The Complexity of Finding Fair Many-To-One Matchings

Niclas Boehmer
Algorithmics and Computational Complexity, Technische Universität Berlin, Germany

Tomohiro Koana
Algorithmics and Computational Complexity, Technische Universität Berlin, Germany

Abstract

We analyze the (parameterized) computational complexity of "fair" variants of bipartite many-to-one matching, where each vertex from the "left" side is matched to exactly one vertex and each vertex from the "right" side may be matched to multiple vertices. We want to find a "fair" matching, in which each vertex from the right side is matched to a "fair" set of vertices. Assuming that each vertex from the left side has one color modeling its attribute, we study two fairness criteria. In one of them, we deem a vertex set fair if for any two colors, the difference between the numbers of their occurrences does not exceed a given threshold. Fairness is relevant when finding many-to-one matchings between students and colleges, voters and constituencies, and applicants and firms. Here colors may model sociodemographic attributes, party memberships, and qualifications, respectively.

We show that finding a fair many-to-one matching is NP-hard even for three colors and maximum degree five. Our main contribution is the design of fixed-parameter tractable algorithms with respect to the number of vertices on the right side. Our algorithms make use of a variety of techniques including color coding. At the core lie integer linear programs encoding Hall like conditions. To establish the correctness of our integer programs, we prove a new separation result, inspired by Frank's separation theorem [Frank, Discrete Math. 1982], which may also be of independent interest. We further obtain complete complexity dichotomies regarding the number of colors and the maximum degree of each side.

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms

Keywords and phrases Graph theory, polynomial-time algorithms, NP-hardness, FPT, ILP, color coding, submodular and supermodular functions, algorithmic fairness

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.27

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: https://arxiv.org/abs/2206.06988

Funding *Niclas Boehmer*: Supported by the Deutsche Forschungsgemeinschaft (DFG) projects MaMu (NI 369/19) and ComSoc-MPMS (NI 369/22).
Tomohiro Koana: Supported by the Deutsche Forschungsgemeinschaft (DFG) project DiPa (NI 369/21).

Acknowledgements We are grateful to the anonymous *ICALP 2022* reviewers for their thoughtful, constructive, and helpful comments.

1 Introduction

A many-to-one matching in a bipartite graph $G = (U \cup V, E)$ is an edge subset $M \subseteq E$ such that each vertex in U is incident to at most one edge in M. We study the computational complexity of finding a "fair" many-to-one matching and call this problem Fair Matching: Given a bipartite graph $G = (U \cup V, E)$ in which every vertex in U is colored, it asks for a many-to-one matching M such that for each $v \in V$ the vertices matched to v meet a

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).
Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;
Article No. 27; pp. 27:1–27:18
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

fairness criterion. In this work, we require that M is "left-perfect", i.e., every vertex in U is incident to exactly one edge in M. Using a slightly different formulation, Stoica et al. [24] recently studied this problem in terms of a fairness requirement derived from *margin of victory* (MoV). Generally, the margin of victory of a multiset is defined as the number of occurrences of the most frequent value minus the number of occurrences of the second most frequent value. (Given a set of colored vertices, we obtain a multiset from the occurrences of colors in it.) By requiring that the margin of victory of a set of colored vertices shall not exceed a given threshold, we prevent one color from becoming a dominating majority (see Stoica et al. [24] for a more extensive motivation of this concept). As an alternative simple fairness measure, we consider Max-Min, which is defined as the difference between the number of occurrences of the most frequent value and the number of occurrences of the least frequent value in a multiset. In a set of colored vertices with a small value of Max-Min all colors appear more or less equally often.

Which of MoV or Max-Min is more appropriate depends not only on the specific application (as discussed in the next two paragraphs) but also on the underlying data. Suppose that we have $2n$ red, $2n$ blue, and one green vertex. Then, it would be natural to deem a subset consisting of n red, n blue, and one green vertex fair (as it is in some sense the best we can hope for). Now, MoV seems to be a better fit because the MoV of the described subset is zero whereas the Max-Min value is $n-1$. In contrast, if there are $2n$ red, $2n$ blue, and $2n$ green vertices, then the same subset with n red, n blue, and one green vertex should be considered as unfair, rendering Max-Min more suitable for this color distribution than MoV. In general, Max-Min seems to be a natural choice for homogeneous data. The first example illustrates, however, that in some scenarios, MoV may serve as a viable relaxation of Max-Min.[1]

A notable application of Fair Matching emerges in the context of district-based elections. In such elections, voters (modeled by vertices in U) are divided into constituencies (modeled by vertices in V), and then each constituency elects its own representative. Here, colors can represent various attributes. For instance, colors may represent political standings. A small margin of victory is particularly desirable in this case because it will lead to close elections, holding politicians accountable for their job. One could also strive for "fair" representation of different ethnic groups or age groups by modeling ethnicity or age with colors. Other applications include the assignment of school children to schools (where colors may model sociodemographic attributes) or the assignment of reviewers to academic papers (where colors may model the level of expertise or academic background of reviewers).

Similar fairness considerations also arise in modern online systems (see, e.g., [23] for a survey). For instance, fairness is a pressing issue to counter targeted advertising or to improve recommender systems. Here one task is to ensure that the content (each perhaps represented by multiple vertices in U) falling into different categories (colors) is assigned to users (vertices in V) in a way that each user is presented with a "diverse" selection of content. Lastly, we mention that (Max-Min) Fair Matching has also applications outside of the "fairness" context: Imagine a centralized job market for companies (vertices in V) and applicants (vertices in U), each having a specific skillset (color). Firms may wish to balance between applicants with different skillsets so that employees with various skillsets may be placed into teams. For instance, it may be desirable for a software company to hire roughly the same number of frontend and backend developers.

[1] Note that the Max-Min value is at least the MoV value for any multiset.

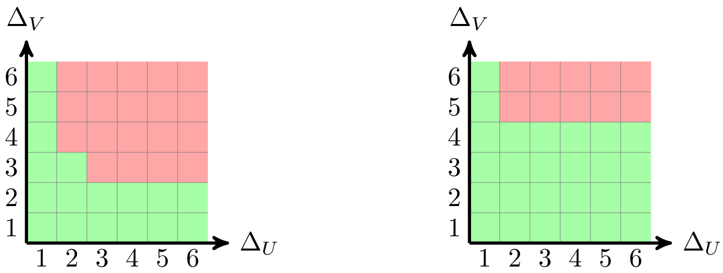

Figure 1 The complexity landscape of MAX-MIN FAIR MATCHING (left) and MOV FAIR MATCHING (right) for the maximum degree Δ_U (resp., Δ_V) in U (resp., V). Green denotes polynomial-time solvability and red denotes NP-hardness. All NP-hardness results hold for three colors.

Our Contributions. We perform a refined complexity analysis of the NP-hard MAX-MIN FAIR MATCHING and MOV FAIR MATCHING problems in terms of the size k of V, the number $|C|$ of colors, and the maximum degree Δ_U and Δ_V among vertices in U and V, respectively. Our main contribution are arguably involved FPT algorithms for the parameterization k (Section 4). At the heart of the design of our algorithms lies an integer linear program (ILP) of bounded dimension. We essentially determine whether Hall-like conditions that guarantee the existence of a fair matching are fulfilled by formulating these in a system of linear inequalities. In order to establish the correctness of our ILP formulations, we prove what we call *touching separation theorem*, getting inspiration from Frank's separation theorem on submodular and supermodular functions [14]. For MOV FAIR MATCHING, we apply our approach in conjunction with the color coding technique [5]. To familiarize ourselves with the ideas underlying our ILPs, in Section 3, we start with a warm-up where we present ILP-based fixed-parameter tractable algorithms for the larger parameter $k + |C|$. To sum up, as it is straightforward to see that MAX-MIN/MOV FAIR MATCHING are FPT with respect to the size n of U[2], we establish the fixed-parameter tractability of both problems for the two natural parameters n and k.[3] We then in Section 5 study the computational complexity of MAX-MIN/MOV FAIR MATCHING with respect to Δ_U, Δ_V, and $|C|$. We show that MAX-MIN/MOV FAIR MATCHING is polynomial-time solvable for $|C| = 2$ and that it becomes NP-hard for $|C| \geq 3$. Moreover, we settle all questions concerning the problems' classical complexity in terms of Δ_U and Δ_V, revealing a complete complexity landscape in this regard (see Figure 1). Finally, in Section 6, we show that MAX-MIN/MOV FAIR MATCHING are linear-time solvable when every vertex in U can be matched to any vertex in V. Notably, all our algorithmic results hold even if we require that each vertex from V is matched to at least one vertex from U. This further constraint may appear when we need to divide the vertices into exactly k non-empty fair subsets. Although this constraint is seemingly simple, sometimes (e.g., in our FPT algorithm for MAX-MIN FAIR MATCHING for k) non-trivial adaptions are needed.

Related Work. Stoica et al. [24] introduced three problems where the task is to partition a set of colored vertices into subsets with a small margin of victory satisfying some global size constraints. Among these three, the most general is FAIR CONNECTED REGROUPING,

[2] We can enumerate all fair subsets of U in $2^n \cdot (n+k)^{O(1)}$ time. Then, Knapsack-like dynamic programming solves FAIR MATCHING in $3^n \cdot (n+k)^{O(1)}$ time.

[3] In most described applications k is typically quite small and much smaller than n. For instance, Stoica et al. [24] performed some experiments for MOV FAIR MATCHING to assign voters to districts with $n = 50,000$ and $k = 10$, and to assign students to schools with $n = 41,834$ and $k = 61$.

where one is given a vertex-colored graph G, an integer k, and a function that determines for each vertex which subsets it can be part of. The task is then to find a partitioning of G into k fair districts (i.e., connected components). In a follow-up work, Boehmer et al. [9] analyzed how the structure of G influences the (parameterized) computational complexity of Fair Connected Regrouping. The other two problems Stoica et al. [24] considered are special cases of Fair Connected Regrouping: One is Fair Regrouping where the connectivity constraints are dropped (corresponding to MoV Fair Matching). The other is Fair Regrouping_X, which is a special case of Fair Regrouping, where any vertex can belong to any district (corresponding to MoV Fair Matching on complete bipartite graphs; we study this special case in Section 6). They proved that Fair Regrouping is NP-hard for three colors (without any constraints on the degree of the graph) and is XP with respect to the number k of districts (i.e., polynomial-time solvable for constant k). They also showed that Fair Regrouping_X is XP with respect to the number of colors.

Coming back to Fair Matching as a matching problem, Ahmed et al. [4] proposed a global supermodular objective to model the fairness (which they call diversity) of a bipartite weighted many-to-many matching and developed a polynomial-time greedy heuristic for it. Ahmadi et al. [1] extended the work of Ahmed et al. [4] by generalizing the problem to the case where vertices can have multiple different colors and presented a pseudo-polynomial-time algorithm for it. Moreover, Dickerson et al. [13] applied this formulation of fairness to an online setting where vertices from the left side arrive over time and Ahmed et al. [3] applied it to the task of forming teams.

Fairness is also a popular topic when finding a stable many-to-one matching of vertices that have preferences over each other. Here, fairness constraints are typically modeled by imposing for each vertex from the right side certain lower and upper bounds on the number of vertices of each color that can be matched to it [7, 8, 10, 16, 18].

More broadly speaking, fairness has recently also been frequently applied to a variety of different problems from the area of combinatorial optimization. For instance, in the context of the Knapsack [22] or Maximum Coverage [6] problem, fairness means that all types are represented equally in the selected solution. For clustering, each cluster is considered fair when each type accounts for a certain fraction of vertices in it [2, 15].

The proof (or their completion) of all results marked by ($\star$) are omitted.

2 Preliminaries

For two integers $i < j \in \mathbb{N}$, let $[i, j] = \{i, i+1, \ldots, j-1, j\}$ and let $[i] = [1, i]$. For a set S and an element $x \in S$, we sometimes write $S - x$ to denote $S \setminus \{x\}$.

Let $G = (U \cupdot V, E)$ be a bipartite graph, where U is the *left side* and V is the *right side* of G. Let $n := |U|$ and $k := |V|$ be the number of vertices in the left side and right side, respectively. For a vertex $w \in U \cupdot V$ and an edge set $M \subseteq E$, let $M(w)$ be the set of vertices *matched to w* in M, i.e., $M(w) = \{w' \in U \cupdot V \mid \{w, w'\} \in M\}$. We say that $M \subseteq E$ is a *many-to-one matching* in G if $|M(u)| \leq 1$ for every $u \in U$. A many-to-one matching M is *left-perfect* if $|M(u)| = 1$ for every $u \in U$. Note that we require M to be left-perfect as otherwise an empty set would constitute a trivial solution for our problem. When clear from context, we refer to a left-perfect many-to-one matching as a matching. For a vertex $w \in U \cupdot V$, let $N_G(w)$ be the set of its neighbors in G, i.e., $N_G(w) = \{w' \in U \cupdot V \mid \{w, w'\} \in E\}$. For $W \subseteq U \cupdot V$, let $N_G(W) = \bigcup_{w \in W} N_G(w)$ be the joint neighborhood of vertices from W and let $\nu_G(W) = \{w' \in U \cupdot V \mid N_G(w') \subseteq W\}$ be the set of vertices which are only adjacent to vertices in W. We drop the subscript $\cdot_G$ when it is clear from context.

Let C be the set of colors and let $\text{col}\colon U \to C$ be a function that assigns a color to every vertex of U. For $U' \subseteq U$, let $U'_c \subseteq U'$ be the set of vertices $u \in U'$ of color c. For instance, given a matching M and a vertex $v \in V$, $M(v)_c$ denotes the set of vertices matched to v in M that have color c. We denote by $G_c = G[U_c \cup V]$ the graph G restricted to vertices from $U_c \cup V$. We use the shorthand $N_c(W)$ (resp., $\nu_c(W)$) for $N_{G_c}(W)$ (resp., $\nu_{G_c}(W)$).

Throughout the paper, we assume that the set C of colors is equipped with some linear order $\leq_C$, which serves as a tie breaker. So $\arg\max$ over C is well-defined. We write $\max^1$ for max and $\max^2$ for the second largest element.

We now define our two fairness measures. For a subset of vertices $U' \subseteq U$, let $\text{MoV}(U') := \max^1_{c\in C} |U'_c| - \max^2_{c\in C} |U'_c|$ be the difference between the number of occurrences of the most and second most frequent color in U'. Similarly, for a subset of vertices $U' \subseteq U$, let $\text{MmM}(U') := \max_{c\in C} |U'_c| - \min_{c\in C} |U'_c|$ be the difference between the number of occurrences of the most and least frequent color in U'. A subset of vertices $U' \subseteq U$ is ℓ-fair according to MoV (resp., Max-Min) if $\text{MoV}(U') \leq \ell$ (resp., $\text{MmM}(U') \leq \ell$).[4] A many-to-one matching M in G is ℓ-fair according to MoV (resp., Max-Min) if $M(v)$ is ℓ-fair according to MoV (resp., Max-Min) for all $v \in V$. The considered fairness notion will always be clear from context. We now define our central problem Π Fair Matching for some fairness measure Π:

> Π Fair Matching
> **Input:** A bipartite graph $G = (U \uplus V, E)$, a set C of colors, a function $\text{col}\colon U \to C$, and an integer $\ell \in \mathbb{N}$.
> **Question:** Is there a left-perfect many-to-one matching $M \subseteq E$ which is ℓ-fair according to the fairness measure Π?

We also sometimes consider Π Fair Matching with size constraints where additionally given two integers p and q, we require that the matching M to be found satisfies $p \leq |M(v)| \leq q$ for all $v \in V$. We refer to the case with $p = 1$ and $q = n$ as the non-emptiness constraint. This constraint is arguably crucial for some applications when we want to partition the vertices in the left side into exactly k non-empty subsets.

Let $\mathcal{I}$ be an instance of some problem and let $\mathcal{P}(\mathcal{I})$ be an integer linear program (ILP) constructed from $\mathcal{I}$. We say that $\mathcal{P}$ is *complete* if $\mathcal{P}(\mathcal{I})$ is feasible whenever $\mathcal{I}$ is a yes-instance. Conversely, we say that $\mathcal{P}$ is *sound* if $\mathcal{I}$ is a yes-instance whenever $\mathcal{P}(\mathcal{I})$ is feasible. In this work, we will make use of Lenstra's algorithm [19, 20] that decides whether an ILP of size L with p variables is feasible using $O(p^{2.5p+o(p)} \cdot |L|)$ arithmetic operations.

We assume that the reader is familiar with basic concepts in parameterized complexity (see for instance [12]). As a reminder, an FPT algorithm for a parameter k is an algorithm whose running time on input $\mathcal{I}$ is $f(k) \cdot |\mathcal{I}|^{O(1)}$ for some computable function f.

3 Warmup: FPT Algorithms for $k + |C|$

We prove that both Fair Matching problems are fixed-parameter tractable with respect to $k+|C|$: We present an integer linear programming (ILP) formulation of these problems whose number of variables is bounded in a function of the parameter $k+|C|$ and subsequently employ

[4] Notably, the definition of margin of victory of Stoica et al. [24] differs slightly from ours in that in their definition sets of vertices where the two most frequent colors have the same number of occurrences have a margin of victory of one (and not of zero). We chose our definition in accordance with Boehmer et al. [9] to be able to distinguish a tie between two colors from one color being one vertex ahead of another.

Lenstra's algorithm [19, 20]. Notably, one can upper-bound the number of "types" (according to their neighborhoods and colors) of vertices in U by $2^k \cdot |C|$. Using this observation, it is straightforward to give an ILP formulation of MAX-MIN/MOV FAIR MATCHING using $O(2^k \cdot |C|)$ variables. Instead, we follow a theoretically more involved but more efficient approach. For this, we use a structural property of our problem related to Hall's theorem, which decreases the number of variables in our ILP to $O(|C| \cdot k)$. We reuse some of the results from this section in Section 4, where we prove that MAX-MIN/MOV FAIR MATCHING are actually fixed-parameter tractable with respect to k.

To prove that the ILP we present in the following is complete, we use the following:

▶ **Lemma 1** ($\star$). *Let $G = (U \cup V, E)$ be a bipartite graph and let M be a left-perfect many-to-one matching. Then, $|\nu(W)| \leq \sum_{v\in W} |M(v)| \leq |N(W)|$ for every $W \subseteq V$.*

For proving that our ILP is sound, we use the following:

▶ **Lemma 2.** *Let $G = (U \cup V, E)$ be a bipartite graph and let $\{z_v \in \mathbb{N} \mid v \in V\}$ be a set of integers. Suppose that $\sum_{v\in W} z_v \geq |\nu_G(W)|$ for every $W \subseteq V$ and that $\sum_{v\in V} z_v = |U|$. Then, there is a left-perfect many-to-one matching M such that $|M(v)| = z_v$ for every $v \in V$.*

Proof. Assume that the conditions stated in the lemma hold. We prove the existence of such a matching M by making use of Hall's theorem [17]. To do so, we introduce an auxiliary bipartite graph G' as follows: In G', the vertices on one of the two sides are the vertices from U. The vertices on the other side are $V' := \bigcup_{v\in V} Z_v$, where Z_v is a set of z_v vertices. There is an edge between $u \in U$ and $v' \in Z_v \subseteq V'$ if and only if $\{u, v\} \in E$. In order to apply Hall's theorem, we show that $|U'| \leq |N_{G'}(U')|$ for every $U' \subseteq U$.

Fix some $U' \subseteq U$ and let $W' = N_G(U')$. The construction of G' gives us $|N_{G'}(U')| = \sum_{v\in W'} z_v$. By the assumption of the lemma, we have $\sum_{v\in W'} z_v \geq |\nu_G(W')|$. Moreover, we have $\nu_G(W') = \nu_G(N_G(U')) \supseteq U'$. Consequently, we obtain $|N_{G'}(U')| = \sum_{v\in W'} z_v \geq |\nu_G(W')| \geq |U'|$. Hall's theorem then implies that G' admits a one-to-one matching M' which matches all vertices from U. In fact, M' is a perfect matching since $|V'| = \sum_{v\in V} z_v = |U|$ by our assumption. Now consider the matching M in G where a vertex $u \in U$ is matched to $v \in V$ if u is matched to a vertex from Z_v in M'. Then, M is a left-perfect many-to-one matching with $|M(v)| = z_v$ for every $v \in V$. ◀

Using Lemmas 1 and 2, we give an ILP formulation of FAIR MATCHING with $O(|C| \cdot k)$ variables.

ILP formulation. Introduce a variable $z_v^c \in \mathbb{N}$ for every $v \in V$ and every $c \in C$. The variable z_v^c represents the number of vertices in U of color c that are matched to v. Suppose that the given instance admits a left-perfect ℓ-fair many-to-one matching M respecting the values of z_v^c. Then, for each $c \in C$, there is a matching M_c in G_c with $|M_c(v)| = z_v^c$ for $v \in V$ and $|M_c(u)| = 1$ for $u \in U_c$. As shown in Lemma 1, from this one can conclude that the following constraints must be fulfilled:

$$|\nu_c(W)| \leq \sum_{v\in W} z_v^c \leq |N_c(W)| \text{ for all } W \subseteq V, c \in C.$$

Next, we encode the fairness requirement. For MAX-MIN, we need to have that for every pair of colors the number of vertices of these two colors assigned to some vertex $v \in V$ differ by at most ℓ. Thus, we add the constraint

$$z_v^{c'} - z_v^c \leq \ell \text{ for all } v \in V \text{ and } c, c' \in C.$$

To model ℓ-fairness for MoV Fair Matching, we introduce two new binary variables $a_v^c, b_v^c \in \{0,1\}$ for all $v \in V$ and $c \in C$. Informally speaking, the intended meaning of these variables is that $a_v^c = 1$ (resp., $b_v^c = 1$) when c is the most (resp., second most) frequent color among vertices matched to v. We ensure that the values of these variables are set accordingly as follows:

$$z_v^c - z_v^{c'} \geq n(a_v^c + b_v^{c'} - 2), \quad z_v^c - z_v^{c'} \geq n(b_v^{c'} - a_v^c - 1), \text{ and } a_v^c + b_v^c \leq 1 \quad \forall v \in V, c, c' \in C;$$
$$\sum_{c \in C} a_v^c = \sum_{c \in C} b_v^c = 1 \quad \forall v \in V.$$

For the first constraint, note that it becomes $z_v^c \geq z_v^{c'}$ if $a_v^c = b_v^{c'} = 1$ and that it is always fulfilled otherwise. Similarly for the second constraint, observe that it becomes $z_v^c \geq z_v^{c'}$ if $a_v^{c'} = 0$ and $b_v^c = 1$ and that it is always fulfilled otherwise.

Finally, using a_v^c and b_v^c (and their meaning as proven above), we add the following constraint that encodes the ℓ-fairness in terms of margin of victory of $M(v)$ for all $v \in V$:

$$z_v^c - z_v^{c'} - n(2 - a_v^c - b_v^{c'}) \leq \ell \qquad \forall v \in V$$

Lastly, we can also add linear constraints ensuring that the number of vertices from U matched to each vertex $v \in V$ is between p and q: $p \leq \sum_{c \in C} z_v^c \leq q$ for all $v \in V$.

▶ **Theorem 3.** *Max-Min/MoV Fair Matching with arbitrary size constraints can be solved in $O^\star((|C| \cdot k)^{O(|C| \cdot k)})$ time.*

Proof. We show that an instance $(G = (U \cup V, E), C, \mathrm{col}, \ell, p, q)$ of Max-Min/MoV Fair Matching with size constraints admits an ℓ-fair left-perfect many-to-one matching if and only if the constructed ILP constructed is feasible. As described above, if the given instance is a yes-instance, there is an assignment to z_v^c satisfying all the constraints. Conversely, suppose that the ILP admits a solution $\{z_v^c \mid v \in V, c \in C\}$. Then, from our first set of constraints, we have for every $c \in C$ that $\sum_{v \in W} z_v^c \geq |\nu_c(W)|$ for every $W \subseteq V$ and $\sum_{v \in V} z_v^c = |U_c|$ (the later part follows from our first set of constraints for $W = V$). By Lemma 2, it follows that G_c has a matching M_c in which every vertex in U_c is matched and the values of z_v^c for $v \in V$ and $c \in C$ are respected. Aggregating M_c for all colors c yields a left-perfect ℓ-fair many-to-one matching for G respecting the given size constraints. Using Lenstra's algorithm [19, 20], the feasibility of the ILP can be determined in the claimed time. ◀

4 FPT Algorithms for k

In this section, we develop FPT algorithms for Max-Min/MoV Fair Matching for the parameterization $|V| = k$. We start with a discussion of the challenges for our algorithms. Afterwards, we obtain a new result on submodular and supermodular functions. Using this, in Section 4.1 we present the algorithm for Max-Min Fair Matching, and in Section 4.2 the slightly more involved algorithm for MoV Fair Matching.

The crux of our algorithms is an ILP as in Section 3. However, since we look into the parameterization without $|C|$, it would be too costly to introduce variables for each color. To illustrate our idea to work around this issue, take Max-Min Fair Matching as an example. One of the straightforward ideas how to formulate this problem as an ILP would be to introduce two variables $x_v \leq y_v$ for every vertex $v \in V$, where x_v (resp., y_v) encodes the minimum (resp., maximum) number of vertices of some color $c \in C$ matched to v. Informally speaking, to encode the ILP constraints from Section 3, we could now replace every occurrence of z_v^c with x_v (resp., y_v) if the constraint in which z_v^c occurs imposes an upper (resp., lower) bound on z_v^c: the first set of constraints added in Section 3 translates to $\sum_{v\in W} y_v \geq \max_{c\in C} |\nu_c(W)|$ and $\sum_{v\in W} x_v \leq \min_{c\in C} |N_c(W)|$ for every $W \subseteq V$. Moreover, we add the constraint $y_v - x_v \leq \ell$ for all $v \in V$. Let $\mathcal{P}$ denote the ILP obtained this way. (See Section 4.1 for the formal construction of $\mathcal{P}$.) Although it is easy to see that $\mathcal{P}$ is complete, it turns out to be nontrivial to prove its soundness. To highlight this challenge, suppose that $\mathcal{P}$ is feasible for $\{x_v, y_v \mid v \in V\}$. To show that this implies that there is an ℓ-fair matching M, we have to show that for every color $c \in C$, the graph G_c has a matching M_c such that $x_v \leq |M_c(v)| \leq y_v$ for all $v \in V$. Unfortunately, we cannot directly apply Lemma 2 (as we were able to do in Section 3): To construct an ℓ-fair matching from $\{x_v, y_v \mid v \in V\}$ using Lemma 2, we have to show that for every color $c \in C$, there always exists a set of integers $\{z_v^c \mid v \in V\}$ with the following properties:

(i) $\sum_{v\in W} z_v^c \geq |\nu_c(W)|$, $\forall W \subseteq V$
(ii) $\sum_{v\in V} z_v^c = |U_c|$
(iii) $x_v \leq z_v^c \leq y_v$, $\forall v \in V$.

Finding an assignment of variables z_v^c that fulfill (i) and (iii) is trivial; setting $z_v^c = y_v$ suffices because $\mathcal{P}$ dictates that $\sum_{v\in W} y_v \geq \max_{c'\in C} |\nu_{c'}(W)| \geq |\nu_c(W)|$ for every $W \subseteq V$. However, integers satisfying (i), (ii), and (iii) simultaneously are not trivially guaranteed to exist. To nevertheless prove their existence, we prove what we call the *touching separation theorem*, inspired by Frank's separation theorem [14] on submodular and supermodular functions. Our theorem implies that if there is a solution to $\mathcal{P}$, then there is an assignment of variables z_v^c satisfying (i), (ii), and (iii). Submodular and supermodular functions are defined as follows:

▶ **Definition 4.** *Let $f\colon 2^S \to \mathbb{N}$ be a set function over a set S. We say that f is submodular if $f(X)+f(Y) \geq f(X\cup Y)+f(X\cap Y)$ for every $X, Y \subseteq S$ and supermodular if $f(X)+f(Y) \leq f(X\cup Y) + f(X\cap Y)$ for every $X, Y \subseteq S$. Moreover, f is modular if f is both submodular and supermodular.*

Modular functions admit a simpler characterization, which will be useful in several proofs:

▶ **Lemma 5** (folklore, $\star$). *Let $f\colon 2^S \to \mathbb{N}$ be a set function. Then, f is modular if and only if $f(X) = f(\emptyset) + \sum_{x\in X}(f(x) - f(\emptyset))$ for every $X \subseteq S$.*

Frank's separation theorem [14] can be stated as follows: If $f\colon 2^S \to \mathbb{N}$ is submodular and $g\colon 2^S \to \mathbb{N}$ is supermodular, and $g(X) \leq f(X)$ for every $X \subseteq S$, then there exists a modular set function $h\colon 2^S \to \mathbb{N}$ such that $g(X) \leq h(X) \leq f(X)$ for every $X \subseteq S$. We prove an analogous separation theorem where the "lower-bound" function is the maximum of two functions – one being modular and the other being supermodular – and the "upper-bound" function is a modular function. (In general, the "lower-bound" function arising this way is neither submodular nor supermodular.) We additionally require that the function separating the two functions "touches" the lower-bound function. Our proof uses a rather involved induction on the size of S.

▶ **Theorem 6** (Touching separation theorem, $\star$). *Let $f, f'\colon 2^S \to \mathbb{N}$ be two modular set functions and let $g\colon 2^S \to \mathbb{N}$ be a supermodular set function. Suppose that $f(\emptyset) = f'(\emptyset) = g(\emptyset) = 0$ and $\max(f(X), g(X)) \leq f'(X)$ for every $X \subseteq S$. Then, there is a modular set function $h\colon 2^S \to \mathbb{N}$ such that $\max(f(X), g(X)) \leq h(X) \leq f'(X)$ for every $X \subseteq S$ and $h(S) = \max_{X\subseteq S} f(X) + g(S \setminus X)$.*

To apply Theorem 6, we use a supermodular function arising from a bipartite graph:

▶ **Lemma 7** ($\star$). *Let $G = (U \cupdot V, E)$ be a bipartite graph and let $g\colon 2^V \to \mathbb{N}$ be a set function such that $g(W) = |\nu(W)|$ for each $W \subseteq V$. Then, g is supermodular.*

4.1 Max-Min Fair Matching

We now present an FPT algorithm for Max-Min Fair Matching. As mentioned, our algorithm, which follows a similar approach as used in Section 3, builds an ILP. The difference is that the ILP constructed here involves $O(k)$ and not $\Theta(k \cdot |C|)$ variables. We prove the correctness of the ILP in Theorem 8, crucially relying on Theorem 6.

ILP formulation. We add two variables $x_v \leq y_v \in \mathbb{N}$ for every $v \in V$. The intended meaning for these variables is that for every color $c \in C$, the number of vertices from U_c matched to v is between x_v and y_v. We obtain the following constraints from Lemma 1:

$$\sum_{v \in W} y_v \geq \max_{c \in C} |\nu_c(W)| \text{ and } \sum_{v \in W} x_v \leq \min_{c \in C} |N_c(W)| \quad \forall W \subseteq V. \tag{1}$$

To encode ℓ-fairness, we add: $y_v - x_v \leq \ell$ for all $v \in V$.

▶ **Theorem 8.** *Max-Min Fair Matching can be solved in $O^\star(k^{O(k)})$ time.*

Proof. Our ILP uses $O(k)$ variables and $O(2^k)$ constraints; the construction takes $O^\star(2^k)$ time. Using Lenstra's algorithm [19, 20], it is possible to check whether the constructed ILP is feasible in $O^\star(k^{O(k)})$ time. It remains to prove the correctness of our ILP.

The completeness of our ILP follows from Lemma 1. For the soundness, suppose that $\{x_v, y_v \mid v \in V\}$ is a feasible solution for the ILP. For every color $c \in C$, we show the following: For the set $U_c \subseteq U$ of vertices of color c, there is a matching M_c in the bipartite graph $G_c = G[U_c \cup V]$ such that $x_v \leq |M_c(v)| \leq y_v$ for each $v \in V$, from which the soundness of the ILP directly follows. To show the existence of such a matching M_c, we will rely on Lemma 2. Note, however, that Lemma 2 asks for integers $\{z_v \mid v \in V\}$ meeting certain constraints, which we cannot choose arbitrarily as we have to respect the constraint $x_v \leq z_v \leq y_v$ for every $v \in V$. Let us fix some color $c \in C$. We find integers $\{z_v \mid v \in V\}$ via Theorem 6:

Let $f, f', g\colon 2^V \to \mathbb{N}$ be set functions such that $f(W) = \sum_{v \in W} x_v$ and $f'(W) = \sum_{v \in W} y_v$, and $g(W) = |\nu_c(W)|$ for each $W \subseteq V$. Note that f and f' are modular by Lemma 5 and that g is supermodular by Lemma 7. The constraints in the ILP ensure that $\max(f(W), g(W)) \leq f'(W)$ for every $W \subseteq V$. Consequently, Theorem 6 yields a set modular function $h\colon 2^V \to \mathbb{N}$ such that $\max(f(W), g(W)) \leq h(W) \leq f'(W)$ for each $W \subseteq V$ and $h(V) = \max_{W \subseteq V} f(W) + g(V \setminus W)$. Let $z_v = h(v)$ for every $v \in V$. Note that as h is modular and fulfills the constraints stated above, $\sum_{v \in W} z_v = h(W) \geq g(W) = |\nu_c(W)|$ for each $W \subseteq V$. Hence, in order to apply Lemma 2 on the integers $\{z_v \mid v \in V\}$, it remains to show that $\sum_{v \in V} z_v = h(V) = |U_c|$. Since our ILP requires that $f(W) = \sum_{v \in W} x_v \leq |N_c(W)|$ for $W \subseteq V$, we have that

$$\begin{aligned} f(W) + g(V \setminus W) &\leq |N_c(W)| + |\nu_c(V \setminus W)| \\ &= |\{u \in U_c \mid N_c(u) \cap W \neq \emptyset\}| + |\{u \in U_c \mid N_c(u) \subseteq V \setminus W\}| = |U_c|. \end{aligned}$$

Moreover, we have $f(\emptyset) + g(V) = |U_c|$, resulting in $h(V) = \max_{W \subseteq V} f(W) + g(V \setminus W) = |U_c|$ by Theorem 6. Therefore, the graph G_c admits a matching M_c with $x_v \leq |M_c(v)| \leq y_v$ for each $v \in V$ by Lemma 2. Combining these matchings yields an ℓ-fair matching in G. ◀

Non-emptiness constraint. For a matching M found by our algorithm, we may have $M(v) = \emptyset$ for some $v \in V$. So the non-emptiness constraint may be violated. Unfortunately, there is seemingly no simple linear constraint to ensure that $M(v) \neq \emptyset$ for each $v \in V$.[5]

[5] Adding $x_v > 0$ ensures that there is at least one vertex of each color matched to v. Adding $y_v > 0$ is not enough, as the ILP only ensures that for each color *at most* y_v vertices are matched to v.

Overcoming this challenge, we now develop an FPT algorithm for Max-Min Fair Matching with the non-emptiness constraint. We will build upon the ILP formulation for Theorem 8. First, observe that for $\ell = 0$, it suffices to add the constraint $y_v > 0$ for all $v \in V$ because this ensures that every vertex $v \in V$ is matched to at least $x_v = y_v > 0$ vertices of each color.

For $\ell > 0$, we develop a more involved algorithm. We will only describe its main ideas here. Suppose that there is an ℓ-fair many-to-one matching M such that $M(v) \neq \emptyset$ for each $v \in V$. By choosing an arbitrary element u_v of $M(v)$ for each $v \in V$, we obtain a matching $M^\star := \{\{u_v, v\} \mid v \in V\} \subseteq M$ with $|M^\star(v)| = 1$ for each $v \in V$. Since there are possibly $n^{\Omega(k)}$ choices for $M^\star$, we cannot assume that $M^\star$ is given. Instead, our algorithm only "guesses" some structural properties of $M^\star$, which we will incorporate into the above ILP for Max-Min Fair Matching by adding further constraints. (Our algorithm avoids guessing objects marked with a star.)

For each $v \in V$, let $\chi(v)$ be the color of $M^\star(v)$. By iterating over all possible partitions $\mathcal{V}$ of V, we first guess a partition of V according to $\chi(v)$, i.e., two vertices $v, v' \in V$ belong to the same subset of $\mathcal{V}$ if and only if $\chi(v) = \chi(v')$. For each $S \in \mathcal{V}$, let $\chi(S)$ be the color of $\chi(v)$ for an arbitrary vertex $v \in S$. Let us fix some $S \in \mathcal{V}$. We will formulate constraints that must be fulfilled when each vertex v in S is matched to $M^\star(v)$ (which has color $\chi(S)$). For this, let $U_S^\star \subseteq U_{\chi(S)}$ be the set of vertices from U of color $\chi(S)$ incident to an edge in $M^\star$ and let $G_S^\star$ be the graph obtained from $G_{\chi(S)}$ by deleting all edges incident to $U_S^\star$ and then adding the edges of $M^\star$ whose endpoint on the left side has color $\chi(S)$. Since every edge in M with an endpoint of color $\chi(S)$ is present in $G_S^\star$, these edges form a left-perfect many-to-one matching in $G_S^\star$. Thus, by Lemma 1, we should have $\sum_{v \in W} y_v \geq |\nu_{G_S^\star}(W)|$ and $\sum_{v \in W} x_v \leq |N_{G_S^\star}(W)|$ for each $W \subseteq V$ if there is a left-perfect matching containing $M^\star$ in $G_S^\star$. We need to evaluate $|\nu_{G_S^\star}(W)|$ and $|N_{G_S^\star}(W)|$ to include these constraints into our ILP. Note, however, that we cannot compute these values without $M^\star$ given. In the following, we explain how we can nevertheless incorporate these constraints by guessing further structural aspects of $M^\star$.

First, we rewrite $|\nu_{G_S^\star}(W)|$ and $|N_{G_S^\star}(W)|$ as follows:

$$|\nu_{G_S^\star}(W)| = |\nu_{\chi(S)}(W)| + |\{v \in S \cap W \mid N_{G_{\chi(S)}}(M^\star(v)) \not\subseteq W\}| \text{ and}$$
$$|N_{G_S^\star}(W)| = |N_{\chi(S)}(W)| - |\{v \in S \setminus W \mid N_{G_{\chi(S)}}(M^\star(v)) \cap W \neq \emptyset\}|.$$

To see why the first equation holds, observe that we have only deleted edges when constructing $G_S^\star$ from $G_{\chi(S)}$. Thus, we have $\nu_{G_{\chi(S)}}(W) \subseteq \nu_{G_S^\star}(W)$. Moreover, we have $u \in \nu_{G_S^\star}(W) \setminus \nu_{G_{\chi(S)}}(W)$ if and only if $u = M^\star(v)$ for some $v \in S \cap W$ (which implies that u is only adjacent to v in $G_S^\star$) and u has a neighbor outside of W in $G_{\chi(S)}$. The second equation follows similarly: As we only deleted edges, we have $N_{G_S^\star}(W) \subseteq N_{G_{\chi(S)}}(W)$. We also have $u \in N_{G_{\chi(S)}}(W) \setminus N_{G_S^\star}(W)$ if and only if $u = M^\star(v)$ for some $v \in S \setminus W$ and u has a neighbor in W in $G_{\chi(S)}$. To evaluate the second term in each of these equations, we guess a function $\mu \colon V \to 2^V$ such that $\mu(v) = N_{G_{\chi(S)}}(M^*(v))$ for each $v \in V$. For $\alpha(S, W) := |\{v \in S \cap W \mid \mu(v) \not\subseteq W\}|$ and $\beta(S, W) := |\{v \in S \setminus W \mid \mu(v) \cap W \neq \emptyset\}|$, we then have:

$$|\nu_{G_S^\star}(W)| = |\nu_{\chi(S)}(W)| + \alpha(S, W) \text{ and } |N_{G_S^\star}(W)| = |N_{\chi(S)}(W)| - \beta(S, W). \tag{2}$$

To incorporate the constraints $\sum_{v \in W} y_v \geq |\nu_{G_S^\star}(W)|$ and $\sum_{v \in W} x_v \leq |N_{G_S^\star}(W)|$, it remains to deal with $|\nu_{\chi(S)}(W)|$ and $|N_{\chi(S)}(W)|$. Note that we cannot compute $|\nu_{\chi(S)}(W)|$ or $|N_{\chi(S)}(W)|$ without $M^\star$ given. Moreover, it would be costly to guess these values (which can be $\Omega(n)$) for every $S \in \mathcal{V}$ and $W \subseteq V$ or the values of $\chi(S)$ (which can be $\Omega(|C|)$) for every $S \in \mathcal{V}$. Instead, we relate these two values to $\max_{c\in C}|\nu_c(W)|$ and $\min_{c\in C}|N_c(W)|$ by guessing their respective differences. Although these differences may be $\Omega(n)$, we discover that we can cap them at k: When they are larger, then the arising constraints are already covered by Constraint (1) from the original ILP. Formally, we guess $X(S,W) = \min(|N_{\chi(S)}(W)| - \min_{c\in C}|N_c(W)|, k)$ and $Y(S,W) = \min(\max_{c\in C}|\nu_c(W)| - |\nu_{\chi(S)}(W)|, k)$ for each $S \in \mathcal{V}$ and $W \subseteq V$ by iterating over all possible $X, Y\colon \mathcal{V} \times 2^V \to \{0,\dots,k\}$ (for each of them we have at most $(k+1)^{k\cdot 2^k} \in O(2^{2^{O(k)}})$ choices).

To account for the constraints $\sum_{v\in W} y_v \geq |\nu_{G^\star_S}(W)|$ and $\sum_{v\in W} x_v \leq |N_{G^\star_S}(W)|$, we now add the following constraint to the ILP for each $S \in \mathcal{V}$ and $W \subseteq V$:

$$\sum_{v\in W} y_v \geq \max_{c\in C}|\nu_c(W)| - Y(S,W) + \alpha(S,W) \text{ and}$$
$$\sum_{v\in W} x_v \leq \min_{c\in C}|N_c(W)| + X(S,W) - \beta(S,W).$$

To see why these constraints must be fulfilled, we consider two cases. If $|N_{\chi(S)}(W)| - \min_{c\in C}|N_c(W)| \leq k$ (resp., $\max_{c\in C}|\nu_c(W)| - |\nu_{\chi(S)}(W)| \leq k$), then by the definition of $X(S,W)$ (resp. $Y(S,W)$) and Equation (2), we have $|N_{G^\star_S}(W)| = \min_{c\in C}|N_c(W)| + X(S,W) - \beta(S,W)$ (resp., $|\nu_{G^\star_S}(W)| = \max_{c\in C}|\nu_c(W)| - Y(S,W) + \alpha(S,W)$). Otherwise, we claim that the constraint $\sum_{v\in W} x_v \leq |N_{G^\star_S}(W)|$ (resp. $\sum_{v\in W} y_v \geq |\nu_{G^\star_S}(W)|$) is already captured by Constraint (1), since $X(S,W) = k \geq \beta(S,W)$ (resp., $Y(S,W) = k \geq \alpha(S,W)$), we have $\sum_{v\in W} x_v \leq \min_{c\in C}|N_c(W)| \leq \min_{c\in C}|N_c(W)| + X(S,W) - \beta(S,W)$ (resp., $\sum_{v\in W} y_v \geq \max_{c\in C}|\nu_c(W)| \geq \max_{c\in C}|\nu_c(W)| - Y(S,W) + \alpha(S,W)$).

Iterating over all described guesses and for each solving the constructed ILP, we get:

▶ **Theorem 9** (⋆). Max-Min Fair Matching *with the non-emptiness constraint can be solved in* $O^*(2^{2^{O(k)}})$ *time.*

4.2 MoV Fair Matching

We now develop an FPT algorithm for MoV Fair Matching for the parameter k, which also works with the non-emptiness constraint. Our algorithm has two parts. In the first part, we give an FPT algorithm (using an ILP) for an auxiliary problem called Targeted MoV Fair Matching. This is a variant of MoV Fair Matching, where for each $v \in V$, the two most frequent colors appearing in $M(v)$ are given as part of the input. We establish the soundness of the ILP for Targeted MoV Fair Matching using again Theorem 6 and Lemma 2. In the second part, we present a (randomized) parameterized reduction from MoV Fair Matching to Targeted Mov Fair Matching using the color coding technique [5]. To apply this technique, we show that the colors that appear (second) most frequently in $M(v)$ for some $v \in V$ "stand out" in a fair matching that fulfills certain conditions. Then, the color coding technique essentially allows us to determine these colors.

Part I. First, we define an auxiliary problem, which we call Targeted MoV Fair Matching. The input for MoV Fair Matching is also part of the input for Targeted MoV Fair Matching. Moreover, Targeted MoV Fair Matching takes as input two functions $\mu^1, \mu^2\colon V \to C$. In Targeted MoV Fair Matching, we ask for an ℓ-fair matching M such that for every vertex $v \in V$, $\mu^1(v)$ (resp., $\mu^2(v)$) is the most (resp., second most) frequent color among the vertices $M(v)$ matched to v in M.

We now develop an FPT algorithm for Targeted MoV Fair Matching by means of an ILP. Let $C_{1,2} = \{\mu^1(v), \mu^2(v) \mid v \in V\}$ be the set of colors that appear (second) most frequent among the vertices matched to some vertex in V and let $C' = C \setminus C_{1,2}$ be the set of other colors. Notably, the size of $C_{1,2}$ is linearly bounded in our parameter k. For every $v \in V$, we introduce a variable y_v which represents the number of vertices of color $\mu^2(v)$ matched to v. The values of y_v need to be chosen in a way such that, for every color $c' \in C'$, there is a matching $M_{c'}$ in $G_{c'}$ such that $|M_{c'}(v)| \leq y_v$ for all $v \in V$. By Lemma 1, we obtain the following constraint which must be fulfilled and add it to the ILP:

$$\sum_{v \in W} y_v \geq \max_{c' \in C'} |\nu_{c'}(W)| \quad \forall W \subseteq V.$$

For the vertices of colors in $C_{1,2}$, we impose constraints in the same way as in Section 3. For every $c \in C_{1,2}$ and $v \in V$, we introduce a variable z_v^c which represents the number of vertices of color c matched to v. Then, we add the following constraints according to Lemma 1.

$$|\nu_c(W)| \leq \sum_{v \in W} z_v^c \leq |N_c(W)| \quad \forall W \subseteq V, c \in C_{1,2}.$$

In case we have a lower bound $p = 1$, we additionally require that $z_v^{\mu^1(v)} \geq 1$ for all $v \in V$. Finally, we encode the ℓ-fairness: $y_v = z_v^{\mu^2(v)}$ and $z_v^{\mu^2(v)} \leq z_v^{\mu^1(v)} \leq z_v^{\mu^2(v)} + \ell$ for all $v \in V$.

In order to show the correctness of the ILP, with the help of Theorem 6, we prove the following adaptation of Lemma 2, in which we show that there exists a matching of the vertices of colors from C' to vertices from V respecting y_v for all $v \in V$.

▶ **Lemma 10** ($\star$). *Let $G = (U \cup V, E)$ be a bipartite graph and let $\{z_v \in \mathbb{N} \mid v \in V\}$ be a set of integers. Suppose that $\sum_{v \in W} z_v \geq |\nu_G(W)|$ for every $W \subseteq V$. Then, there is a left-perfect many-to-one matching M such that $M(v) \leq z_v$ for every $v \in V$.*

Using Lemmas 2 and 10, we can now show that the above constructed ILP is feasible if and only if the given Targeted MoV Fair Matching is a yes-instance.

▶ **Proposition 11** ($\star$). *Targeted MoV Fair Matching can be solved in $O^\star(k^{O(k^2)})$ time even with the non-emptiness constraint.*

Part II. We will employ the color coding technique to reduce MoV Fair Matching to Targeted Mov Fair Matching. To do so, we introduce another auxiliary problem called $\mathcal{Q}$-MoV Fair Matching. To define the problem, we first introduce additional notation. Suppose that M is a matching in the input graph $G = (U \cup V, E)$. Let $\mu_M^1, \mu_M^2 : V \to C$ be mappings such that for every vertex $v \in V$, $\mu_M^1(v)$ (resp., $\mu_M^2(v)$) is the most (resp., second most) frequent color among the vertices $M(v)$ matched to v in M. When the maximum or second maximum is achieved by more than one color, we break ties according to a fixed linear order $\leq_C$ on C. Let $\mathcal{V} = \{v^1, v^2 \mid v \in V\}$ be a set containing $2|V|$ elements and let $\mathcal{P}(M)$ be a partition of $\mathcal{V}$ into subsets $S \subseteq \mathcal{V}$ where for every $S \in \mathcal{P}(M)$, $v^i, v'^j \in S$ if and only if $\mu_M^i(v) = \mu_M^j(v')$ for $v, v' \in V$ and $i, j \in [2]$.

Using this, we define $\mathcal{Q}$-MoV Fair Matching. Here, $\mathcal{Q}$ is a partition of $\mathcal{V}$ and we assume that $\mathcal{Q}$ is fixed. The input of $\mathcal{Q}$-MoV Fair Matching is identical to the input of MoV Fair Matching. The difference is that $\mathcal{Q}$-MoV Fair Matching asks for a left-perfect ℓ-fair many-to-one matching M in G *consistent* with $\mathcal{Q}$, that is, $\mathcal{P}(M) = \mathcal{Q}$. Clearly, an instance of MoV Fair Matching is a yes-instance if and only if there exists a partition $\mathcal{Q}$ of $\mathcal{V}$ such that the corresponding $\mathcal{Q}$-MoV Fair Matching instance is a yes-instance. Since there are at most $k^{O(k)}$ ways to partition $\mathcal{V}$ (which is a set of $2k$ elements), we can afford to "guess" $\mathcal{Q}$ in our FPT algorithm. We thus focus on solving $\mathcal{Q}$-MoV Fair Matching. In particular, we assume without loss of generality that $\mathcal{Q}$ is cardinality-wise maximum – the input graph G has no ℓ-fair left-perfect many-to-one matching M such that $|\mathcal{P}(M)| > |\mathcal{Q}|$ (this assumption becomes essential in the proof of Lemma 12).

We solve $\mathcal{Q}$-MoV Fair Matching using color coding. We focus on $\ell > 0$ and omit the case $\ell = 0$. To apply the color coding method, we show a rather technical lemma stating that if a color c is among the two most frequent colors in $M(v)$ for some $v \in V$ in a certain fair matching M, then c needs to "stand out" – the number of its occurrence in the neighborhood of v in G is greater than for any other color c', unless c' also "stands out".

▶ **Lemma 12.** *Let $\mathcal{I}$ be a yes-instance of $\mathcal{Q}$-MoV Fair Matching with $\ell > 0$. Then, $\mathcal{I}$ admits an ℓ-fair left-perfect many-to-one matching M such that, for every $v \in V$ and $i \in [2]$, $|N_{\mu^i_M(v)}(v)| \geq |N_{c'}(v)|$ for every $c' \in \{c \in C \mid \forall v \in V, i \in \{1,2\} \colon c \neq \mu^i_M(v)\}$.*

Proof. For a matching M, let $\sigma(M) := \sum_{v \in V, i \in \{1,2\}} |M(v)_{\mu^i_M(v)}|$ be the total number of the occurrences of the two most frequent colors in $M(v)$ over all vertices $v \in V$. Consider an ℓ-fair matching M that maximizes $\sigma(M)$ among all ℓ-fair matchings consistent with $\mathcal{Q}$. Let $C_M := \{\mu^i_M(v) \in C \mid v \in V, i \in \{1,2\}\}$ be the set of colors which are one of the two most frequent colors in $M(v)$ for some $v \in V$ and let $C'_M := C \setminus C_M$ be the set of other colors.

Assume for the sake of contradiction that there there is some $v \in V$ and $i \in [2]$ with $\mu^i_M(v) = c \in C_M$ and color $c' \in C'_M$ such that $|N_{c'}(v)| > |N_c(v)|$. Recall that by the definition of C'_M it holds that $|M(v)_c| \geq |M(v)_{c'}|$. Consider a left-perfect matching M' obtained from M as follows. Initially, let $M' := M$. We then repeat the following procedure $|M(v)_c| - |M(v)_{c'}| + 1$ times: We delete from M' an arbitrary edge $\{u, v'\} \in M'$ such that $u \in N(v) \subseteq U$ is a vertex of color c' and v' is some vertex in $V - v$, and then add an edge $\{u, v\}$. Note that this is always possible, as $|M(v)_{c'}| \leq |M(v)_c| \leq |N_c(v)| < |N_{c'}(v)|$. We claim that M' is a left-perfect ℓ-fair matching. It is easy to verify the left-perfectness of M'. For the ℓ-fairness, since $c' \in C'_M$, for every $v' \in V - v$, the number of occurrences of the two most frequent colors in $M'(v')$ has not changed compared to $M(v')$, i.e., $\max^1_{c'' \in C} |M'(v')_{c''}| = \max^1_{c'' \in C} |M(v')_{c''}|$ and $\max^2_{c'' \in C} |M'(v')_{c''}| = \max^2_{c'' \in C} |M(v')_{c''}|$. Thus, $M'(v')$ is ℓ-fair for each $v' \in V - v$. It thus remains to show that $\max^1_{c'' \in C} |M'(v)_{c''}| - \max^2_{c'' \in C} |M'(v)_{c''}| \leq \ell$. We consider three cases:

- If $c = \mu^1_M(v)$, then as we added $|M(v)_c| - |M(v)_{c'}| + 1$ vertices of color c' to $M(v)$, we have $\max^1_{c'' \in C} |M'(v)_{c''}| = |M'(v)_{c'}| = |M(v)_c| + 1$ and $\max^2_{c'' \in C} |M'(v)_{c''}| = |M(v)_c|$, and thus $\max^1_{c'' \in C} |M'(v)_{c''}| - \max^2_{c'' \in C} |M'(v)_{c''}| = 1 \leq \ell$.
- Suppose that $c = \mu^2_M(v)$ and $\max^1_{c \in C} |M(v)_c| > \max^2_{c \in C} |M(v)_c|$. Then, we have $\max^1_{c \in C} |M'(v)_c| = \max^1_{c \in C} |M(v)_c|$ and $\max^2_{c \in C} |M(v)_c| = \max^2_{c \in C} |M(v)_c| + 1$, and thus $\max^1_{c'' \in C} |M'(v)_{c''}| - \max^2_{c'' \in C} |M'(v)_{c''}| \leq \ell - 1 \leq \ell$.
- Suppose that $c = \mu^2_M(v)$ and $\max^1_{c \in C} |M(v)_c| = \max^2_{c \in C} |M(v)_c|$. Then, we have $\max^1_{c \in C} |M'(v)_c| = \max^1_{c \in C} |M(v)_c| + 1$ and $\max^2_{c \in C} |M(v)_c| = \max^2_{c \in C} |M(v)_c|$, and thus $\max^1_{c'' \in C} |M'(v)_{c''}| - \max^2_{c'' \in C} |M'(v)_{c''}| \leq 1 \leq \ell$.

This proves the ℓ-fairness of M.

Recall that $\mu^i_M(v) = c$. We now show that the existence of M' contradicts one of our assumptions. To that end, we consider the following two cases:

- Suppose that $\mu^j_M(v') = c$ for some $v' \in V \setminus \{v\}$ and $j \in [2]$. This case contradicts the fact that $\mathcal{Q}$ is cardinality-wise maximum: To see why, note that for the set S of $\mathcal{Q}$ containing v^i (which is of size at least two, since either v'^1 or v'^2 is in S), we have $\mathcal{P}(M') = (\mathcal{P}(M) \setminus \{S\}) \cup \{S - v, \{v\}\}$, implying that $|\mathcal{P}(M')| > |\mathcal{P}(M)|$.
- Suppose that $\mu^j_M(v') \neq c$ for each $v' \in V \setminus \{v\}$ and $j \in [2]$. Observe that $\{v^i\} \in \mathcal{Q}$. The matching M' is thus consistent with $\mathcal{Q}$. Moreover, we have $\sigma(M') > \sigma(M)$, which is a contradiction. ◀

With Lemma 12 at hand, we are ready to give a randomized reduction from $\mathcal{Q}$-MoV Fair Matching to Targeted MoV Fair Matching.

▶ **Lemma 13** ($\star$). *Let $\mathcal{I}$ be an instance of $\mathcal{Q}$-MoV Fair Matching. We can compute in polynomial time an instance $\mathcal{J}$ of Targeted MoV Fair Matching such that (i) $\mathcal{J}$ is a yes-instance with probability at least δ (where $1/\delta \in k^{O(k)}$) if $\mathcal{I}$ is a yes-instance and (ii) $\mathcal{J}$ is a no-instance if $\mathcal{I}$ is no-instance.*

Proof for $\ell > 0$. We describe a polynomial-time procedure to construct an instance $\mathcal{J}$ of Targeted MoV Fair Matching from a given instance $\mathcal{I} = (G = (U \cup V, E), C, \mathrm{col}, \ell)$ of $\mathcal{Q}$-MoV Fair Matching. We randomly assign each color $c \in C$ to one of the subsets in $\mathcal{Q}$ with the intended meaning that if we assign color $c \in C$ to $S \in \mathcal{Q}$, then in a matching M in $\mathcal{I}$ one of the following holds: (i) for $v^1 \in S$, c appears as the most frequent color in $M(v)$ and for $v^2 \in S$, c appears as the second most frequent color in $M(v)$ and c appears nowhere else as the most or second most frequent color or (ii) c does not appear as the most or second most frequent color in $M(v)$ for any $v \in V$. Formally let $\lambda\colon C \to \mathcal{Q}$ be a function assigning each color to a subset in $\mathcal{Q}$, where each assignment is chosen uniformly and independently at random.[6] For every $v \in V$ and $i \in [2]$ with $v^i \in S$ for some $S \in \mathcal{Q}$, we find $c_v^i = \arg\max |N_c(v)|$, where $\arg\max$ is taken over all colors c with $\lambda(c) = S$. Then, we construct an instance $\mathcal{J}$ of Targeted MoV Fair Matching, where $\mu^i(v) = c_v^i$ for every $v^i \in \mathcal{V}$. If a hypothetical matching M such that $\mu_M^i(u) = \mu^i(u) = c_v^i$ for $v \in V$ and $i \in [2]$ is not consistent with $\mathcal{Q}$, we let $\mathcal{J}$ be a trivial no-instance. Clearly, the construction of $\mathcal{J}$ takes polynomial time.

Suppose that $\mathcal{I}$ is a yes-instance. We show that in this case $\mathcal{J}$ is a yes-instance with probability at least δ with $\delta^{-1} \in k^{O(k)}$. Let M be an ℓ-fair matching for $\mathcal{I}$ that fulfills the properties described in Lemma 12 (which by Lemma 12 always exists). Assume that $\lambda(\mu_M^i(v)) = S_v^i$ holds for every $v \in V$ and $i \in [2]$, where $S_v^i \in \mathcal{Q}$ denotes the subset in $\mathcal{Q}$ to which v^i belongs. We claim that under this assumption on λ, the instance $\mathcal{J}$ constructed by our procedure is a yes-instance. In fact, we show that M is a solution of $\mathcal{J}$ (which also directly implies that there is a solution consistent with $\mathcal{Q}$ and thus that no trivial no-instance is returned). Since M is an left-perfect ℓ-fair matching in G, we only have to show that $\mu_M^i(v) = \mu^i(v) = c_v^i$ for every $v \in V$ and $i \in [2]$.

Let $C_M := \{\mu_M^i(v) \in C \mid v \in V, i \in [2]\}$ be the set of colors which are among the two most frequent colors in $M(v)$ for some $v \in V$ and let $C_M' := C \setminus C_M$ be the set of other colors. By Lemma 12, for every $v \in V$, $c' \in C_M'$, and $i \in [2]$, we have $|N_{\mu_M^i(v)}(v)| \geq |N_{c'}(v)|$. Since we always break ties according to a fixed linear order (including when we find $\mu_M^i(v)$), we have $\mu_M^i(v) = \arg\max_{c \in C_M' \cup \{\mu_M^i(v)\}} |N_c(v)|$. Recall that when constructing $\mathcal{J}$ we have defined $c_v^i = \arg\max |N_c(v)|$, where the maximum is taken over the set $C_v^i := \{c \in C \mid \lambda(c) = S_v^i\}$. By our assumption that $\lambda(\mu_M^i(v)) = S_v^i$ for every $v \in V$ and $i \in [2]$, we have $\mu_M^i(v) \in C_v^i$. We also have $c \notin C_v^i$ for every $c \in C_M \setminus \{\mu_M^i(v)\}$, which implies that $C_v^i \subseteq C_M' \cup \{\mu_M^i(v)\}$. Consequently, we obtain $c_v^i = \arg\max_{c \in C_v^i} |N_c(v)| = \mu_M^i(v)$ for every $v \in V$ and $i \in [2]$. It follows that M is a solution of $\mathcal{J}$.

Finally, observe that the probability that $\lambda(\mu_M^i(v)) = S_v^i$ for every $v \in V$ and $i \in [2]$ is at least $|\mathcal{Q}|^{-|\mathcal{Q}|} \geq (2k)^{-2k}$, since λ is chosen uniformly and independently at random. Thus, if $\mathcal{I}$ is a yes-instance, $\mathcal{J}$ is a yes-instance with probability at least $(2k)^{-2k}$.

If $\mathcal{J}$ is a yes-instance, then $\mathcal{I}$ is also a yes-instance, as, by construction of $\mathcal{J}$, a solution M for $\mathcal{J}$ needs to satisfy $\mathcal{P}(M) = \mathcal{Q}$ and is thus also a solution for $\mathcal{I}$. ◀

[6] In the language of color coding, the function λ is often described as assignments of "colors" (the "colors" in color coding are different from the colors used here).

We repeat the algorithm of Lemma 13 independently $\delta^{-1} \in k^{O(k)}$ times on the given instance $\mathcal{I}$ of $\mathcal{Q}$-Fair Matching. If $\mathcal{I}$ is a yes-instance, then at least one instance of Targeted Mov Fair Matching returned by the algorithm is a yes-instance with probability at least $1-(1-\delta)^{\delta^{-1}} \geq 1-1/e$. By Proposition 11, a yes-instance of Targeted Mov Fair Matching can be recognized in $O^\star(k^{O(k^2)})$ time. We thus have a randomized algorithm to solve $\mathcal{Q}$-MoV Fair Matching in $O^\star(k^{O(k^2)})$ time. Recall that $\mathcal{Q}$ is a partition of a $2k$-element set. So there are at most $k^{O(k)}$ choices for $\mathcal{Q}$, which gives us the following theorem (we remark that our algorithm can be derandomized using a standard method [12]):

▶ **Theorem 14.** *There is a randomized $O^\star(k^{O(k^2)})$-time algorithm to solve MoV Fair Matching even with the non-emptiness constraint.*

5 Complexity Dichotomies with respect to $|C|$ and Maximum Degree

In this section, we study the computational complexity of Max-Min/MoV Fair Matching for fixed values of Δ_U, Δ_V, and $|C|$. Recall that Δ_U (resp., Δ_V) is the maximum degree of all vertices in U (resp., V). We identify several computational dichotomies regarding these parameters (see Figure 1). We first show a dichotomy on $|C|$. In particular, Max-Min/MoV Fair Matching is polynomial-time solvable for $|C| = 2$ (even if there are arbitrary lower size constraints), while it is NP-hard for $|C| = 3$.

▶ **Theorem 15** ($\star$)**.** *Max-Min/MoV Fair Matching is polynomial-time solvable for $|C| = 2$ and NP-hard for $|C| \geq 3$.*

The next two theorems concern complexity dichotomies with respect to Δ_U and Δ_V. All NP-hardness results here hold for three colors, while all polynomial-time results hold for an arbitrary number of colors.

▶ **Theorem 16** ($\star$)**.** *MoV Fair Matching is polynomial-time solvable if $\Delta_U \leq 1$ or $\Delta_V \leq 4$ (even with the non-emptiness constraint) and NP-hard otherwise.*

As part of the proof of Theorem 16, we give an algorithm that solves MoV in polynomial-time for $\Delta_V \leq 4$. Our algorithm is a polynomial-time reduction to a polynomial-time solvable special case of General Factor, which we will define in the proof.

▶ **Proposition 17** ($\star$)**.** *MoV Fair Matching is polynomial-time solvable for $\Delta_V \leq 4$.*

Proof Sketch. We give a polynomial-time reduction to General Factor: In an instance of General Factor the input is an undirected graph $H = (W, F)$ and a degree list function $L \colon W \to 2^{\mathbb{N}}$ such that $L(w) \subseteq \{0, \ldots, \deg_G(w)\}$ for every vertex $w \in W$. The problem asks for a spanning subgraph $H' = (W, F')$ for $F' \subseteq F$ such that $\deg_{H'}(w) \in L(w)$ for every $w \in W$. By a result of Cornuéjols [11], General Factor is polynomial-time solvable if for every $w \in W$, the degree list $L(w)$ has gaps of size at most one, i.e., $\{\min L(w), \min L(w)+1 \ldots, \max L(w)\} \setminus L(w)$ does not contain any two consecutive integers. Given an instance $\mathcal{I} = (G = (U \cup V, E), C, \mathrm{col}, \ell)$ of MoV Fair Matching with $\Delta_V \leq 4$, we will construct an equivalent instance $\mathcal{J} = (H = (W, F), L)$ of this polynomial-time solvable special case of General Factor.

In the following, we give an overview of our construction. For every $v \in V$, we add a subgraph $H_v = (W_v, F_v)$ to H which contains all vertices from U adjacent to v in G ($N_G(v) \subseteq W_v$) but no other vertices from U. For the construction of H_v, we make extensive case distinctions depending on $N_G(v)$ and ℓ. See Figure 2 for two examples. The construction of H_v for all other cases are deferred to the appendix. The graph H is then the union of H_v

(a) H_v for $\ell = 0$ and $N(v)$ has two vertices of color α, one vertex of color β, and one vertex of color γ.

(b) H_v for $\ell = 1$ and $N(v)$ has two vertices of color α and two vertices of color β.

Figure 2 Exemplary constructions of H_v.

for all $v \in V$; notably, a vertex u from U may appear in multiple subgraphs H_v, in which case we identify all occurrences of u and merge them into one vertex. Moreover, for each $u \in U$, we set $L(u) = \{1\}$, which ensures that u is "matched" in every solution of $\mathcal{J}$.

To show that $\mathcal{I}$ and $\mathcal{J}$ are equivalent, it suffices to show the following for every $v \in V$:

A vertex set $S \subseteq N_G(v)$ is ℓ-fair if and only if there is a spanning subgraph $H'_v = (W_v, F'_v)$ of H_v such that $\deg_{H'_v}(u) = 1$ for every $u \in S$, $\deg_{H'_v}(u) = 0$ for every $u \in N(v) \setminus S$, and $\deg_{H'_v}(v') \in L(v')$ for every $v' \in W_v \setminus U$. $(\star)$

To see why $(\star)$ is sufficient, assume that $(\star)$ holds true. For the forward direction, suppose that $\mathcal{I}$ is a yes-instance, i.e., there is a ℓ-fair left-perfect many-to-one matching M in G. Then, $M(v) \subseteq N_G(v)$ is ℓ-fair for every $v \in V$. Hence, as we assume that $(\star)$ holds, we have a subgraph H'_v for every $v \in V$ that satisfies the degree constraints of $(\star)$. Consider a spanning subgraph H' whose edge set is the union of the edge set of H'_v over all v. It is easy to verify that H' constitutes a solution for $\mathcal{J}$. For the converse direction, suppose that $\mathcal{J}$ is a yes-instance, i.e., there is a spanning subgraph $H' = (W, F')$ with $\deg(w) \in L(w)$ for every $w \in W$. Then, the subgraph of H' induced by W_v satisfies all the degree constraints of $(\star)$. For $v \in V$, let $S_v \subseteq N_G(v)$ be the set of vertices that have a neighbor in $H'[W_v]$. By $(\star)$, S_v is ℓ-fair. Moreover, since $L(u) = \{1\}$ for every $u \in U$, every vertex appears in S_v for exactly one vertex $v \in V$. It follows that a matching M with $M(v) = S_v$ for every $v \in V$ is a ℓ-fair left-perfect many-to-one matching in $\mathcal{I}$. We remark that we can adapt our algorithm to handle the non-emptiness constraint. ◀

For Max-Min Fair Matching, we prove that fewer cases are polynomial-time solvable than for MoV Fair Matching:

▶ **Theorem 18** $(\star)$. *Max-Min Fair Matching is polynomial-time solvable if $\Delta_U \leq 1$, $\Delta_V \leq 2$, or $(\Delta_U, \Delta_V) = (2, 3)$ (even with the non-emptiness constraint) and NP-hard otherwise.*

6 Fair Matching on Complete Bipartite Graphs

A natural special case of Fair Matching is when the underlying graph is complete, i.e., each vertex in U can be assigned to any vertex in V. This special case is also among the three problems introduced by Stoica et al. [24] (they called it Fair Regrouping_X). Stoica et al. [24] presented a straightforward XP algorithm for MoV Fair Matching with size constraints parameterized by $|V|$ but left open the classical complexity. We partially settle this open question by proving that MoV Fair Matching on complete bipartite graphs is polynomial-time solvable even with the non-emptiness constraint. In fact, we find a precise characterization of yes-instances, which turns out to be surprisingly simple. However, it requires an intricate analysis to prove this, especially when the non-emptiness constraint is present. To simplify notation, we assume that $C = \{c_1, \ldots, c_{|C|}\}$ and that $|U_{c_i}| \geq |U_{c_{i+1}}|$ for each $i \in [1, |C| - 1]$ and set $|U_{c_i}| := 0$ for $i > |C|$.

▶ **Theorem 19** (⋆). *A* MoV Fair Matching *instance* $\mathcal{I} = (G = (U \cup V, E), C, \text{col}, \ell)$ *with* G *being a complete bipartite graph is a yes-instance if and only if* $|U_{c_1}| \leq \ell k + \sum_{i \in [k]} |U_{c_{i+1}}|$. *With the non-emptiness constraint,* $\mathcal{I}$ *is a yes-instance if and only if it additionally satisfies:* $\ell > 0$ *and* $n \geq k$*, or* $\ell = 0$ *and* $n \geq 2k$.

▶ **Theorem 20** (⋆). *A* Max-Min Fair Matching *instance* $\mathcal{I} = (G = (U \cup V, E), C, \text{col}, \ell)$ *with* G *being a complete bipartite graph is a yes-instance if and only if* $|U_{c_1}| \leq \ell k + |U_{c_{|C|}}|$. *With the non-emptiness constraint,* $\mathcal{I}$ *is a yes-instance if and only if it additionally satisfies:* $\ell > 0$ *and* $n \geq k$*, or* $\ell = 0$ *and* $|U_{c_1}| \geq k$.

Theorems 19 and 20 imply that Max-Min/MoV Fair Matching on a complete bipartite graph are solvable in linear time even with the non-emptiness constraint.

7 Conclusion

In this work, we have investigated the (parameterized) computational complexity of the Fair Matching problem. Two concrete directions of open questions are:

- We have provided algorithms that solve Fair Matching even if we require that every vertex in the right side is matched to at least one vertex. Can we extend our algorithms to handle arbitrary size constraints? In particular, does Fair Matching remain fixed-parameter tractable with respect to k? We have shown in Section 3 that Fair Matching is indeed FPT with respect to $k + |C|$ even for arbitrary size constraints. However, it does not seem straightforward to incorporate arbitrary size constraints in the ILPs given in Section 4. The complexity of Fair Matching on complete bipartite graphs (Section 6) is also open when arbitrary size constraints are present.
- Is Fair Matching solvable in $O^\star(2^k)$ time? Note that the ILP presented in Section 4.1 (which solves Max-Min Fair Matching without the non-emptiness constraint) is an ILP where the constraint matrix involves only zeros and ones when $\ell = 0$. Can we exploit such a structure in the constraint matrix to obtain a faster algorithm?

For future research, it would also be natural to study other variants of the Fair Matching problem. For instance, we may relax the left-perfect constraint studied in this work and consider a variant where the objective is to maximize the matching size under a fairness constraint. One may also look into other fairness notions such as proportionality constraints [21].

References

1 Saba Ahmadi, Faez Ahmed, John P. Dickerson, Mark Fuge, and Samir Khuller. An algorithm for multi-attribute diverse matching. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI '20)*, pages 3–9. ijcai.org, 2020.

2 Sara Ahmadian, Alessandro Epasto, Ravi Kumar, and Mohammad Mahdian. Fair correlation clustering. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS '20)*, pages 4195–4205. PMLR, 2020.

3 Faez Ahmed, John Dickerson, and Mark Fuge. Forming diverse teams from sequentially arriving people. *J. Mech. Design*, 142(11):111401, 2020.

4 Faez Ahmed, John P. Dickerson, and Mark Fuge. Diverse weighted bipartite b-matching. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI '17)*, pages 35–41. ijcai.org, 2017.

5 Noga Alon, Raphael Yuster, and Uri Zwick. Color-coding. *J. ACM*, 42(4):844–856, 1995. doi:10.1145/210332.210337.

6 Abolfazl Asudeh, Tanya Y. Berger-Wolf, Bhaskar DasGupta, and Anastasios Sidiropoulos. Maximizing coverage while ensuring fairness: a tale of conflicting objective. *CoRR*, abs/2007.08069, 2020. arXiv:2007.08069.

7 Haris Aziz, Serge Gaspers, Zhaohong Sun, and Toby Walsh. From matching with diversity constraints to matching with regional quotas. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*, pages 377–385. IFAAMAS, 2019.

8 Inacio Bo. Fair implementation of diversity in school choice. *Games Econ. Behav.*, 97:54–63, 2016.

9 Niclas Boehmer, Tomohiro Koana, and Rolf Niedermeier. A refined complexity analysis of fair districting over graphs. *CoRR*, abs/2102.11864, 2021. Accepted as an extended abstract at AAMAS '22. arXiv:2102.11864.

10 Jiehua Chen, Robert Ganian, and Thekla Hamm. Stable matchings with diversity constraints: Affirmative action is beyond NP. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI '20)*, pages 146–152. ijcai.org, 2020.

11 Gérard Cornuéjols. General factors of graphs. *J. Comb. Theory, Ser. B*, 45(2):185–198, 1988. doi:10.1016/0095-8956(88)90068-8.

12 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms.* Springer, 2015. doi:10.1007/978-3-319-21275-3.

13 John P. Dickerson, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. Balancing relevance and diversity in online bipartite matching via submodularity. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI '19)*, pages 1877–1884. AAAI Press, 2019.

14 András Frank. An algorithm for submodular functions on graphs. *Discrete Math.*, 16:97–120, 1982.

15 Zachary Friggstad and Ramin Mousavi. Fair correlation clustering with global and local guarantees. In *Proceedings of the 17th International Symposium on Algorithms and Data Structures (WADS '21)*, pages 414–427. Springer, 2021.

16 Isa E Hafalir, M Bumin Yenmez, and Muhammed A Yildirim. Effective affirmative action in school choice. *Theor. Econ.*, 8(2):325–363, 2013.

17 Philip Hall. On representatives of subsets. *Classic Papers in Combinatorics*, pages 58–62, 1987.

18 Chien-Chung Huang. Classified stable matching. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '10)*, pages 1235–1253. SIAM, 2010.

19 Ravi Kannan. Minkowski's convex body theorem and integer programming. *Math. Oper. Res.*, 12(3):415–440, 1987.

20 Hendrik W. Lenstra. Integer programming with a fixed number of variables. *Math. Oper. Res.*, 8:538–548, 1983.

21 Thành Nguyen and Rakesh Vohra. Stable matching with proportionality constraints. *Oper. Res.*, 67(6):1503–1519, 2019.

22 Deval Patel, Arindam Khan, and Anand Louis. Group fairness for knapsack problems. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '21)*, pages 1001–1009. ACM, 2021.

23 Dana Pessach and Erez Shmueli. Algorithmic fairness. *CoRR*, abs/2001.09784, 2020. arXiv:2001.09784.

24 Ana-Andreea Stoica, Abhijnan Chakraborty, Palash Dey, and Krishna P. Gummadi. Minimizing margin of victory for fair political and educational districting. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '20)*, pages 1305–1313. IFAAMAS, 2020.