Group Fairness in Set Packing Problems

Sharmila Duppala, **Juan Luque**, **John Dickerson** and **Aravind Srinivasan**
University of Maryland, College Park
{sduppala, jluque, johnd, asriniv1}@umd.edu

Abstract

Kidney exchange programs (KEPs) typically seek to match incompatible patient-donor pairs based on a utilitarian objective where the number or overall quality of transplants is maximized—implicitly penalizing certain classes of difficult to match (e.g., highly-sensitized) patients. Prioritizing the welfare of highly-sensitized (hard-to-match) patients has been studied as a natural *fairness* criterion. We formulate the KEP problem as k-set packing with a probabilistic group fairness notion of proportionality fairness—namely, fair k-set packing (FAIRSP). In this work we propose algorithms that take arbitrary proportionality vectors (i.e., policy-informed demands of how to prioritize different groups) and return a probabilistically fair solution with provable guarantees. Our main contributions are randomized algorithms as well as hardness results for FAIRSP variants. Additionally, the tools we introduce serve to audit the price of fairness involved in prioritizing different groups in realistic KEPs and other k-set packing applications. We conclude with experiments on synthetic and realistic kidney exchange FAIRSP instances.

1 Introduction

Fielded kidney exchange programs (KEPs) often employ a utilitarian objective, where the overall utility of successful transplants is maximized. This approach, however, can penalize certain classes of hard-to-match patients (e.g., highly-sensitized, older, O-type blood, and others), as highlighted by many studies from the AI/ML [Dickerson *et al.*, 2014; Li *et al.*, 2014; Farnadi *et al.*, 2021; Sun *et al.*, 2021], economics [Roth *et al.*, 2005; Ashlagi *et al.*, 2019], and medical communities [Ashlagi *et al.*, 2011]. Subsequent research by [McElfresh and Dickerson, 2018] explored two group-fairness criteria—lexicographical fairness and weighted fairness—under random-graph models, and determined that the tradeoff between efficiency and fairness, known as the "price of fairness," is relatively small. However, it should be noted that these guarantees apply only under the assumption of stochastic models in kidney exchange and there is currently a lack of fair algorithms with provable guarantees for the general problem—*directed cycle packing*—as originally defined [Abraham *et al.*, 2007; Biro *et al.*, 2009]. Several works studied approximation algorithms for KEPs under this model [Biro *et al.*, 2009; Lin *et al.*, 2019; Xiao and Wang, 2018] but none of those algorithms consider fairness; therefore the problem of group fairness in kidney exchange on arbitrary patient-donor graphs presents a unique and intriguing challenge, which can be approached from both theoretical computer science and economics perspectives.

KEPs are often formulated as cycle-packing problems with bounded length cycles, say cycle length at most k. This, in turn, can be reduced to k-set packing[1] by naively enumerating the cycles in $\mathcal{O}(n^k)$ worst-case time [Blum *et al.*, 2015]; each cycle, which has l (less than or equal to k) vertices, represents a subset of l elements in the k-set packing formulation.

For the rest of the paper, we lean into kidney exchange to illustrate the need for fairness in the well-known k-set packing problem, though other suitable applications exist such as crew scheduling and barter exchanges. We conduct experiments over a realistic kidney-exchange dataset.

We study here the notion of group-fairness in k-set packing problems. In our model, each of the n elements in the universe $\mathcal{U}$ is assigned a color (category), and our goal is to compute a valid packing such that (i) each group is *fairly* represented satisfying exogenous proportionality constraints and (ii), subject to (i), the overall weight of the packing is maximized. On the flip side, our work is not solely meant to correct so-called unfair algorithms. Rather, the tools we develop also serve as a basis to understand and audit the group-level behavior of algorithms. That is to say, the lack of studies of group-level fairness in k-set packing is representative of missing, but necessary, work towards understanding group outcomes in kidney exchange, equitable crew scheduling, and k-set packing's other applications. Thus, we view our work as an important step towards understanding group-level fairness in kidney exchange.

Let $[t]$ denote the set $\{1, 2, \ldots, t\}$. In *fair* k-set packing, denoted FAIRSP, we are given a set of elements $\mathcal{U} = [n]$ and a collection $\mathcal{S} = \{S_1, \ldots, S_m\}$ of subsets of $\mathcal{U}$ where

[1]The k-set packing problem is set packing over a universe $\mathcal{U}$ and collection of subsets $\mathcal{S}$ of $\mathcal{U}$ restricted to each $S \in \mathcal{S}$ having at most k elements. For $k \geq 3$, it is NP-hard [Karp, 1972].

$|S_i| \leq k$ for each S_i. Each subset $S_i \in \mathcal{S}$ has weight $w_i \geq 0$. Moreover, each $u \in \mathcal{U}$ is assigned one of c colors by $\mathcal{C} : \mathcal{U} \rightarrow [c]$; i.e., each u belongs to one of c groups. We are also given a proportionality vector $\vec{p} \in [0,1]^c$ such that $\sum_{\ell \in [c]} p_\ell = 1$. The objective of $\overline{\text{FAIRSP}}$ is to select a packing[2] of maximum weight subject to a p_ℓ fraction of the packing consisting of elements with color ℓ: i.e., letting $y_i = 1$ if S_i is selected and $y_i = 0$ otherwise, we aim to maximize $\sum_{i \in [m]} w_i y_i$ subject to the proportionality constraint.

However, many instances become immediately infeasible under these additional constraints. For example, let us consider an instance with $\mathcal{U} = \{r, b_1, b_2, g_1, g_2\}$, $S_1 = \{r, b_1, b_2\}$, $S_2 = \{r, g_1, g_2\}$, and $\vec{p} = [1/3, 1/3, 1/3]$; here r, b_i's, and g_i's are colored red, blue, and green, respectively. Clearly, there does not exist a feasible (non-empty) fair packing. Instead, the fairness constraints can be satisfied in expectation by switching our target to a *distribution* over packings (say, pick each set with probability $1/2$).

Observation 1. *The integrality gap of the LP corresponding to $\overline{\text{FAIRSP}}$ is unbounded.*

Therefore, we shift our attention to randomized notions of fairness. Such probabilistic fairness guarantees have been previously studied by [Asudeh *et al.*, 2022; Bastani *et al.*, 2019]. This motivates our work's exploration of randomized algorithms and probabilistic fairness.

Any randomized packing algorithm ALG is captured by a distribution $\mathcal{D}$ over the set of all packings; i.e., running ALG is akin to sampling from the distribution $\mathcal{D}$. Let $Z_i = 1$, $i \in [m]$, indicate that S_i is selected in the packing returned by ALG ($Z_i = 0$ otherwise). Letting $v_{i\ell}$ be the number of elements of color ℓ in S_i, we have $N_\ell := \sum_{i \in [m]} v_{i\ell} Z_i$ and $N := \sum_{\ell \in [c]} N_\ell$. We focus our study on FAIRSP, defined as follows: Given $\mathcal{U}$, $\mathcal{S}$, $w_i \geq 0$, $\mathcal{C} : \mathcal{U} \rightarrow [c]$ (as defined in $\overline{\text{FAIRSP}}$), FAIRSP seeks a distribution $\mathcal{D}$ such that $\mathbb{E}_{Z \sim \mathcal{D}}[\sum_{i \in [m]} w_i Z_i]$ is maximized while satisfying the proportional constraints in *expectation* i.e., for any color ℓ, $\mathbb{E}[N_\ell] = p_\ell \mathbb{E}[N]$ holds. We require $\mathbb{E}[N] > 0$ to rule out the distribution with support only over the empty set.

In this paper, we study randomized algorithms for FAIRSP, with a focus on two specific randomized fairness notions outlined in Section 3. We drop the word *probabilistic* while referring to the proportionality constraints of FAIRSP for brevity.

2 Related Work

k-set packing is a well-studied problem in combinatorial optimization. For $k = 2$ the problem is *maximum-weight matching*, which can be solved in polynomial time, and for $k = 3$ the problem becomes APX-hard [Hazan *et al.*, 2006]. To the best of our knowledge, the problem closest to our formulation is k-set packing. In the past, the best algorithms for both the weighted and unweighted variants depend on local-search techniques with an approximation ratio of $(\frac{k+1}{2} + \epsilon)$ [Berman, 2000] and $(\frac{k+1}{3} + \epsilon)$ [Fürer and Yu, 2014] respectively. However, recently, [Neuwohner, 2021] improved upon the weighted variant, achieving an approximation ratio of $\frac{k+\epsilon_k}{2}$[3]. KEPs were formulated as instances of k-set packing in [Biro *et al.*, 2009] for the study of approximation algorithms. Specifically, [Biro *et al.*, 2009] developed a $k - 1 + \epsilon$ approximation algorithm based on the local search technique utilizing augmenting paths.

It is known that the natural LP-relaxation for k-set packing has an integrality gap at least $k - 1 + \frac{1}{k}$ [Füredi *et al.*, 1993]. [Brubach *et al.*, 2019] gives a $k + \epsilon_k$–approximate algorithm based on LP rounding that almost matches this lower bound. [Anegg *et al.*, 2021] simplified the analysis of [Brubach *et al.*, 2019] and improved the approximation factor to $k - \epsilon_k$. The problem is $\Omega(\frac{k}{\log k})$-hard to approximate [Hazan *et al.*, 2006] (i.e., even an efficient $O(\frac{k}{\log k})$–approximation seems unlikely).

To our knowledge, *fair* k-set packing has not been studied before for a general k, although the special case of $k = 2$ (i.e., maximum-weight matching) has been explored extensively under both group—[Sankar *et al.*, 2021; Ma *et al.*, 2022; Nanda *et al.*, 2020] and individual– [García-Soriano and Bonchi, 2020] fairness. Fair Matching has been studied under various group-fairness notions such as lexicographical fairness [García-Soriano and Bonchi, 2020] and Rawlsian (max-min) group fairness [Ma *et al.*, 2022; Esmaeili *et al.*, 2023]; both have been studied in online and offline settings.

A great deal of work in fair algorithm design uses the framework of modeling the problem as an *optimization problem* with precisely defined *fairness constraints* and vice-versa [Sankar *et al.*, 2021]. Group-level fairness has been studied in various settings [Diana *et al.*, 2021; Tsang *et al.*, 2019; Ma *et al.*, 2022]. [Asudeh *et al.*, 2022] considered group-level proportionality fairness (both deterministic and randomized) for the *fair maximum coverage* problem which is closely related to the set packing problem.

3 Preliminaries and Main Contributions

Throughout this paper, we denote $[n] = \{1, \ldots, n\}$ for any positive integer n; we use OPT to denote both an optimal algorithm and its objective value; and analogously ALG for both a generic algorithm and its objective value. We use ϵ_n to denote a vanishing term i.e., $\lim_{n \to 0} \epsilon_n = 0$. We use either $\{Y_i\}$ or $\{Y_i\}_{i \in [m]}$ to refer to the set of variables $\{Y_1, Y_2, \ldots, Y_m\}$ interchangeably.

Approximation Factor. For NP-hard combinatorial optimization problems, we utilize the powerful framework of approximation algorithms, where the aim is to design efficient algorithms (polynomial running time) with provably near-optimal objectives values.

We say ALG achieves an approximation factor of $\alpha \geq 1$ if $\alpha \mathbb{E}[\text{ALG}] \geq \text{OPT}$ across all feasible FAIRSP instances.

Parameters. The complexity of FAIRSP can depend on the following parameters; number of colors c; maximum element frequency $f \in [m]$, where the frequency of an element is the number of sets it belongs to; and $k \in [n]$. The k-restrained variant is *necessary* in applications such as crew scheduling

[2]A packing is a collection of pairwise disjoint subsets of $\mathcal{S}$.

[3]We use ϵ_n to denote a vanishing term i.e., $\lim_{n \to 0} \epsilon_n = 0$.

and kidney exchange, as motivated by prior work [Biró *et al.*, 2019]. Moreover, the natural LP for k-set packing (LP (3) without the constraints (3c)) has integrality gap of at least $k-1+\frac{1}{k}$ and there exists a $(k-\epsilon_k)$-approximate algorithm by [Anegg *et al.*, 2021]. Note our algorithmic results are parameterized on k, c and f.

We study algorithms with approximation guarantees on the packing objective and fairness of FAIRSP. We make this precise via two related fairness guarantees for randomized algorithms of FAIRSP.

Randomized approximate Fairness (RF). For any two colors $\ell, q \in [c]$,

$$\mathbb{E}[N_\ell] / \mathbb{E}[N_q] \le \gamma\Big(\frac{p_\ell}{p_q}\Big). \tag{1}$$

The guarantee above is fair only in expectation. It is possible for a single realized packing to be very unfair itself. We address this next via a stronger notion of randomized proportional fairness.

Strongly Randomized approximate Fairness (SRF).

$$\Pr\Big[\bigwedge_{\ell,q\in[c]} \Big(\frac{N_\ell}{N_q} \le \gamma\Big(\frac{p_\ell}{p_q}\Big)\Big)\Big] \ge 0.99 \tag{2}$$

If a randomized algorithm ALG, satisfies the inequalities (1) and (2), then it is said to have a fairness factor γ with respect to RF and SRF respectively. RF and SRF, respectively, are ex-ante and ex-post fairness guarantees for FAIRSP algorithms with respect to the proportionality vector and underlying groups (colors). In Section 4 we give algorithms for FAIRSP with varied approximation guarantees on the packing objective and RF or SRF guarantees with *fairness factor* γ.

Connections to existing works. For $k \ge 3$, the problem of k-set packing is well-studied, and several previous works (such as [Bansal *et al.*, 2010; Brubach *et al.*, 2019; Anegg *et al.*, 2021]) have focused on developing algorithms based on the natural LP relaxation of k-set packing. The primary motivation for using LP rounding techniques is the ease of incorporating proportionality constraints as *additional* constraints in the LP, as demonstrated in LP (3). Our main *technical challenge* is to adapt these rounding algorithms in such a way that they produce packings with a small approximation factor and fairness ratio γ simultaneously.

Choosing a proportionality vector. Previous studies such as [Chierichetti *et al.*, 2017; Bei *et al.*, 2022] have examined a concept of proportional fairness that is broader than the one used in our work. They apply upper and lower bound constraints on the desired proportional representation of each group in the solution. In contrast, we use a special case where the upper and lower bounds coincide, meaning we insist on the final packing solution exactly preserving the proportional constraints.

Our algorithms take an input parameter $\vec{p}$ and output a corresponding fair solution, assuming that the given instance is feasible. Determining an appropriate $\vec{p}$ is intentionally deferred to stakeholders and policymakers. The question of how to fairly allocate resources to groups is an ongoing topic of

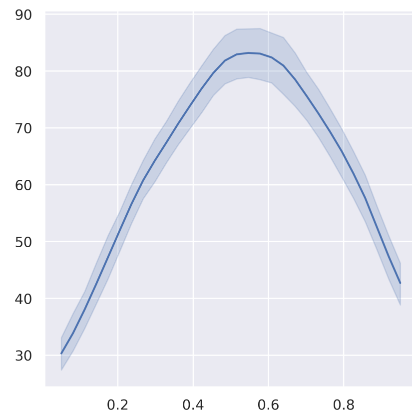

Figure 1: $\vec{p} = [p_1, p_2]$ vs. packing objective achieved by Algorithm 1. Full description in Section 5.

debate which we make no attempt to conclude (e.g. see [Roth, 2015]). Nevertheless, we view our work as progress toward formalizing and facilitating discussion of group-level fairness in kidney exchange [McElfresh and Dickerson, 2018] and other k-set packing applications. For example, the $\vec{p}$ corresponding to an optimal packing has optimal Price of Fairness (PoF); i.e., there is no cost to the objective. On the other hand, other choices of $\vec{p}$ may come at a great PoF. For a policymaker to advocate for proportional group fairness, it is imperative they be able to articulate this trade-off. Our work is a step towards formalizing these necessary tools. Indeed, we explicitly view our work as descriptive, not prescriptive—we can map out a form of Pareto frontier (as in Figure 1) balancing the "tightness" imposed by the proportionality vector (x-axis) against the overall utility of the solution (y-axis). Then, a domain expert could use this for decision support when choosing the proper balance.

3.1 Main Contributions

We study the problem of FAIRSP, a probabilistic fair variant of the k-set packing problem. We motivate the probabilistic fairness notions (RF and SRF as defined in (1) and (2) resp.) by first showing that the deterministic fair variant $\overline{\text{FAIRSP}}$ has unbounded integrality gap for the natural LP relaxation of the problem.

We provide two algorithms FAIRSAMPLE and FAIRELIMINATE that respectively guarantee RF (with $\gamma \approx 1$ w.h.p. and approximation factor $\alpha \le k + \epsilon_n$) and SRF (with $\gamma = \mathcal{O}(1)$ and $\alpha = \mathcal{O}(ck^2)$ where c is the number of colors). We say *sufficiently large packing size* to mean the fractional number of elements selected in the optimal LP (3) solution, i.e., $\sum_{i\in[m]}\sum_{\ell\in[c]} v_{i\ell} y_i^*$, is sufficiently large.

Theorem 1. *For any instance of* FAIRSP, FAIRSAMPLE *is a randomized, polynomial-time algorithm with an approximation factor* $(k+\epsilon_k)$ *on the packing objective. Moreover, with probability* $1-\epsilon_L$, FAIRSAMPLE *guarantees RF with* $\gamma = 1+\epsilon_L$ *where L is a tunable parameter.*

Theorem 2. *For any instance of* FAIRSP *satisfying* $f = o(\sqrt{n})$ *and* $p_\ell = \Theta(1/c)$ *for all* $\ell \in [c]$, FAIRELIMINATE *is a randomized, polynomial-time algorithm with an approximation factor $\mathcal{O}(ck^2)$ on the packing objective. Moreover, for instances with sufficiently large packing size,* FAIRELIMINATE *guarantees SRF with $\gamma = \mathcal{O}(1)$.*

Theorem 3. *For any instance of* FAIRSP *satisfying $f = o(n/k^6c^4)$ and $p_\ell = \Theta(1/c)$ for all $\ell \in [c]$,* FAIRELIMINATE *is a randomized, polynomial-time algorithm with an approximation factor of $\mathcal{O}(ck^2)$ on the packing objective. Moreover, for instances with sufficiently large packing size,* FAIRELIMINATE *guarantees SRF with $\gamma = \mathcal{O}(1)$.*

Remark 1. *Theorem 3 is a stronger result than Theorem 1 since the former allows for instances with f (maximum frequency of any element) asymptotically larger than that of the latter.*

Note that stronger fairness guarantees of FAIRELIMINATE holds for the family of FAIRSP instances satisfying: (i) the frequency of any element is bounded by $n^{1-o(1)}$, e.g., $f = o(n)$ for $k = \mathcal{O}(\log n)$ and $c = \mathcal{O}(\log n)$, (ii) the proportionality vector satisfies $p_\ell = \Theta(1/c)$ for all ℓ, and (iii) the packing size is sufficiently large.

The input parameter a to FAIRELIMINATE is carefully chosen to balance both γ and α. For the family of instances with constant number of colors c, and sufficiently large packing size we prove that $\gamma = \mathcal{O}(1)$ and $\alpha = \mathcal{O}(k^2)$. For many practical real-world applications, like kidney-exchange markets, k is a very small number, say 2 or 3, thus resulting in a reasonable (constant) approximation factor for both FAIRSAMPLE and FAIRELIMINATE, along with the desired fairness guarantees.

Notice that k-set packing is a special case of FAIRSP with $c = 1$, hence the integrality gap of LP (3) is at least $k-1+1/k$ [Füredi *et al.*, 1993]. Therefore, our approximation factor for FAIRSAMPLE almost matches the lower bound and further guarantees RF with γ arbitrarily close to optimal i.e., 1.

Whereas, FAIRSAMPLE only guarantees fairness in expectation—with no promises on the ex-post fairness of the allocation—FAIRELIMINATE guarantees the stronger SRF with $\gamma = \mathcal{O}(1)$ and approximation factor $\alpha = \mathcal{O}(ck^2)$.

Finally, in Section 5, we support these theoretical results with experiments on synthetic datasets, including a realistic dataset drawn from a real-world kidney exchange.

4 Randomized Algorithms

In this section we focus on randomized algorithms for FAIRSP with RF and SRF guarantees. Our algorithms are based on rounding optimal Linear Program (LP) solutions. The LP formulation for FAIRSP, found in LP (3), is the standard LP formulation of set packing with added fairness constraints. Due to space constraints, we omit the proofs from this section.

$$\max \quad \sum_{i\in[m]} w_i y_i \tag{3a}$$

$$\text{subject to} \quad \sum_{i:u_j\in S_i} y_i \leq 1, \qquad j \in [n] \tag{3b}$$

$$\sum_{i\in[m]} v_{i\ell} y_i = p_\ell \sum_{t\in[c]} \sum_{i\in[m]} v_{it} y_i, \qquad \ell \in [c] \tag{3c}$$

$$y_i \in [0,1], \qquad i \in [m]. \tag{3d}$$

Theorem 4. *The optimal value of LP* (3) *is an upper bound on the objective of* FAIRSP.

Proof. Consider an optimal randomized algorithm OPT[4] for FAIRSP. For each $i \in [m]$, let y_i be the probability that S_i is packed by the randomized algorithm OPT. We can verify that the LP objective (3a) captures the exact value of OPT (i.e., the max expected weight of the packing). Therefore, to prove our claim it suffices to show that $\{y_i\}_{i\in[m]}$ is a feasible point of LP (3). For each $i \in [m]$, let Y_i denote the indicator random variable that set S_i is selected by OPT. OPT can be viewed as a distribution over all feasible deterministic algorithms, any realization $\{Y_i\}$ satifsies $\sum_{j\in S_i} Y_i \leq 1$ for each $j \in [n]$. Therefore, $\mathbb{E}[\sum_{j\in S_i} Y_i] \leq 1$ which satisfies (3b). Since OPT is optimal for FAIRSP, for each color $\ell \in [c]$, $\{Y_i\}$ must satisfy proportional constraints in *expectation* i.e., $\mathbb{E}[N_\ell] = p_\ell \mathbb{E}[N]$, so (3c) is satisfied. Lastly, since $\{y_i\}$ are probabilities, (3d) is satisfied. □

Lemma 1. *An instance of* FAIRSP *is feasible if and only if its corresponding LP* (3) *has a non-zero feasible solution.*

By Lemma 1 we can efficiently verify whether the underlying FAIRSP instance is feasible. Moreover, this implies our algorithms abort *only when* the underlying FAIRSP instance is infeasible.

Algorithmic challenges and techniques. As highlighted in the introduction, randomized algorithms can be substantially more powerful than deterministic ones on the objective studied here; however, the design and analysis of a randomized algorithm can be technically challenging mainly due to the craft involved in using randomness and analysis of such random variables. Several existing dependent rounding methods [Bansal *et al.*, 2010; Brubach *et al.*, 2019; Anegg *et al.*, 2021] for k-set packing are based on *randomized rounding* (which samples subsets based on an LP solution) followed by *suitable alterations* (which drops some sets to ensure a packing is returned) that lead to a feasible packing with good approximation guarantees. However, the packing returned by these algorithms is far from satisfying any fairness constraints. Thus the main challenge we address is designing packing algorithms guaranteeing RF and SRF. Our primary contribution is the two algorithms, FAIRSAMPLE and FAIRELIMINATE, which are tailored to guarantee RF and SRF, respectively. Our algorithms utilize a randomized dependent rounding method, drawing inspiration from the works of [Bansal *et al.*, 2010] and [Brubach *et al.*, 2019], to optimize the packing objective function. Additionally, we modify existing algorithms to further provide probabilistic fairness guarantees (minimizing γ). In rest of the paper we use simulation to refer to Monte Carlo simulation.

1. We augment the rounding method of [Brubach *et al.*, 2019] with a simulation based method to tightly bound the probability of selecting any subset S_i, $i \in [m]$. Thus, guaranteeing RF with near-optimal γ (i.e., $\gamma = 1 + \epsilon_L$[5]). (See Theorem 1)

[4]We use OPT for both the optimal algorithm and the optimal objective value interchangeably.

[5]L is the number of simulations in FAIRSAMPLE.

2. Although FAIRSAMPLE provides a better approximation on the packing objective than FAIRELIMINATE, it is hard to provide a useful lowertail bound of N_ℓ since the random variables $\{Z_i\}_{i\in[m]}$ are highly dependent. We use a much more intricate analysis on the simpler FAIRELIMINATE using concentration bounds from [Janson and Ruciński, 2002] which guarantee SRF with $\gamma = \mathcal{O}(1)$ (a small constant), under reasonable assumptions on the input instances (See Theorems 2 and 3).

4.1 FAIRSAMPLE for RF Guarantees

We obtain FAIRSAMPLE, Algorithm 1, by augmenting the rounding algorithm from [Brubach *et al.*, 2019] with an additional simulation step (step 13). The purpose behind the simulation-based sampling is to "tightly" bound $\mathbb{E}[N_\ell]$ for each color ℓ; the importance of this step is discussed in the "Intuition behind the analysis of FAIRSAMPLE". The probability of satisfying RF with $\gamma = 1+\epsilon_L$ depends on the number of simulations L. Nevertheless, L needs only be polynomial in the problem size for high-quality results. For more details on the simulation-based sampling method see the works of [Ma, 2014; Adamczyk *et al.*, 2015].

Steps 1-4 FAIRSAMPLE solves the LP of the given FAIRSP instance or aborts if the LP is infeasible. **Steps 7-10** Next, it runs randomized rounding on each variable using a carefully chosen function $f(y_i^*)$ on the optimal LP solution $\{y_i^*\}$. In the alteration step, the variables rounded to 1 are dropped based on a randomly ordered permutation generated in the previous step. These steps are repeated for L rounds. **Steps 13- 15** Finally, with the previous L simulations, we have estimated $q_i' \approx \Pr(S_i \in \mathcal{F}_2)$ for each $i \in [m]$. This ensures that $\Pr(S_i \in \mathcal{F}) \approx \frac{y_i^*}{k+\frac{k}{\exp k}}$, which is crucial to the analysis. Theorem 1 formalizes the approximation guarantees provided by FAIRSAMPLE for the objective and RF fairness ratio.

Intuition behind the analysis of FAIRSAMPLE. The first part of FAIRSAMPLE, steps 1-11, is similar to that of [Brubach *et al.*, 2019], thus leading to $k + \epsilon_k$ approximation on packing objective. The remaining steps of FAIRSAMPLE ensure a tight upper bound on $\mathbb{E}[N_\ell]$ for each color ℓ. This is crucial to upper bounding $\mathbb{E}[N_\ell]/\mathbb{E}[N_{\ell'}]$ for each pair of colors ℓ and ℓ', thus giving a suitable γ.

Theorem 1 demonstrates that FAIRSAMPLE has an approximation factor of $(k + \frac{k}{\exp(k)})$ for the packing objective, and it guarantees RF with a $\gamma = 1 + \epsilon_L$, which can be made arbitrarily close to optimal ($\gamma = 1$) with more simulations. However, it should be noted that the fairness guarantees of FAIRSAMPLE only hold with high probability, as they depend on the simulation step. Nevertheless, the experimental results in Section 5 show FAIRSAMPLE consistently achieves $\gamma \approx 1$ even with only a few simulations.

Note that Theorem 1 only guarantees fairness in expectation—there are no guarantees on the ex-post fairness of the allocation. This shifts our attention to FAIRELIMINATE.

Algorithm 1 FAIRSAMPLE

Require: Number of Monte Carlo simulations L.
1: Solve LP (3) to get an optimal fractional solution $\{y_i^*\}_{i\in[m]}$.
2: **if** LP (3) is infeasible **then**
3:     **return** "Infeasible"
4: **end if**
5: For all $i \in [m]$, set $L_i = 0$.
6: **for** $j = 1, \ldots, L$ **do**
7:     Construct $\mathcal{F}_1$ by sampling each set S_i with probability $f(y_i^*)$ where $f(x) = x\big(1 - \frac{x}{2}\big)$.
8:     For each $S_i \in \mathcal{F}_1$, sample $x_i \sim U(0, 1)$.
9:     Build the valid packing $\mathcal{F}_2$ as follows. In increasing order of x_i, pack S_i if it does not intersect a previously selected set.
10:     For each $S_i \in \mathcal{F}_2$, increment L_i by 1.
11: **end for**
12: Keep the packing $\mathcal{F}_2$ from the last run of the above loop.
13: Set $q_i' = L_i/L$ for each $i \in [m]$.
14: Shrink $\mathcal{F}_2$ into $\mathcal{F}$ as follows. For each $S_i \in \mathcal{F}_2$, keep S_i with probability $\frac{y_i^*}{(q_i'+\epsilon)\left(k+\frac{k}{\exp(k)}\right)}$.
15: **return** The packing $\mathcal{F}$.

4.2 FAIRELIMINATE for SRF Guarantees

FAIRELIMINATE, Algorithm 2, incorporates *randomized rounding* followed by *alterations*, similar to [Bansal *et al.*, 2010]. In the *randomized rounding* step, we select each S_i independently based on a attenuation function $\frac{ay_i^*}{k}$ where a is a small constant. The main difference between FAIRELIMINATE and existing algorithms is allowing for an appropriate parameter a that results in a collection of sets with acceptable "expected packing weight" and "expected cardinality[6]". The *alteration* step ensures a packing is returned by discarding intersecting sets while maintaining good expected weight and not losing too many elements of any color.

To guarantee SRF, we need concentration bounds on the upper and lower tails of N_ℓ. Our randomized rounding procedure lets us upper bound N_ℓ as a sum of independent random variables, making it easy to apply Hoeffding bounds on the upper tail. However, the critical part of the analysis is establishing a *concentration on the lowerbound* of N_ℓ. We solve this problem by a clever application of the *Deletion Method* by [Janson and Ruciński, 2002].

Steps 1-4 are common between FAIRSAMPLE and FAIRELIMINATE. **Steps 5-6** Each variable Y_i is rounded w.r.t function $\frac{ay_i^*}{k}$ for a small constant $a > 0$; let $\mathcal{F}$ denote the set of all subsets that are sampled in this step. **Steps 7-12** Then, in the alteration step, the algorithm removes any sets that intersect with other sets in $\mathcal{F}$, and returns the remaining sets as the final packing. Theorems 2 and 3 state the approximation guarantees provided by FAIRELIMINATE on both the objective and SRF.

Intuition behind the analysis of FAIRELIMINATE. The key idea behind FAIRELIMINATE is to choose an appropriate value for a[7] such that *randomized rounding* results in a set of variables $\{Y_i\}$ with an acceptable "expected weight".

[6]Expected cardinality refers to the expected number of elements selected in this step i.e., $\mathbb{E}[\sum_{i\in[m]} |S_i|Y_i]$.

[7]The value of $a > 0$ is selected via the analysis of the algorithm.

Algorithm 2 FAIRELIMINATE

Require: $a < p^* \frac{1-\delta}{6k}$ where $\delta = \frac{1}{n^{0.25}}$ and $p^* = \min_{\ell \in [c]} p_\ell$
Solve LP (3) to get an optimal fractional solution $\{y_i^*\}$
if LP (3) is infeasible **then**
    return return "Infeasible"
end if
For each $i \in [m]$ sample $Y_i \sim \text{Ber}(\frac{a y_i^*}{k})$ for a small constant $a > 0$
$\mathcal{F} = \{S_i \in \mathcal{S} \mid Y_i = 1\}$
$\mathcal{F}_1 \leftarrow \varnothing$
for S_i and S_j such that $S_i \cap S_j \neq \varnothing$ **do**
    if $Y_i = 1$ and $Y_j = 1$ **then**
        $\mathcal{F}_1 \leftarrow \mathcal{F}_1 \cup \{S_i, S_j\}$
    end if
end for
return $\mathcal{F} - \mathcal{F}_1$

We then prove that the expected number of variables dropped in the alteration step is minimal, resulting in a high-quality approximation of the packing objective.

It's important to note the $\{Z_i\}$ are dependent random variables with both negative and positive correlations. As a result, we cannot directly use the Chernoff-Hoeffding bounds that are applied to sums of negatively correlated random variables, as outlined in [Panconesi and Srinivasan, 1997]. This creates a challenge in achieving concentration on the desired value of N_ℓ. Instead we express Z_i as a function of the independent $\{Y_i\}$, allowing us to utilize various results on the polynomials of independent random variables such as [Kim and Vu, 2000; Janson and Ruciński, 2002; Schudy and Sviridenko, 2011]. We know that the random variable Z_i takes the value either 1 or 0 based on two events: (i) the random variable $Y_i = 1$ in the *randomized rounding* step, and (ii) none of its neighbors $j \in N(i)$[8] have $Y_j = 1$. Therefore, $Z_i = Y_i(1 - \max_{j \in N(i)} Y_j)$, allowing us to express N_ℓ as follows:

$$N_\ell = \sum_{i \in [m]} v_{i\ell}(Y_i - Y_i \max_{j \in N(i)} Y_j) \tag{4}$$

$$\geq \sum_{i \in [m]} v_{i\ell}(Y_i - Y_i \sum_{j \in N(i)} Y_j) \tag{5}$$

$$\geq \underbrace{\sum_{i \in [m]} v_{i\ell} Y_i}_{A_\ell} - 2k \underbrace{\sum_{\substack{1 \leq i < j \leq m \\ S_i \cap S_j \neq \varnothing}} Y_i Y_j}_{B} . \tag{6}$$

By applying the Hoeffding bound [Hoeffding, 1994], we can show that $A_\ell = \sum_{i \in [m]} v_{i\ell} Y_i$ is concentrated around its mean with probability p. Additionally, we can use the Deletion method by [Janson and Ruciński, 2002] to prove that B is also concentrated around its mean with probability q. The crucial step in applying the Deletion Method is the construction of a hypergraph on the random variables $\{Y_i\}$, which can be challenging. Finally, by a union bound and suitable p and q, we find concentration on the lower bound of $A_\ell - 2kB$.

Theorem 2 is a direct outcome of this approach. Theorem 3 is an enhancement of Theorem 2 resulting from a more refined analysis of the concentration of B in the Deletion Method. The complete justification for these theorems can be found in the full version of the paper.

[8] $N(i) = \{S_j : S_i \cap S_j \neq \varnothing\}$.

5 Experiments

Our experiments run on commodity hardware using Python 3.6 with NumPy [Harris *et al.*, 2020] as well as IP and LP solvers in Gurobi 9.1.2. The empirical evaluation of our algorithms uses both purely synthetic and realistic kidney exchange FAIRSP instances. To benchmark our algorithms, we include i) ALG-BLIND, the LP-rounding (fairness agnostic) k-set packing randomized algorithm from [Brubach *et al.*, 2019] as well as ii) IP-FSP and iii) IP-BLIND, which solve the integer restrictions of LP (3) with and without fairness constraints, respectively. The integer restrictions of IP-FSP and IP-BLIND make the fairness guarantee be satisfied deterministically as opposed to in expectation like RF.

Evaluation of randomized algorithms. FAIRSAMPLE and ALG-BLIND provide objective guarantees in expectation thus their reported objective values are the mean objectives of 300 independent LP solution roundings. FAIRSAMPLE uses $L = 300$ to estimate q_i' values. Since FAIRELIMINATE provides ex-post fairness guarantees the reported objective values correspond to a *single* LP rounding per FAIRSP instance.

Kidney exchange datasets. We use realistic kidney exchange instances drawn from the US-based United Network for Organ Sharing (UNOS) fielded exchange program; our synthetic instances are generated from this data in the standard way [Dickerson *et al.*, 2019]. In real programs,[9] both *blood type* and *patient sensitization*, a measure that is high when a patient's body has built up many antibodies to potential donor organs, are taken into account when deciding what type of participant to prioritize. Thus, in our experiments, we treat sensitization as a binary sensitive attribute (taking value 1 if the patient is sensitized and 0 otherwise), and separately treat patient blood type as a different binary sensitive attribute (taking value 1 if the patient is a hard-to-match O-type blood, and 0 otherwise). These binary attributes give rise to two datasets, each with $c = 2$, which we call the *blood group* and *sensitization* datasets. We generate graphs of size 64, 128, 256, and 512 from that real data, noting that the largest exchanges in the world hover at or slightly below 256.[11] Recall as mentioned in the introduction,

[9] For example, the OPTN/UNOS[10] exchange transparently reports its prioritization points in Table 13-2; high values for the sensitization metric CPRA are prioritized extensively, and O-type candidates are prioritized as well. For summaries of countries' exchanges in the European Union and the United Kingdom, we refer the reader to [Biró *et al.*, 2019].

[10] https://optn.transplant.hrsa.gov/media/3390/optn-policies-effective-as-of-dec-4-2019-board-adoptions.pdf

[11] As examples, we direct the reader to [Biró *et al.*, 2019], a recent survey of 17 European countries' national exchange programs, stating that the UK "has become the largest operating [kidney exchange program] in Europe, with 250 recipient-donor pairs registered per matching run," followed by the Netherlands and then Spain, with 110 recipient-donor pairs registered per run, and a long tail of countries afterward. In the US, programs tend to match at a faster cadence

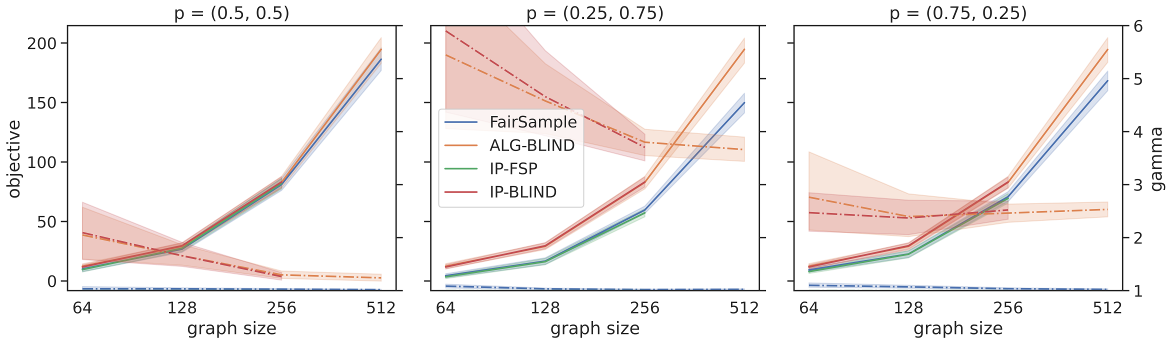

Figure 2: Plots for the *sensitization* dataset. Left (right) y-axes and solid (dashed) lines correspond to packing objectives (RF γ values).

the FAIRSP instances have $\mathcal{S}$ consisting of all 2- and 3-cycles in said graphs. As a result, the largest FAIRSP instances have $n = 512$ and $m \approx 28000$. We omit IP solutions for these largest instances as solving them takes too long and they only serve to motivate probabilistic fairness. Each graph size has 32 instances for which we create fair instances with proportionality vector $\vec{p} \in \{(0.25, 0.75), (0.5, 0.5), (0.75, 0.25)\}$. $w_i = |S_i|$ for all $S_i \in \mathcal{S}$, and $k = 3$.

Experiments on kidney exchange datasets. Altogether, the experiments solve a total of 864 IPs with fairness constraints, 960 LPs with fairness constraints, 96 IPs without fairness constraints and 128 LPs without fairness constraints. Both FAIRSAMPLE and FAIRELIMINATE produce solutions swiftly, taking up to 7 minutes for the largest instances. In both cases the bottleneck is solving the LP. Given there are multiple FAIRSP instances per graph size, the plots report the mean over all instances and the shaded regions capture 95% of the instances.

Figure 1 uses all 32 instances with $n = 256$ from the *blood group* dataset over each of 40 choices of $\vec{p}$; hence running FAIRSAMPLE 1280 times in 9 hours. Figure 2 shows FAIRSAMPLE achieves remarkably improved fairness at little cost to the objective compared to its fairness agnostic counterparts ALG-BLIND and IP-BLIND. IP-FSP also ran into several infeasible instances while this was never an issue for FAIRSAMPLE, further making the case for probabilistic fairness.

Experiments in Figure 3 use FAIRELIMINATE to round the 960 LPs with fairness constraints with 25 evenly spaced choices of the parameter $a \in [0.1, 0.99]$. Both the approximation and fairness factors improve as a increases. Although the worst-case guarantees in Theorem 3 are valid for $a < (1-\delta)p^*/6k$ (for any $\delta > 0$), in this real-world example the fairness guarantees continue to improve as a approaches 1. One explanation is that as a increases, the number of initially sampled sets increases but so does the number of sets dropped in the alterations step; however, since our instances are sparse the number of sets dropped is not too high. For comparison, we include the approximation and fairness factors from a *single* rounding of FAIRSAMPLE. Interestingly, here FAIRSAMPLE outperforms FAIRELIMINATE in ex-post fairness. One possible explanation is that the realistic instances are far from worst-case.

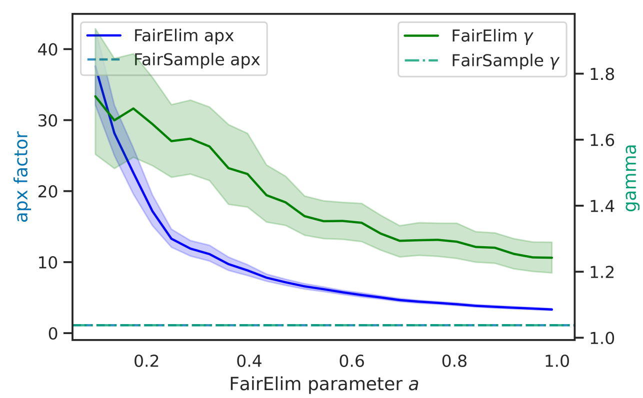

Figure 3: Blue (green) lines show approximation (fairness) factors and correspond to the left (right) y-axes. For comparison, we plot the average approximation factor ($= 1.12$) and γ ($= 1.04$) achieved by FAIRSAMPLE. Approximation factors are the LP objective divided by the rounded solution's objective.

DatasetMaxK. We generate 200 purely synthetic FAIRSP instances with $n = 50$, $m = 1000$, $c \in \{2, 3\}$, and $\vec{p}$ generated uniformly at random. Each instance has a parameter $\max\_k \in \{3, 8, 13, \ldots, 49\}$. For each parameter configuration above, we independently generate 10 instances. Then this parameter $\max\_k$ is used to sample sets. For each $i \in [m]$, $S_i \in \mathcal{S}$ is a random subset of $\mathcal{U}$ with cardinality $k_i \sim \mathcal{N}(\max\_k, 2)$ (while ensuring k_i stays in between 2 and 50). For each $S_i \in \mathcal{S}$, w_i is an integer drawn uniformly between 5 and 35. Element colors are chosen uniformly at random. As expected, the objectives decrease as k increases, thus the approximation factor of FAIRSAMPLE increases alongside k.

which can keep pool sizes low; for example, arguably the most successful program in the US, which is run by the National Kidney Registry (NKR), had a pool size hovering around 150 recipient-donor pairs in Q1 2022 [National Kidney Registry, 2022].

Acknowledgments

We thank the anonymous reviewers for their insightful comments that helped improve the paper. This research was supported in part by NSF CAREER Award IIS-1846237, NSF Award CCF-1918749, NSF Award CCF-1852352, NSF Award SMA-2039862, NIST MSE Award #20126334, DARPA GARD #HR00112020007, DARPA SI3-CMD #S4761, DoD WHS Award #HQ003420F0035, ARPA-E DIFFERENTIATE Award #1257037, ARL Award W911NF2120076, LTS MAGICAPPLE, and gifts by research awards from Amazon, and Google.

References

[Abraham *et al.*, 2007] David J Abraham, Avrim Blum, and Tuomas Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *ACM Conference on Electronic Commerce (EC)*, pages 295–304, 2007.

[Adamczyk *et al.*, 2015] Marek Adamczyk, Fabrizio Grandoni, and Joydeep Mukherjee. Improved approximation algorithms for stochastic matching. In *Algorithms-ESA 2015*, pages 1–12. Springer, 2015.

[Anegg *et al.*, 2021] Georg Anegg, Haris Angelidakis, and Rico Zenklusen. Simpler and stronger approaches for non-uniform hypergraph matching and the füredi, kahn, and seymour conjecture. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 196–203. SIAM, 2021.

[Ashlagi *et al.*, 2011] Itai Ashlagi, Duncan S Gilchrist, Alvin E Roth, and Michael A Rees. Nonsimultaneous chains and dominos in kidney-paired donation—revisited. *American Journal of Transplantation (AJT)*, 11(5):984–994, 2011.

[Ashlagi *et al.*, 2019] Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi. On matching and thickness in heterogeneous dynamic markets. *Operations Research*, 67(4):927–949, 2019.

[Asudeh *et al.*, 2022] Abolfazl Asudeh, Tanya Berger-Wolf, Bhaskar DasGupta, and Anastasios Sidiropoulos. Maximizing coverage while ensuring fairness: A tale of conflicting objectives. *Algorithmica*, pages 1–45, 2022.

[Bansal *et al.*, 2010] Nikhil Bansal, Nitish Korula, Viswanath Nagarajan, and Aravind Srinivasan. On k-column sparse packing programs. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 369–382. Springer, 2010.

[Bastani *et al.*, 2019] Osbert Bastani, Xin Zhang, and Armando Solar-Lezama. Probabilistic verification of fairness properties via concentration. *Proceedings of the ACM on Programming Languages*, 3(OOPSLA):1–27, 2019.

[Bei *et al.*, 2022] Xiaohui Bei, Shengxin Liu, Chung Keung Poon, and Hongao Wang. Candidate selections with proportional fairness constraints. *Autonomous Agents and Multi-Agent Systems*, 36, 2022.

[Berman, 2000] Piotr Berman. A d/2 approximation for maximum weight independent set in d-claw free graphs. In *Scandinavian Workshop on Algorithm Theory*, pages 214–219. Springer, 2000.

[Biro *et al.*, 2009] Péter Biro, David F Manlove, and Romeo Rizzi. Maximum weight cycle packing in directed graphs, with application to kidney exchange programs. *Discrete Mathematics, Algorithms and Applications*, 1(04):499–517, 2009.

[Biró *et al.*, 2019] Péter Biró, Bernadette Haase-Kromwijk, et al. Building kidney exchange programmes in Europe: an overview of exchange practice and activities. *Transpl.*, 103(7), 2019.

[Blum *et al.*, 2015] Avrim Blum, John P Dickerson, Nika Haghtalab, Ariel D Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. In *Conference on Economics and Computation (EC)*, pages 325–342, 2015.

[Brubach *et al.*, 2019] Brian Brubach, Karthik A. Sankararaman, Aravind Srinivasan, and Pan Xu. Algorithms to approximate column-sparse packing problems. *ACM Trans. Algorithms*, 16(1), November 2019.

[Chierichetti *et al.*, 2017] Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. In *Neural Information Processing Systems (NeurIPS)*, 2017.

[Diana *et al.*, 2021] Emily Diana, Wesley Gill, Michael Kearns, Krishnaram Kenthapadi, and Aaron Roth. Minimax group fairness: Algorithms and experiments. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2021.

[Dickerson *et al.*, 2014] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Price of fairness in kidney exchange. In *Int. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2014.

[Dickerson *et al.*, 2019] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Failure-aware kidney exchange. *Management Science*, 65(4):1768–1791, 2019.

[Esmaeili *et al.*, 2023] Seyed A. Esmaeili, Sharmila Duppala, Davidson Cheng, Vedant Nanda, Aravind Srinivasan, and John P. Dickerson. Rawlsian fairness in online bipartite matching: Two-sided, group, and individual. In *Conf. on Artificial Intelligence (AAAI)*, 2023.

[Farnadi *et al.*, 2021] Golnoosh Farnadi, William St-Arnaud, Behrouz Babaki, and Margarida Carvalho. Individual fairness in kidney exchange programs. In *Conf. on Artificial Intelligence (AAAI)*, volume 35, pages 11496–11505, 2021.

[Füredi *et al.*, 1993] Zoltán Füredi, Jeff Kahn, and Paul D. Seymour. On the fractional matching polytope of a hypergraph. *Combinatorica*, 13(2):167–180, 1993.

[Fürer and Yu, 2014] Martin Fürer and Huiwen Yu. Approximating the k-set packing problem by local improvements. In *International Symposium on Combinatorial Optimization*, pages 408–420. Springer, 2014.

[García-Soriano and Bonchi, 2020] David García-Soriano and Francesco Bonchi. Fair-by-design matching. *Data*

Mining and Knowledge Discovery, 34(5):1291–1335, 2020.

[Harris *et al.*, 2020] Charles Harris, K Jarrod Millman, et al. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.

[Hazan *et al.*, 2006] Elad Hazan, Shmuel Safra, and Oded Schwartz. On the complexity of approximating k-set packing. *Computational Complexity*, 15(1):20–39, 2006.

[Hoeffding, 1994] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. In *The collected works of Wassily Hoeffding*, pages 409–426. Springer, 1994.

[Janson and Ruciński, 2002] Svante Janson and Andrzej Ruciński. The infamous upper tail. *Random Structures & Algorithms*, 20(3):317–342, 2002.

[Karp, 1972] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer, 1972.

[Kim and Vu, 2000] Jeong Han Kim and Van H Vu. Concentration of multivariate polynomials and its applications. *Combinatorica*, 20(3):417–434, 2000.

[Li *et al.*, 2014] Jian Li, Yicheng Liu, Lingxiao Huang, and Pingzhong Tang. Egalitarian pairwise kidney exchange: fast algorithms vialinear programming and parametric flow. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, pages 445–452, 2014.

[Lin *et al.*, 2019] Mugang Lin, Jianxin Wang, Qilong Feng, and Bin Fu. Randomized parameterized algorithms for the kidney exchange problem. *Algorithms*, 12(2):50, 2019.

[Ma *et al.*, 2022] Will Ma, Pan Xu, and Yifan Xu. Group-level fairness maximization in online bipartite matching. In *Int. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2022.

[Ma, 2014] Will Ma. Improvements and generalizations of stochastic knapsack and multi-armed bandit approximation algorithms. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1154–1163. SIAM, 2014.

[McElfresh and Dickerson, 2018] Duncan C McElfresh and John P Dickerson. Balancing lexicographic fairness and a utilitarian objective with application to kidney exchange. In *Conf. on Artificial Intelligence (AAAI)*, 2018.

[Nanda *et al.*, 2020] Vedant Nanda, Pan Xu, Karthik Abhinav Sankararaman, John Dickerson, and Aravind Srinivasan. Balancing the tradeoff between profit and fairness in rideshare platforms during high-demand hours. In *Conf. on Artificial Intelligence (AAAI)*, volume 34, pages 2210–2217, 2020.

[National Kidney Registry, 2022] National Kidney Registry. National Kidney Registry quarterly report: Q1 2022. https://www.kidneyregistry.org/wp-content/uploads/2022/05/nkr-report-2022-Q1_v16.pdf, 2022. [Online; accessed 19 May 2022].

[Neuwohner, 2021] Meike Neuwohner. An improved approximation algorithm for the maximum weight independent set problem in d-claw free graphs. *arXiv preprint arXiv:2106.03545*, 2021.

[Panconesi and Srinivasan, 1997] Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff–hoeffding bounds. *SIAM Journal on Computing*, 26(2):350–368, 1997.

[Roth *et al.*, 2005] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Pairwise kidney exchange. *Journal of Economic Theory (JET)*, 125(2):151–188, 2005.

[Roth, 2015] Alvin E Roth. *Who gets what–and why: the new economics of matchmaking and market design*. Houghton Mifflin Harcourt, 2015.

[Sankar *et al.*, 2021] Govind S Sankar, Anand Louis, Meghana Nasre, and Prajakta Nimbhorkar. Matchings with group fairness constraints: Online and offline algorithms. *arXiv preprint arXiv:2105.09522*, 2021.

[Schudy and Sviridenko, 2011] Warren Schudy and Maxim Sviridenko. Bernstein-like concentration and moment inequalities for polynomials of independent random variables: multilinear case. *arXiv preprint arXiv:1109.5193*, 2011.

[Sun *et al.*, 2021] Zhaohong Sun, Taiki Todo, and Toby Walsh. Fair pairwise exchange among groups. In *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 419–425, 2021.

[Tsang *et al.*, 2019] Alan Tsang, Bryan Wilder, Eric Rice, Milind Tambe, and Yair Zick. Group-fairness in influence maximization. In *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, 2019.

[Xiao and Wang, 2018] Mingyu Xiao and Xuanbei Wang. Exact algorithms and complexity of kidney exchange. In *Int. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 555–561, 2018.