, J., Yi, K., Zhang, Q.: Clustering with diversity. In: Abramsky, S., Gavaille, C., Kirchner, C., Meyer auf der Heide, F., Spirakis, P.G. (eds.) *Automata, Languages and Programming (ICALP 2010)*, Lecture Notes in Computer Science (LNCS), vol. 6198, pp. 188–200. Springer Berlin Heidelberg (2010). https://doi.org/10.1007/978-3-642-14165-2_17
28. Li, S.: An improved approximation algorithm for the hard uniform capacitated k -median problem. In: Jansen, K., Rolim, J.D.P., Devanur, N.R., Moore, C. (eds.) *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 28, pp. 325–338. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2014). <https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2014.325>
29. Li, S.: A 1.488 approximation algorithm for the uncapacitated facility location problem. *Information and Computation* **222**, 45–58 (2013). <https://doi.org/10.1016/j.ic.2012.01.007>
30. Li, S.: Approximating capacitated k -median with $(1 + \epsilon)k$ open facilities. In: Krauthgamer, R. (ed.) *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016)*. pp. 786–796. SIAM (2016). <https://doi.org/10.1137/1.9781611974331.ch56>
31. Li, S.: On uniform capacitated k -median beyond the natural LP relaxation. *ACM Transactions on Algorithms* **13**, 22:1–22:18 (2017). <https://doi.org/10.1145/2983633>
32. Li, S., Svensson, O.: Approximating k -median via pseudo-approximation. *SIAM Journal on Computing* **45**, 530–547 (2016). <https://doi.org/10.1137/130938645>
33. Romei, A., Ruggieri, S.: A multidisciplinary survey on discrimination analysis. *The Knowledge Engineering Review* **29**, 582–638 (2014). <https://doi.org/10.1017/S0269888913000039>
34. Rösner, C., Schmidt, M.: Privacy Preserving Clustering with Constraints. In: Chatzigiannakis, I., Kaklamanis, C., Marx, D., Sannella, D. (eds.) *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Leibniz International Proceedings in Informatics (LIPIcs), vol. 107, pp. 96:1–96:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2018). <https://doi.org/10.4230/LIPIcs.ICALP.2018.96>
35. Svitkina, Z.: Lower-bounded facility location. *ACM Transaction on Algorithms* **6**, 69:1–69:16 (2010). <https://doi.org/10.1145/1824777.1824789>
36. Žliobaitė, I., Kamiran, F., Calders, T.: Handling Conditional Discrimination. In: *Proceedings of the 2011 IEEE 11th International Conference on Data Mining*. pp. 992–1001. ICDM '11, IEEE Computer Society (2011). <https://doi.org/10.1109/ICDM.2011.72>
37. Zemel, R., Wu, Y., Swersky, K., Pitassi, T., Dwork, C.: Learning fair representations. In: Dasgupta, S., McAllester, D. (eds.) *Proceedings of the 30th International Conference on Machine Learning (ICML 2013)*. Proceedings of Machine Learning Research, vol. 28, pp. 325–333. PMLR (2013), <http://proceedings.mlr.press/v28/zemel13.html>

Appendix

Proof (Proof of Lemma 2). Recall that for k -center and k -supplier, speaking of the cost of an LP solution is a bit sloppy; we mean that (\hat{x}, \hat{y}) is a feasible solution in the LP with threshold \hat{c} .

The definition of (\hat{x}, \hat{y}) means the following. For every (fractional) assignment from a point j to a center i' , we look at the cluster with center $i = \bar{\phi}(i')$ to which i' is assigned to by $\bar{\phi}$. We then shift this assignment to i . So from the perspective of i , we collect all fractional assignments to centers in $\bar{C}(i)$ and consolidate them at i . Notice that the (fractional) number of points assigned to i after this process may be less than one since (\bar{x}, \bar{y}) may include centers that are very close together.

Since that \hat{y} is simply \bar{y} it is integral as well and has the same number of centers, thus \hat{y} also satisfies (7) if the problem uses it. Next, we observe that (\hat{x}, \hat{y}) satisfies fairness, i.e., respects (12). This is true because (x^{LP}, y^{LP}) satisfies them, and because we move *all* assignment from a center i' to the same center $\bar{\phi}(i')$. This shifting operation preserves the fairness. Inequality (5) is true because we only move assignments to centers that are fully open in (\bar{x}, \bar{y}) , i.e., the inequality can not be violated as long as (4) is true (which it is for (x^{LP}, y^{LP}) since it is a feasible LP solution). Equality (4) is true for (\hat{x}, \hat{y}) since all assignment of j is moved to some fully open center. Thus (\hat{x}, \hat{y}) is a feasible solution for the LP. It remains to show that \hat{c} is small enough, which depends on the objective.

k -median and k -means. We start by showing this for k -median (where the distances are a metric, i.e., $\beta = 1$ in the β -triangle inequality (1)) and k -means (where the distances are a semi-metric with $\beta = 2$). We observe that here, the cost of (\hat{x}, \hat{y}) is

$$\hat{c} = \sum_{j \in P} \sum_{i \in L} \hat{x}_{ij} d(i, j) = \sum_{j \in P} \sum_{i \in L} \sum_{i' \in \bar{C}(i)} x_{i'j}^{LP} d(i, j).$$

Now fix $i \in L$, $i' \in \bar{C}(i)$ and $j \in P$ arbitrarily. By the β -relaxed triangle inequality, $d(i, j) \leq \beta \cdot d(i', j) + \beta \cdot d(i', i)$. Furthermore, we know that $i' \in \bar{C}(i)$, i.e., $\bar{\phi}(i') = i$ and $d(i', i) \leq d(i', \bar{\phi}(j))$. We can use this to relate $d(i', i)$ to the cost that j pays in (\bar{x}, \bar{y}) :

$$d(i', i) \leq d(i', \bar{\phi}(j)) \leq \beta \cdot d(j, i') + \beta \cdot d(j, \bar{\phi}(j)).$$

Adding this up yields

$$\begin{aligned} & \sum_{j \in P} \sum_{i \in L} \sum_{i' \in \bar{C}(i)} x_{i'j}^{LP} d(i, j) \\ & \leq \sum_{j \in P} \sum_{i \in L} \sum_{i' \in \bar{C}(i)} (\beta + \beta^2) x_{i'j}^{LP} d(i', j) + \sum_{j \in P} \sum_{i \in L} \sum_{i' \in \bar{C}(i)} \beta^2 \cdot x_{i'j}^{LP} d(j, \bar{\phi}(j)) \\ & = (\beta + \beta^2) \cdot c^{LP} + \beta^2 \cdot \bar{c}. \end{aligned}$$

For $\beta = 1$ (k -median), this is $2c^{LP} + \bar{c}$, for $\beta = 2$ (k -means), we get $12c^{LP} + 8\bar{c}$

Facility location. For facility location, we have to include the facility opening costs. We open the facilities that are open in (\bar{x}, \bar{y}) , which incurs a cost of $\sum_{i \in L} \bar{y}_i f_i$. The distance costs are the same as for k -median, so we get a total cost of

$$\sum_{j \in P} \sum_{i \in L} \sum_{i' \in \bar{C}(i)} 2x_{i'j}^{LP} d(i', j) + \sum_{j \in P} \sum_{i \in L} \sum_{i' \in \bar{C}(i)} x_{i'j}^{LP} d(j, \bar{\phi}(j)) + \sum_{i \in L} \bar{y}_i f_i \leq 2c^{LP} + \bar{c}.$$

k -center and k -supplier. For the k -center and k -supplier proof, we again fix $i \in L, i' \in \bar{C}(i)$ and $j \in P$ arbitrarily and use that $d(i, j) \leq d(i, i') + d(i', j)$. Now for k -center, we know that $d(i, i') \leq \bar{c}$ since $i' \in \bar{C}(i)$, and we know that $d(i', j) \leq c^{LP}$ for all j where $x_{i'j}^{LP}$ is strictly positive. Thus, if \hat{x}_{ij} is strictly positive, then $d(i, j) \leq \bar{c} + c^{LP}$. For k -supplier, we have no guarantee that $d(i, i') \leq \bar{c}$ since i' is not necessarily an input point. Instead, $i' \in \bar{C}(i)$ means that the point j' in P which is closest to i' is assigned to i by \bar{x} . Since j' is the closest to i' in P , we have $d(i', j') \leq d(i', j)$. Furthermore, since $j' \in \bar{C}(i)$, $d(i, j') \leq \bar{c}$. Thus, we get for k -supplier that

$$d(i, j) \leq d(i, i') + d(i', j) \leq d(i, j') + d(i', j') + d(i', j) \leq \bar{c} + 2 \cdot c^{LP}.$$

□

A figure for the rounding of the x variables

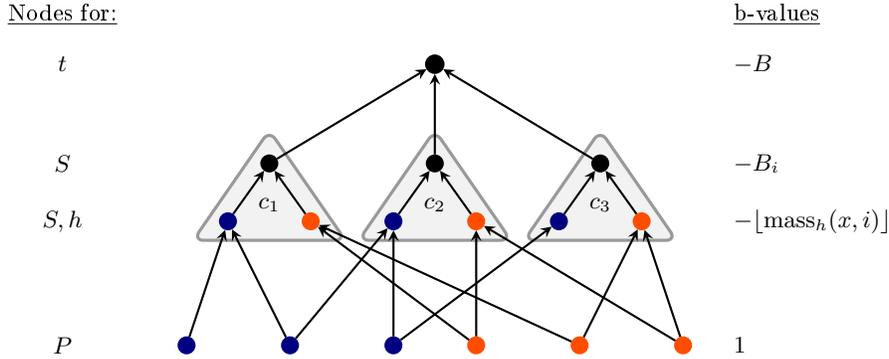


Fig. 1. Example for the graph G used in the rounding of the x -variables. $B_i = \lfloor \text{mass}(x, i) \rfloor - \sum_{h \in Col} \lfloor \text{mass}_h(x, i) \rfloor$ and $B = |P| - \sum_{i \in S} \lfloor \text{mass}(x, i) \rfloor$.