Fair Pairwise Exchange among Groups

Zhaohong Sun[1] , **Taiki Todo**[2] and **Toby Walsh**[1]
[1]UNSW Sydney
[2]Kyushu University
{zhaohong.sun, t.walsh}@unsw.edu.au, todo@inf.kyushu-u.ac.jp

Abstract

In this paper, we study the pairwise organ exchange problem among groups motivated by real-world applications. We consider two types of group formulation: either each group represents a certain type of patient-donor pairs who are compatible with the same set of organs, or a set of patient-donor pairs who reside in the same region. We consider a natural research question how to match maximum number of pairwise compatible patient-donor pairs in a fair and individually rational way. We first propose a natural fairness concept that is applicable to both types of group formulation. We then design a polynomial-time algorithm that checks whether there exists a matching satisfying optimality, individual rationality and fairness. Finally, we present several running time upper bounds for computing such matchings for different graph structures.

1 Introduction

Due to a shortage of organ transplantation from deceased donors, living donations have become a significant approach to saving lives of patients who suffer from serious organ dysfunction. One issue of living donations is that organs from willing donors may be medically incompatible with intended patients. This can be partially overcome with the introduction of *organ exchange*, which allows patients to swap their donors with others to obtain compatible organs [Roth *et al.*, 2004]. Since transplant operations are usually conducted simultaneously, some fixed upper bound is always imposed on the length of exchange cycles. Pairwise organ exchange is the most common form in real-life, involving just two pairs of patient and donor [Roth *et al.*, 2005]. The mechanism design of the organ exchange market has attracted considerable attention from both economics and computer science [Abraham *et al.*, 2007; Bertsimas *et al.*, 2013; Dickerson *et al.*, 2014; Hajaj *et al.*, 2015; Dickerson and Sandholm, 2017; Dickerson *et al.*, 2019; Ergin *et al.*, 2020; Freedman *et al.*, 2020].

In this paper, we study the organ exchange problem from a different perspective. We concentrate on a pairwise organ exchange market where all patient-donor pairs are partitioned into disjoint groups. The idea of dividing patient-donor pairs into groups is motivated by applications. Next, we describe two types of group formulation that arises in recent literature.

In the first type of group formulation, a group of agents represents a certain type of patient-donor pairs where patients are compatible with the same set of organs in the market. For instance, Dickerson *et al.* [2017] introduce a new model for kidney exchange that classifies all participating patient-donor pairs into a fixed number of types, based on a common set of attributes, e.g., blood type, tissue type, age, insurance, willingness to travel. Ergin *et al.* [2017] study the dual-donor organ exchange problem w.r.t. lung and liver exchange. They consider a simplified model for theoretical analysis, by grouping patient-donor pairs together according to their blood types without taking tissue-type or size compatibility into account.

In the second form of group formulation, each group represents a set of patient-donor pairs who locate in a certain area. In other words, groups are formed geographically. For instance, each group can represent a set of patient-donor pairs in a certain hospital [Ashlagi *et al.*, 2015]. Each group can also represent a state or territory in the national kidney exchange market, e.g. Australian Organ and Tissue Authority [Mattei *et al.*, 2017]. Or each group can represent some country in the European organ exchange program [Biró *et al.*, 2019].

Although the idea of dividing patient-donors into groups has been proposed in the previous literature, scant attention was paid to the fair allocation of patient-donor pairs among groups. One exception is [Biró *et al.*, 2019] that considers core allocations, i.e. outcomes that cannot be improved upon by a coalition of agents. However, the core may be empty and it is co-NP-hard to check its existence. In contrast, we focus on a fundamental research question *how to design efficient algorithms that match a maximum number of pairwise compatible patient-donor pairs in a fair and individual rational way*.

The contributions of this paper are summarized as follows. First, we propose a straightforward fairness concept for organ exchange among groups based on a notion of *selection ratio*, which provides us flexibility to capture different ideas such as egalitarianism and proportionality. Second, we introduce a general polynomial-time algorithm that finds a matching satisfying maximality, individual rationality and fairness whenever it exists. Third, we provide several running time upper bounds of finding such matchings for different graph structures w.r.t. two forms of group formulation.

2 Model

We consider a pairwise organ exchange problem via a compatibility graph $G = (V, E)$ where each vertex $v \in V$ represents a patient-donor pair. We assume that for each vertex, the patient is incompatible with his donor, otherwise they will conduct the transplant immediately rather than participating in the exchange program [Roth *et al.*, 2007].

The edge set E is specified as follows: There is an edge between two vertices $i, j \in V$ if the donor of pair i is compatible with the patient of pair j and the donor of pair j is compatible with the patient of pair i. The vertex set V is partitioned into k disjoint sets where each set $V_i \subseteq V$ represents a group of patient-donor pairs. Let $P = \{V_1, \cdots, V_k\}$ denote the *partition* of vertices, i.e., $V_i \cap V_j = \emptyset$ for any $i \neq j$ and $\bigcup_{V_i \in P} V_i = V$.

We consider two types of compatible graph structures depending on whether or not all vertices within the same group have the same set of neighbors, corresponding to two types of group formulation. If all vertices within the same group have the same set of neighbors, then we refer to such graphs as *compatible graphs with identical neighbors*.

A matching M in G is a set of edges without common vertices, and a maximum matching is a matching that contains the largest possible number of edges. Let $\mathbb{M}$ denote the set of all possible matchings in G and let $\mathbb{M}^*$ denote the set of all maximum matchings in G. Given a matching M in G, let $|M_i|$ denote the number of matched vertices from group V_i. For each group V_i, let $\max(V_i) = \max_{M \in \mathbb{M}^*} |M_i|$ denote the *maximum bound* of matched vertices from group V_i among $\mathbb{M}^*$, and let $\min(V_i) = \min_{M \in \mathbb{M}^*} |M_i|$ denote the *minimum bound* of matched vertices from group V_i among $\mathbb{M}^*$.

Let $G[V_i]$ denote a subgraph of G induced by vertices from group V_i only. Consider any maximum matching M' in the subgraph $G[V_i]$. Let $\widetilde{\min}(V_i) = 2 \cdot |M'|$ denote the *modified minimum bound* of group V_i, which is equal to twice the size of the maximum matching M' in $G[V_i]$. Note that the modified minimum bound of group V_i is the largest number of matched pairs from group V_i if only exchanges between vertices within group V_i are allowed. The modified minimum bound is critical for the definition of individual rationality.

3 Desirable Properties

For our problem, an algorithm takes as input a compatibility graph G, and outputs a matching M as the outcome. Next, we introduce three desirable properties that an outcome should satisfy. The first property is *optimality*, which requires that a matching should maximize the number of exchanges among compatible patient-donor pairs.

Definition 1 (Optimality)**.** *Given a compatible graph G, a matching M in G satisfies optimality if M is a maximum matching in G.*

The second property, *individual rationality*, requires that in an individually rational matching M, the number of matched pairs $|M_i|$ of each group V_i should not be smaller than its modified minimum bound $\widetilde{\min}(V_i)$. This property guarantees that, each group V_i will receive at least the same number of matched pairs as group V_i conducts exchanges by itself.

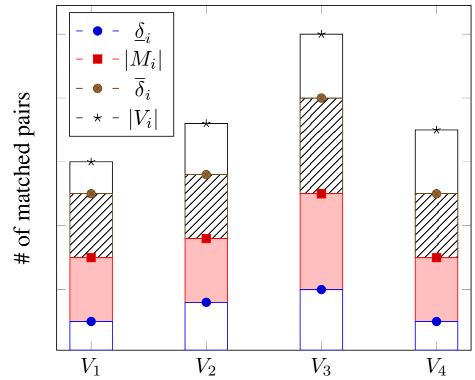

Figure 1: Given a matching M, for each group V_i, the colored region represents $|M_i| - \underline{\delta}_i$ and the shaded region represents $\overline{\delta}_i - |M_i|$. The height of each bar equals the size $|V_i|$ of each group.

Definition 2 (Individual Rationality)**.** *Given a compatibility graph $G = (V, E)$ and a partition P of vertices V, a matching M in G satisfies individual rationality, if for each group $V_i \in P$, we have $|M_i| \geq \widetilde{\min}(V_i)$.*

Note that for compatible graphs with identical neighbors, individual rationality is trivially satisfied, because each group forms an independent set and the modified minimum bound is zero for each group.

Next, we introduce an important notion called *selection ratio* that is key to the definition of our third property *fairness*. Let $\overline{\delta}_i$ and $\underline{\delta}_i$ denote some upper and lower bound for group V_i with $\overline{\delta}_i \geq \underline{\delta}_i$. Intuitively, these two bounds represent two target quotas s.t. we expect the number of matched pairs from group V_i to fall into the range of these two bounds.

Given a matching M, the *selection ratio* $\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i)$ of group V_i w.r.t. $\overline{\delta}_i$ and $\underline{\delta}_i$ is represented as a fraction, where the numerator is the difference between the number of matched pairs $|M_i|$ from V_i in the matching M and the lower bound $\underline{\delta}_i$, and the denominator is the difference between the upper bound $\overline{\delta}_i$ and the lower bound $\underline{\delta}_i$. Formally,

Definition 3 (Selection Ratio)**.** *Given a matching M in G, two quotas $\overline{\delta}_i$ and $\underline{\delta}_i$ with $\overline{\delta}_i \geq \underline{\delta}_i$, the selection ratio of group V_i w.r.t. $\overline{\delta}_i$ and $\underline{\delta}_i$ is*

$$\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = \frac{|M_i| - \underline{\delta}_i}{\overline{\delta}_i - \underline{\delta}_i} \quad \textit{for } \overline{\delta}_i > \underline{\delta}_i$$

When $\overline{\delta}_i = \underline{\delta}_i$, we assume $\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = -\infty$ if $|M_i| < \overline{\delta}_i$ and $\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = \infty$ if $|M_i| \geq \overline{\delta}_i$.

The selection ratio measures to what extent, the number of matched pairs $|M_i|$ from group V_i surpasses the lower bound $\underline{\delta}_i$ on the scale of the lower bound to the upper bound, as shown in Figure 1. Note that it is flexible to choose upper and lower bounds in the formula of selection ratio, which allows us to capture different reasonable ideas, e.g. egalitarianism and proportionality. We will discuss different choices of upper and lower bounds in Section 4.

The third property, *fairness among groups*, requires that the minimum selection ratio among all groups should be maximized among all matchings. This is a natural and unified fairness concept for different compatibility graph structures.

Definition 4 (Fairness among Groups)**.** *Given the set of all matchings $\mathbb{M}$ in G, a partition P of vertices V and two vectors of quotas $\overline{\delta} = (\overline{\delta}_i)_{V_i \in P}$ and $\underline{\delta} = (\underline{\delta}_i)_{V_i \in P}$ with $\overline{\delta}_i \geq \underline{\delta}_i$ for each group $V_i \in P$, a matching M in G is fair (among groups) w.r.t. $\overline{\delta}$ and $\underline{\delta}$ if it maximizes the minimal selection ratio among all groups:*

$$M \in \arg\max_{M' \in \mathbb{M}} \min_{V_i \in P} \alpha(|M'_i|, \overline{\delta}_i, \underline{\delta}_i)$$

W.L.O.G, we assume that $\overline{\delta}_i \geq \underline{\delta}_i$ for each group $V_i \in P$ for the rest of the paper.

4 Choice of Upper and Lower Bounds

In this section, we propose several choices of upper and lower bounds and assign different names to corresponding selection ratios as summarized in Table 1.

Egalitarianism

For each group V_i, if we set $\overline{\delta}_i = |V|$ and $\underline{\delta}_i = 0$, then the *egalitarian-selection-ratio* (Egalitarian) $\alpha(\cdot)$ represents the ratio between the number of matched pairs $|M_i|$ from group V_i and the total number of vertices $|V|$:

$$\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = \frac{|M_i| - \underline{\delta}_i}{\overline{\delta}_i - \underline{\delta}_i} = \frac{|M_i|}{|V|}$$

A fair matching w.r.t. *Egalitarian* tries to equalize the numbers of matched pairs among all groups. Note that it is possible to choose another positive constant as $\overline{\delta}_i$ instead of $|V|$.

Proportional to Group Sizes (Group-Size)

For each group V_i, if we set $\overline{\delta}_i = |V_i|$ and $\underline{\delta}_i = 0$, then the *group-size-selection-ratio* (Group-Size) $\alpha(\cdot)$ represents the ratio between the number of matched pairs $|M_i|$ from group V_i and its group size $|V_i|$:

$$\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = \frac{|M_i| - \underline{\delta}_i}{\overline{\delta}_i - \underline{\delta}_i} = \frac{|M_i|}{|V_i|}$$

A fair matching w.r.t. *Group-Size* tries to ensure that for each group, the number of matched pairs $|M_i|$ from group V_i is proportional to its group size $|V_i|$.

Proportional to Maximum Bounds (Maximum)

For each group V_i, if we set $\overline{\delta}_i = \max(V_i)$ and $\underline{\delta}_i = 0$, then the *maximum-bound-selection-ratio* (Maximum) $\alpha(\cdot)$ represents the ratio between the number of matched pairs $|M_i|$ from group V_i and its maximum bound $\max(V_i)$:

$$\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = \frac{|M_i| - \underline{\delta}_i}{\overline{\delta}_i - \underline{\delta}_i} = \frac{|M_i|}{\max(V_i)}$$

A fair matching w.r.t. *Maximum* tries to ensure that for each group, the number of matched pairs $|M_i|$ from group V_i is proportional to its maximum bound $\max(V_i)$.

	upper bound	lower bound
Egalitarian	$\|V\|$	0
Group-Size	$\|V_i\|$	0
Maximum	$\max(V_i)$	0
Minimum	$\min(V_i) + 1$	0
Max-Min	$\max(V_i)$	$\min(V_i)$
Max-M-Min	$\max(V_i)$	$\widetilde{\min}(V_i)$

Table 1: Different Choices of Upper and Lower Bounds

Proportional to Minimum Bounds (Minimum)

For each group V_i, if we set $\overline{\delta}_i = \min(V_i)+1$ and $\underline{\delta}_i = 0$, then the *minimum-bound-selection-ratio* (Minimum) $\alpha(\cdot)$ represents the ratio between the number of matched pairs $|M_i|$ from group V_i and its minimum bound $\min(V_i)$ plus 1:

$$\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = \frac{|M_i| - \underline{\delta}_i}{\overline{\delta}_i - \underline{\delta}_i} = \frac{|M_i|}{\min(V_i) + 1}$$

A fair matching w.r.t. *Minimum* tries to ensure that for each group, the number of matched pairs $|M_i|$ from group V_i is proportional to its minimum bound $\min(V_i)$ plus 1. Note that we consider minimum bound $\min(V_i)$ plus 1 to avoid the case that $\min(V_i) = 0$. Similarly, we can replace the minimum bound $\min(V_i)$ by the modified minimum bound $\widetilde{\min}(V_i)$.

Proportional to the Range of Maximum and Minimum Bounds (Max-Min)

For each group V_i, if we set $\overline{\delta}_i = \max(V_i)$ and $\underline{\delta}_i = \min(V_i)$, then the *maximum-minimum-selection-ratio* (Max-Min) $\alpha(\cdot)$ represents the ratio between the number of matched pairs $|M_i|$ from group V_i minus its minimum bound, and the difference between its maximum and minimum bounds.

$$\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = \frac{|M_i| - \underline{\delta}_i}{\overline{\delta}_i - \underline{\delta}_i} = \frac{|M_i| - \min(V_i)}{\max(V_i) - \min(V_i)}$$

A fair matching w.r.t. *Max-Min* tries to ensure that for each group V_i, the number of matched pairs $|M_i|$ from group V_i minus it minimum bound is proportional to the difference between its maximum and minimum bounds.

Proportional to the Range of Maximum and Modified Minimum Bounds (Max-M-Min)

For each group V_i, if we set $\overline{\delta}_i = \max(V_i)$ and $\underline{\delta}_i = \widetilde{\min}(V_i)$, then the *maximum-modified-minimum-selection-ratio* (Max-M-Min) $\alpha(\cdot)$ represents the ratio between the number of matched pairs $|M_i|$ from group V_i minus its minimum bound, and the difference between its maximum and modified minimum bounds.

$$\alpha(|M_i|, \overline{\delta}_i, \underline{\delta}_i) = \frac{|M_i| - \underline{\delta}_i}{\overline{\delta}_i - \underline{\delta}_i} = \frac{|M_i| - \widetilde{\min}(V_i)}{\max(V_i) - \widetilde{\min}(V_i)}$$

A fair matching w.r.t. *Max-M-Min* tries to ensure that for each group, the number of matched pairs $|M_i|$ from group V_i minus its modified minimum bound is proportional to the difference between its maximum and modified minimum bounds.

5 Compatibility between Optimality, Individual Rationality and Fairness

In this section, we discuss whether or not there always exists a matching that satisfies optimality, individual rationality and fairness w.r.t. different choices of upper and lower bounds in Table 1. The results are different for compatibility graphs with identical and different neighbors, as summarized in Theorem 1 and 2, respectively.

Theorem 1. *For compatibility graphs with identical neighbors, there always exists a matching that satisfies optimality, individual rationality, and fairness w.r.t. any choice of upper and lower bounds in Table 1.*

Proof. Given a compatibility graph G with identical neighbors, let M' denote any fair matching w.r.t. any choice of upper and lower bounds in Table 1. If M' is not a maximum matching in G, by Berge's theorem [Berge, 1957], then there must exist an augmenting path p w.r.t. M' such that we can update M' to be a new matching $M' \oplus p$ in which the edges from $p \setminus M'$ are added while the edges from $M' \cap p$ are removed. Note that the vertices that are matched in M' are still matched in the new matching $M' \oplus p$. Eventually we can obtain a maximum matching M by iteratively finding all augmenting paths w.r.t. M', and for each group V_i, the number of matched pairs $|M_i|$ in M is at least as large as the number of matched pairs $|M_i'|$ in M'. Then the minimum selection ratio in matching M remains the same as the one in M', otherwise M' cannot be a fair matching. By the definition of fairness, the maximum matching M is also a fair matching.

Next, we show that the matching M also satisfies individual rationality. Recall that, we assume the patient-donor pair in each vertex is incompatible, which indicates that there is no edge between two vertices from the same group. Thus for each group V_i, its modified minimum bound $\widetilde{\min}(V_i)$ is always 0. Then for any maximum matching M', we have that for each group V_i, $\min(V_i) \leq |M_i| \leq \max(V_i)$ holds. Since $\widetilde{\min}(V_i) = 0 \leq \min(V_i)$, then any maximum matching in G is also individually rational. This completes the proof. □

Theorem 2. *For compatibility graphs with different neighbors, there always exists a matching satisfying optimality, individual rationality and fairness w.r.t. Max-M-Min. Fairness w.r.t. any other choice of upper and lower bounds in Table 1 is incompatible with optimality and individual rationality.*

Proof. Next, we show that given a compatibility graph G with different neighbors, fairness w.r.t. Max-M-Min is compatible with optimality and individual rationality. Let M denote any fair matching w.r.t. Max-M-Min selection ratio in G. If M is not a maximum matching, then we can update M to be a maximum matching as discussed in the proof for Theorem 1. For the sake of contradiction, suppose M does not satisfy individual rationality. Then there must exist a group V_i such that $|M_i| < \widetilde{\min}(V_i)$ and therefore its selection ratio $\alpha(\cdot)$ is negative (including the case $\widetilde{\min}(V_i) = \max(V_i)$).

Now consider another matching M^* obtained as follows. For each group V_i, let O_i denote a maximum matching in the subgraph $G[V_i]$ induced by vertices from group V_i only. The matching $M^* = \bigcup_{V_i \in P} O_i$ is the combination of each maximum matching O_i in the corresponding subgraph $G[V_i]$. Then we have that, for each group $|M_i^*| = \widetilde{\min}(V_i)$ holds, which satisfies individual rationality. However, this leads to a contradiction that the matching M maximizes the minimal selection ratio among all groups. Thus the assumption is wrong, and a maximum and fair matching w.r.t. Max-M-Min selection ratio is also individually rational.

We prove the second part of Theorem 2 by a counterexample in the Appendix due to space limitation. □

6 Algorithm Design

In this section, we present a general algorithm that finds a matching that achieves optimality, individual rationality and fairness w.r.t. different choices of upper and lower bounds in Table 1. The algorithm works for compatibility graphs with both identical and different neighbors, and it yields either a matching satisfying three properties whenever one exists or a NO-instance otherwise, as described in Algorithm 2.

Computing a Fair Matching

Next, we describe the first step of Algorithm 2 that computes a fair matching w.r.t. some choice of upper and lower quotas in Table 1, as described in Algorithm 1.

In order to develop a good intuition of how Algorithm 1 works, we postpone the following two technical details later: i) Some choices of upper and lower bounds require the maximum and minimum bounds of all groups. ii) Algorithm 1 iteratively invokes an Algorithm Γ that solves the following problem of *Matching with Quotas*. Intuitively, Algorithm Γ checks whether there exists a matching M s.t. for each group V_i, the number of matched pairs $|M_i|$ is not smaller than some target quota λ_i. We show how to design efficient algorithms regarding these two issues in Section 7.

Matching with Quotas	
Input:	A compatibility graph G, a partition P of vertices V, a vector of targets $\lambda = (\lambda_i)_{V_i \in P}$.
Question:	Whether there exists a matching M in G s.t. for each $V_i \in P$, $\lambda_i \leq |M_i|$ holds.

The basic idea of Algorithm 1 is to apply binary search to find the maximal selection ratio α such that there exists a matching where each group has a weakly larger selection ratio than α. Algorithm 1 takes as input a compatibility graph G, a vector of upper bounds $\overline{\delta} = (\overline{\delta}_i)_{1 \cdots k}$ and a vector of lower bounds $\underline{\delta} = (\underline{\delta}_i)_{1 \cdots k}$. During the process of Algorithm 1, we keep track of three variables α, α_ℓ and α_u, representing the current, the lower and the upper selection ratio respectively. In the initialization step, we set the lower selection ratio α_ℓ to be 0 and set the upper selection ratio α_u to be 1. For each group V_i, initialize its target quota λ_i to be 0. During each round of Algorithm 1, first compute the current selection ratio $\alpha = (\alpha_\ell + \alpha_u)/2$. Then for each group V_i, calculate its target quota λ_i w.r.t. α, by rounding up the value $\lceil \alpha \cdot (\overline{\delta}_i - \underline{\delta}_i) + \underline{\delta}_i \rceil$. Then we check whether there exists a matching M s.t. for each group V_i, $|M_i| \geq \lambda_i$ holds through Algorithm Γ. If so, then update the lower selection ratio α_ℓ

Input: $G, \overline{\delta} = (\overline{\delta}_i)_{V_i \in P}, \underline{\delta} = (\underline{\delta}_i)_{V_i \in P}$
Output: a fair matching M w.r.t. $\overline{\delta}$ and $\underline{\delta}$

Initialize a lower selection ratio $\alpha_\ell = 0$ and a upper selection ratio $\alpha_u = 1$
Initialize a target quota $\lambda_i = 0$ for each group V_i
while there exists some target λ_i whose value is different from last round **do** % including the initial round
  Set the current selection ratio $\alpha = (\alpha_\ell + \alpha_u)/2$
  for each group V_i **do**
    Compute the target quota λ_i corresponding to α
    $\lambda_i \leftarrow \lceil \alpha \cdot (\overline{\delta}_i - \underline{\delta}_i) + \underline{\delta}_i \rceil$
  if there exists a matching M in G s.t. for each $V_i \in P$, $|M_i| \geq \lambda_i$ holds **then** % Using Algorithm Γ
    $\alpha_\ell \leftarrow \alpha$ % Search between α and α_u
  else
    $\alpha_u \leftarrow \alpha$ % Search between α_ℓ and α
return a matching M

Algorithm 1: Computing a fair matching w.r.t. $\overline{\delta}$ and $\underline{\delta}$

to be α to search a larger selection ratio in the range $[\alpha, \alpha_u]$ in the next round; otherwise, update the upper selection ratio α_u to be α to search a smaller selection ratio in the range $[\alpha_\ell, \alpha]$ in the next round. Repeat these procedures whenever there exists some target λ_i whose value is different from last round, otherwise Algorithm 2 terminates.

Theorem 3. *Given a compatibility graph G and some choice of upper and lower bounds in Table 1, Algorithm 1 yields a fair matching in polynomial time.*

Computing a Maximum, Individually Rational and Fair Matching

Next, we give a high-level description of Algorithm 2 that checks the existence of a matching achieving optimality, individual rationality and fairness w.r.t. some choice of upper and lower bounds in Table 1.

The input consists of a compatibility graph G, and two vectors of quotas $\overline{\delta}$ and $\underline{\delta}$. The first step is to compute a fair matching M' through Algorithm 1 and let α denote the minimal selection ratio among all groups. The second step employs Algorithm Γ to check whether there exists a fair and individually rational matching M s.t. for each group V_i, we have $\max(\lceil \alpha \cdot (\overline{\delta}_i - \underline{\delta}_i) + \underline{\delta}_i \rceil, \widetilde{\min}(V_i)) \leq |M_i|$. In the final step, if there exists a fair and individually rational matching M, then we can update M to be a maximum matching. Otherwise, there does not exist such a matching.

Theorem 4. *Given a compatibility graph and some choice of upper and lower bounds in Table 1, Algorithm 2 finds a matching that satisfies optimality, individual rationality and fairness whenever it exists in polynomial time.*

7 Running Time Upper Bounds

This section is devoted to two remaining technical details in Section 6: i) how to compute the maximum and minimum bounds for each group and ii) how to design Algorithm Γ. We present different polynomial-time algorithms for different

Input: $G, \overline{\delta} = (\overline{\delta}_i)_{V_i \in P}, \underline{\delta} = (\underline{\delta}_i)_{V_i \in P}$
Output: a matching M that is maximum, individually rational and fair w.r.t. $\overline{\delta}$ and $\underline{\delta}$

Compute a fair matching M' in G w.r.t. $\overline{\delta}$ and $\underline{\delta}$ and let α denote the minimal selection ratio among all groups % by Algorithm 1
if there exists a matching M s.t. for each group V_i, we have $\max(\lceil \alpha \cdot (\overline{\delta}_i - \underline{\delta}_i) + \underline{\delta}_i \rceil, \widetilde{\min}(V_i)) \leq |M_i|$ % by Algorithm Γ **then**
  Update M to be a maximum matching.
  return the matching M
else
  return NO-instance

Algorithm 2: Checking the existence of a matching satisfying optimality, individual rationality and fairness w.r.t. $\overline{\delta}$ and $\underline{\delta}$

	Identical Bipartite	Identical Non-bipartite	Different Non-bipartite
Maximum	$O(k^2)$	$O(k^2)$	$O(\lvert V\rvert \cdot \lvert E\rvert)$
Minimum	$O(k^{3.5})$	$O(\lvert V\rvert \cdot \lvert E\rvert)$	$O(\lvert V\rvert \cdot \lvert E\rvert)$
Algorithm Γ	$O(k^3)$	$O(k^4 \cdot \log k)$	$O(\sqrt{\lvert V\rvert} \cdot \lvert E\rvert)$

Table 2: Running time upper bounds where k, $|V|$, $|E|$ represent the numbers of groups, vertices and edges in G respectively

compatibility graph structures and summarize all results on running time upper bounds in Table 2. Detailed proofs of all theorems are presented in the Appendix.

Note that given a compatible graph with identical neighbors $G = (V, E)$ with a partition P of vertices V, we can create an equivalent and compact graph $G' = (V', E')$ in which each node $V_i \in V'$ represents a group with a capacity $b(V_i) = |V_i|$, i.e., the size of group V_i. We can construct such a compact representation graph G' in polynomial time and we assume the input for any compatibility graph with identical neighbors is its compact representation. The details of constructing compact graphs are presented in the Appendix.

Bipartite Compatible Graphs with Identical Neighbors

First, we consider the simplest model, bipartite compatibility graphs with identical neighbors in which all groups form a bipartite graph and all vertices within each group has the same neighbors [Ergin *et al.*, 2017].

Theorem 5. *Given a bipartite compatibility graph with identical neighbors, the maximum bounds of all groups can be computed in time $O(k^2)$ where k is the number of groups.*

Proof. (Sketch) Let $N_G(V_i)$ denote all neighboring vertices of some vertex from group V_i in G. For each group V_i, its maximum bound $\max(V_i)$ equals $\min(|V_i|, |N_G(V_i)|)$ and we can calculate the total number of its neighboring vertices in $O(k)$. Thus the total running time is $O(k^2)$. □

Theorem 6. *Given a bipartite compatibility graph with identical neighbors, the minimum bounds of all groups can be computed in time $O(k^{3.5})$ where k is the number of groups.*

Proof. (Sketch) Let M denote a maximum matching in G and let M' denote a maximum matching in the subgraph $G[V \setminus V_i]$ induced from all groups excluding group V_i. For each group V_i, its minimum bound $\min(V_i)$ is $|M| - |M'|$. Note that we can compute a maximum matching in the bipartite compact graph G' of G via the Hopcroft–Karp algorithm in time $O(k^{2.5})$ [Hopcroft and Karp, 1973]. Thus the total running time for computing all minimum bounds is $O(k^{3.5})$. □

Theorem 7. *Given a bipartite compatibility graph G with identical neighbors and a vector of quotas $\lambda = (\lambda_i)_{V_i \in P}$ where each element λ_i corresponds to one group V_i, checking whether there is a matching M in G s.t. $\lambda_i \leq |M_i|$ holds for each group V_i can be done in time $O(k^3)$ where k is the number of groups.*

Proof. (Sketch) We prove Theorem 7 by converting the problem of checking whether there is a matching M in G s.t. $\lambda_i \leq |M_i|$ holds for each group V_i, into an equivalent network flow problem with edge capacities in polynomial time, which then can be solved in time $O(k^3)$ [Malhotra *et al.*, 1978]. □

Non-bipartite Graphs with Identical Neighbors

Next, we consider non-bipartite compatibility graphs with identical neighbors in which all vertices within each group have the same neighbors [Dickerson *et al.*, 2017].

Theorem 8. *Given a non-bipartite compatibility graph with identical neighbors, the maximum bounds of all groups can be computed in time $O(k^2)$ where k is the number of groups.*

Proof. (Sketch) The same proof for Theorem 5 works. □

Theorem 9. *Given a non-bipartite compatibility graph $G = (V, E)$ with identical neighbors, the minimum bounds of all groups can be computed in time $O(|V| \cdot |E|)$ where $|V|$ and $|E|$ denote the numbers of vertices and edges in G.*

Proof. (Sketch) Consider any maximum matching M in G. If there exists an alternating path p that starts from some unmatched vertex $u \in V \setminus V_i$ and ends at some matched vertex $v \in V_i$ w.r.t. M, then update M to be $M \oplus p$ by taking their symmetric difference. Repeat this procedure until there is no such alternating path w.r.t. M, and the minimum bound of group V_i is $|M_i|$. For each group V_i, we can compute its minimum bound in time $O(|V_i| \cdot |E|)$. And the total running time of computing all minimum bounds is $O(|V| \cdot |E|)$. □

Theorem 10. *Given a non-bipartite compatibility graph G with identical neighbors and a vector of quotas $\lambda = (\lambda_i)_{V_i \in P}$ where each element λ_i corresponds to one group V_i, checking whether there is a matching M in G s.t. $\lambda_i \leq |M_i|$ holds for each group V_i can be done in time $O(k^4 \cdot \log k)$ where k is the number of groups.*

Proof. (Sketch) We prove Theorem 10 by converting the problem of checking whether there is a matching M in G s.t. $\lambda_i \leq |M_i|$ holds for each group V_i, into an equivalent b-matching problem in polynomial time, which then can be solved in time $O(k^4 \cdot \log k)$ [Anstee, 1987]. □

Compatible Graphs with Different Neighbors

Next, we consider compatibility graphs with different neighbors where vertices from the same group may have different neighbors [Mattei *et al.*, 2017; Biró *et al.*, 2019].

Theorem 11. *Given a compatibility graph $G = (V, E)$ with different neighbors, the maximum bounds of all groups can be computed in time $O(|V| \cdot |E|)$.*

Proof. (Sketch) We use almost the same proof as the one for Theorem 9. The main difference is that to compute the maximum bound, we keep finding an alternating path p that starts from some unmatched vertex $v \in V_i$ and ends at some matched vertex $u \in V \setminus V_i$ w.r.t. M. Thus they have the same running time $O(|V| \cdot |E|)$. □

Theorem 12. *Given a compatibility graph $G = (V, E)$ with different neighbors, the minimum bounds of all groups can be computed in time $O(|V| \cdot |E|)$.*

Proof. (Sketch) The same proof for Theorem 9 works. □

Theorem 13. *Given a compatibility graph G with different neighbors and a vector of quotas $\lambda = (\lambda_i)_{V_i \in P}$ where each element λ_i corresponds to one group V_i, checking whether there is a matching M in G s.t. $\lambda_i \leq |M_i|$ holds for each group V_i can be done in time $O(\sqrt{|V|} \cdot |E|)$.*

Proof. (Sketch) Create a new graph G^* by extending $G = (V, E)$ as follows: for each group V_i, add a new set of $|V_i| - \lambda_i$ vertices, denoted by V_i'. Each newly added vertex $v_i' \in V_i'$ is incident to all vertices of V_i. And all newly added vertices $\bigcup_{V_i \in P} V_i'$ are incident to each other. If the total number of vertices is odd, then add one more vertex v^* that is incident to all newly added vertices $\bigcup_{V_i \in P} V_i'$. There exists a matching M in G s.t. $\lambda_i \leq |M_i|$ holds for each group V_i if and only if the induced graph G^* has a perfect matching M^*. We can check whether G^* admits a perfect matching by computing a maximum matching in time $O(\sqrt{|V^*|} \cdot |E^*|)$ [Micali and Vazirani, 1980] where $|V^*|$ and $|E^*|$ are linear in the number of $|V|$ and $|E|$. □

Conclusion

In this paper, we study the pairwise organ exchange problem among groups which covers many real-world organ allocation markets. We propose a new fairness concept based on the notion of selection ratio which is able to capture different natural ideas. And we present a general polynomial-time algorithm that computes a matching satisfying three desirable properties including optimality, individual rationality and fairness. We also show several running time upper bounds on computing such matchings for different graph structures.

Acknowledgements

The authors appreciate Haris Aziz and Serge Gaspers for their constructive comments. This work is partially supported by JSPS KAKENHI Grant Numbers JP20H00587. Toby Walsh is supported by the Australian Research Council via a Laureate Fellowship FL200100204.

References

D. J. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 295–304, 2007.

R. P. Anstee. A polynomial algorithm for b-matchings: an alternative approach. *Information Processing Letters*, 24(3):153–157, 1987.

I. Ashlagi, F. Fischer, I. A. Kash, and A. D. Procaccia. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*, 91:284–296, 2015.

C. Berge. Two theorems in graph theory. *Proceedings of the National Academy of Sciences of the United States of America*, 43(9):842, 1957.

D. Bertsimas, V. F. Farias, and N. Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.

P. Biró, W. Kern, D. Pálvölgyi, and D. Paulusma. Generalized matching games for international kidney exchange. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 413–421, 2019.

J. P. Dickerson and T. Sandholm. Multi-organ exchange. *Journal of Artificial Intelligence Research*, 60:639–679, 2017.

J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Price of fairness in kidney exchange. *Transplantation*, 98:815, 2014.

J. P. Dickerson, A. M. Kazachkov, A. D. Procaccia, and T. Sandholm. Small representations of big kidney exchange graphs. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 487–493, 2017.

J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Failure-aware kidney exchange. *Management Science*, 65(4):1768–1791, 2019.

H. Ergin, T. Sönmez, and M. U. Ünver. Dual-donor organ exchange. *Econometrica*, 85(5):1645–1671, 2017.

H. Ergin, T. Sönmez, and M. U. Ünver. Efficient and incentive-compatible liver exchange. *Econometrica*, 88(3):965–1005, 2020.

R. B. Freedman, S. Jana, W. Sinnott-Armstrong, J. P. Dickerson, and V. Conitzer. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence*, page 103261, 2020.

C. Hajaj, J. P. Dickerson, A. Hassidim, T. Sandholm, and D. Sarne. Strategy-proof and efficient kidney exchange using a credit mechanism. In *Twenty-Ninth AAAI conference on artificial intelligence*, pages 921–928, 2015.

J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on computing*, 2(4):225–231, 1973.

V. M. Malhotra, M. P. Kumar, and S. N. Maheshwari. An $O(|V|^3)$ algorithm for finding maximum flows in networks. *Information Processing Letters*, 7(6):277–278, 1978.

N. Mattei, A. Saffidine, and T. Walsh. Mechanisms for online organ matching. In Carles Sierra, editor, *Proceedings of the 29st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 345–351, 2017.

S. Micali and V. V. Vazirani. An $O(\sqrt{|V|}|E|)$ algoithm for finding maximum matching in general graphs. In *21st Annual Symposium on Foundations of Computer Science (sfcs 1980)*, pages 17–27. IEEE, 1980.

A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly journal of economics*, 119(2):457–488, 2004.

A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.

A. E. Roth, T. Sönmez, and M. U. Ünver. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *American Economic Review*, 97(3):828–851, 2007.