

Rejection-proof Kidney Exchange Mechanisms

Blom D.*, Smeulders B., Spieksma F.C.R.

Eindhoven University of Technology

April 29, 2022

Abstract

Kidney exchange programs (KEPs) form an innovative approach to increasing the donor pool through allowing the participation of renal patients together with a willing but incompatible donor. The aim of a KEP is to identify groups of incompatible donor-recipient pairs that could exchange donors leading to feasible transplants. As the size of a kidney exchange grows, a larger proportion of participants can be transplanted. Collaboration between multiple transplant centers, by merging their separate kidney exchange pools is thus desirable. As each transplant center has its own interest to provide the best care to its own patients, collaboration requires balancing individual and common objectives. We consider a class of algorithmic mechanisms for multi-center kidney exchange programs we call *rejection-proof mechanisms*. Such mechanisms propose solutions with the property that no player wishes to unilaterally deviate. We provide a mechanism optimizing social value under this restriction, though the underlying optimization problem is Σ_2^P -hard. We also describe a computationally easier but sub-optimal alternative. Experiments show that rejection-proofness can be achieved at limited cost compared to optimal solutions for regular kidney exchange. Computationally, we provide algorithms to compute optimal rejection-proof solutions for small and medium instance sizes.

1 Introduction

For patients suffering from end-stage renal disease, a kidney transplant from a living donor is the preferred treatment option. A donor must be medically compatible with the recipient, which means some recipients can not receive a transplant, even though they have a donor willing to undergo a transplant for them. We refer to such a combination of recipient and donor as a *pair*. Kidney exchange programs (KEPs), introduced by Rapaport (1986) and popularized by Roth et al. (2004, 2005) aim to match the recipients from a pair with compatible donors from other pairs. Donors in a pair only donate if their paired recipient receives a transplant in return, leading to cycles of transplants. Kidney exchange programs may also include so-called non-directed donors (NDDs). These donors are not associated with a recipient, and are willing to donate to any recipient in the program, leading to donation chains. Typically, kidney exchanges have a policy limiting the maximum length of cycles and chains. Figure 1 illustrates a kidney exchange graph and solution. Kidney exchange programs are widely used, and are still expanding. We refer to Biró et al. (2019) for a recent overview of the practice in Europe.




Figure 1: Circles represent recipient-donor *pairs*, while squares represent *non-directed donors*. The presence of an arc (i, j) signifies compatibility of the donor of i and the recipient of j . The bold arcs represent a feasible solution, with a cycle between pairs 3 and 5 and a chain from NDD 1 to pair 2 and on to pair 4.

*Corresponding author

Kidney exchanges with larger pools of pairs and NDDs generally provide better outcomes for all participants, as the chances of identifying compatible donors increase. Agarwal et al. (2020) show that many hospital and regional kidney exchange programs in the US operate well below an efficient size. Within Europe, several countries have started transnational cooperations (known as *multi-agent kidney exchange programs*) to increase pool size (Böhmig et al., 2017; Valentín et al., 2019). While merging separate kidney exchange programs is an appealing idea, cooperation also faces hurdles. These vary from legal considerations and the alignment of medical practices, but also from the risk of unequal sharing of benefits. In extreme cases, unrestricted cooperation may even lead to a loss of transplants for an individual hospital or country involved in the merged pool, as illustrated in Figure 2 (which is taken from Ashlagi and Roth (2014)). We will use the term *agent* to refer to an entity (hospital or country) that controls a set of pairs and non-directed donors, and we will use the term *collaboration* to refer to a central authority that is able to propose solutions, i.e., transplant plans that contain pairs or non-directed donors from multiple agents.




Figure 2: While the social optimum of this 2-agent KEP instance is given by $c = (b, 1, 2)$, the red agent can choose to withholding both its pairs, thereby gaining one more transplant for its own recipients.

The example in Figure 2 shows that maximizing transplants in merged pools gives incentives for strategic behaviour. Within the literature on kidney exchange, this behaviour concerns the sharing and withholding of information. Typically, agents choose which of their pairs and NDD's to reveal to the collaboration. The collaboration then proposes a transplant plan, matching some of the agent's recipients. The agent can then also select transplants for their unmatched (either hidden, or revealed but not part of the collaboration's transplant plan) pairs (Ashlagi and Roth, 2014). We will refer to this behaviour by the agents by *withholding strategies*. In particular, it can be shown that no socially optimal mechanisms that are strategy-proof exist (Sönmez and Ünver, 2013; Ashlagi and Roth, 2014). Work on cooperation in kidney exchange has focused on mechanisms to incentivize agents to reveal their complete pool, mostly through credit mechanisms in multi-period settings (Hajaj et al. (2015), Biró et al. (2019), Agarwal et al. (2020), Biró et al. (2020b), Biró et al. (2020a), Klimentova et al. (2020)) or randomized mechanisms (Ashlagi et al. (2015)).

In this paper, we consider a different setting, where play use what we will call *rejection strategies*. Here, we assume the collaboration has full information about agent pools, and the collaboration proposes a solution. Next, agents decide whether to reject (part of) this proposed solution. This entails withdrawing pairs or NDDs from cycles and chains from the proposed solution. Afterwards, the agents can internally select new cycles and chains with pairs or NDDs that were either withdrawn or unmatched in the proposed solution. An example is given in Figure 3.




Figure 3: The social optimum (left) consists of two 3-cycles $(a, 1, 2)$ and $(c, 3, 4)$. If this is the proposed solution, the red agent has an incentive to reject the cycle $(c, 3, 4)$ and internally match (b, c) , thereby increasing the number of transplants for the red agent.

Rejection strategies are relevant when the collaboration cannot enforce the proposed solution. This is actually the case in a kidney exchange setting, as binding agreements to execute a proposed transplant

plan are impossible. Indeed, any measure forcing patients and donors to undergo surgery are legally and ethically unacceptable. Consequently, any solution proposed by the collaboration may be impacted by (unforeseen) withdrawals. Moreover, from the point of view of the collaboration, it may be unclear what the precise cause of a withdrawal of a pair of a particular agent is – it is conceivable that the withdrawn pair is used in a transplant that is hidden from the sight of the collaboration perhaps in order to increase the number of transplants for an individual agent. All this motivates a concept that we study in this contribution: the concept of solutions that are rejection-proof (see Section 2 for formal definitions). The goal of this paper is to study mechanisms that lead to solutions such that no agent has an incentive to withdraw a pair or NDD from the proposed solution.

The paper unfolds as follows. In Section 2, we introduce necessary notation, we formally define withholding and rejection strategies, and we introduce the *rejection-proof* property. In Section 3 we experimentally compare withholding strategies and rejection strategies by considering the resulting loss in total transplants for these strategies. Sections 4 and 5 introduce two rejection-proof mechanisms. The first mechanism, Maximum Rejection-proof KEP (MaxRP-KEP), always proposes the maximum value rejection-proof solution. However, we show finding this solution is a Σ_2^P -hard problem and the mechanism thus requires solving a complex bilevel optimization problem. The second mechanism, Maximum Internal Transplants (MaxInt) is inspired by current practice. The mechanism is intuitive, and relatively simple from a computational point of view, but proposes sub-optimal solutions. In Section 6, we describe the outcomes of the two mechanisms. We show that MaxRP-KEP achieves rejection-proofness at a negligible cost in terms of transplants. Furthermore, it turns out that both mechanisms reduce the benefits of withholding strategies as well. We round off the paper with a discussion in Section 7.

2 Preliminaries and Notation

Consider a directed graph $G = (V, A)$ and a set of agents $\mathcal{P}$. The vertex set is partitioned into subsets V^i for each agent $i \in \mathcal{P}$, i.e., $V = \bigcup_{i \in \mathcal{P}} V^i$. Furthermore, for each $i \in \mathcal{P}$, the set V^i is divided into a set of non-directed donors N^i and a set of pairs P^i , i.e. $V^i = N^i \dot{\cup} P^i$. Let $\mathcal{C}$ be the set of cycles which can be used for transplants (typically all cycles of length at most K), and for each $i \in \mathcal{P}$, let $\mathcal{C}^i \subseteq \mathcal{C}$ be the subset of these cycles on the induced subgraph $G^i = G[V^i]$. We refer to $\mathcal{C}^i$ as the set of *internal* cycles of agent i , ($i \in \mathcal{P}$). Furthermore, we have $\mathcal{C}^0 = \mathcal{C} \setminus \bigcup_{i \in \mathcal{P}} \mathcal{C}^i$, which consists of cycles containing pairs of multiple agents, referred to as *shared* cycles. Let w_c be the value of the cycle, and w_c^i the value of the cycle to agent $i \in \mathcal{P}$. For chains, we restrict ourselves to transplants between donors and recipients belonging to the same agent. The donor of the final pair in a chain can, and often does, become a NDD in later matching runs, which leads to additional transplants in the long run. For this reason, we assume agents prefer to keep chains internal, so they have more NDDs in the following matching runs. Let $\mathcal{D}^i$ be the set of all allowed chains of an agent $i \in \mathcal{P}$ (typically all acyclic paths of length at most L in G^i originating in a node $v \in N^i$). We denote by $\mathcal{D} = \bigcup_{i \in \mathcal{P}} \mathcal{D}^i$ the set of allowed chains in G . Again, w_d denotes the value of the chain d , and w_d^i the value to agent i . For the sake of brevity, we will use the term *exchange* to refer to either a cycle or a chain, and define the set of exchanges $\mathcal{E}$ by setting $\mathcal{E} = \mathcal{C} \cup \mathcal{D}$. For each agent $i \in \mathcal{P}$, we use $\mathcal{E}^i = \mathcal{C}^i \cup \mathcal{D}^i$ to denote the set of internal exchanges to agent i and shared exchanges are denoted by $\mathcal{E}^0$. For each exchange $e \in \mathcal{E}$, we denote by $V(e)$ and $A(e)$ its sets of vertices and arcs respectively.

Typically, the values w_c (w_c^i) and w_d (w_d^i) are considered to be the number of transplants in the cycle or chain (to recipients controlled by agent i). However, note that we allow $\sum_{i \in \mathcal{P}} w_c^i \neq w_c$, for example, the participating agents may have agreed that the cooperation must strive for the maximum number of transplants, but privately wish to prioritize certain transplants.

The set of feasible kidney exchange solutions is captured by the following set

$$\mathcal{X} = \{X \subseteq \mathcal{E} : V(e) \cap V(e') = \emptyset \quad \forall e, e' \in X, e \neq e'\}.$$

Note that for each $X \in \mathcal{X}$, the set of cycles and chains forming a solution, the vertex sets of the exchanges must be disjoint. This condition is required, as donors can donate at most one kidney and each pair or NDD can thus be involved in at most one exchange. We denote the value of a solution $X \in \mathcal{X}$ by $w(X) = \sum_{e \in X} w_e$, the weight of all selected cycles and chains. The value of a solution to a particular agent is denoted by $w^i(X) = \sum_{e \in X} w_e^i$. Furthermore, let us denote by $V(X)$ the set of all pairs and NDDs matched with respect to X .

The challenge in multi-agent kidney exchange problems is to settle upon a satisfying way to agree upon the interplay between the agents on the one hand, and the collaboration on the other hand. In other words, how to design a mechanism to arrive at a solution; we define a mechanism as follows:

Definition 1. A kidney exchange mechanism is a function $f : G \rightarrow \mathcal{X}$ which, given a graph $G = (V, A)$ and its partition of V into subsets V^i for each $i \in \mathcal{P}$, produces a kidney exchange solution $X = f(G) \in \mathcal{X}$.

2.1 Withholding strategies

The traditional setting, in which only withholding strategies are used, implies the following game:

1. Each agent $i \in \mathcal{P}$ identifies a subset $S^i \subseteq V^i$ of its pairs and NDDs to contribute to the shared kidney exchange pool. Throughout the rest of the paper, we will refer to this as agent $i \in \mathcal{P}$ selecting the withholding strategy $W^i = V^i \setminus S^i$.
2. A mechanism f computes a solution $X = f(G[\bigcup_{i \in \mathcal{P}} S^i])$. We denote by $U^i(X) \subseteq S^i$ the set of contributed pairs and NDDs of agent $i \in \mathcal{P}$ that remain unmatched with respect to X .
3. Each agent $i \in \mathcal{P}$ selects a maximum weight kidney exchange solution X^i on the subgraph of G^i induced by $W^i \cup U^i(X)$.

For each agent $i \in \mathcal{P}$, the value is given by $w^i(X \cup X^i)$. We note that it is computationally hard for an agent to maximize its value. Indeed, Smeulders et al. (2022) show that the problem of finding an optimal withholding strategy for agent $i \in \mathcal{P}$ is Σ_2^P -hard, even when the mechanism f outputs a social optimal kidney exchange solution and the withholding strategies $W^j \subseteq V^j$ of each agent $j \in \mathcal{P} \setminus \{i\}$ are fixed and known to agent i . This difficulty is amplified by the fact that in practice the vertex sets V^j of other agents are not known to agent i , so they are not even aware of the possible withholding strategies other agents could select.

Due to the complexity of computing optimal withholding strategies, we consider an intuitive withholding strategy proposed in Ashlagi and Roth (2014), which we call the *greedy withholding strategy*. With this strategy, agent $i \in \mathcal{P}$ selects $W^i \subseteq V^i$, with W^i the set of pairs and NDDs matched in a kidney exchange solution of maximum weight to agent i using only the exchanges internal to agent i . The vertices $S^i = V^i \setminus W^i$ are then contributed to the shared pool. In a situation where all other agents submit all their pairs and NDDs to the collaboration, Ashlagi and Roth (2014) show experimentally that the greedy withholding strategy increases, on average, the number of transplants for an individual agent employing this strategy.

2.2 Rejection strategies

In this paper, we consider a setting that requires a mechanism to propose a solution first, after which the individual agents can select a *rejection strategy*. This setting assumes that the entire compatibility graph G is known to the mechanism, and not necessarily to the individual agents.

1. A mechanism f selects a solution $X = f(G)$.
2. Each agent $i \in \mathcal{P}$ selects a rejection strategy, i.e. a kidney exchange solution $X^i \in \mathcal{X}$ such that for all exchanges $e \in X^i$, we have that $V(e) \cap V^i \neq \emptyset$ and $(X^i \cap \mathcal{E}^0) \subseteq (X \cap \mathcal{E}^0)$.
3. A solution $X' \in \mathcal{X}$ is found where $e \in X'$ if and only if $e \in X^i$ for every $i \in \mathcal{P}$ for which $V(e) \cap V^i \neq \emptyset$.

We say that agent $i \in \mathcal{P}$ *accepts* a proposed solution X if agent i selects the rejection strategy $X^i = \{e \in X \mid V(e) \cap V^i \neq \emptyset\}$, i.e., X^i consists of all exchanges in the proposed solution X involving at least one pair or NDD of V^i . If this is not the case, then we say that agent i *rejects* the solution.

The goal of an agent $i \in \mathcal{P}$ is to select a rejection strategy X^i that has maximum value to agent i itself. We call this problem the Reject Kidney Exchange Problem (RKEP). A mixed-integer programming formulation $\text{RKEP}(X, i)$ of this optimization problem is described in (1). It considers a binary variable x_e^i for each exchange $e \in \mathcal{E}$, which has value 1 if and only if e is selected in the rejection strategy X^i . The agent is free to choose any set of pairwise disjoint internal exchanges ($\mathcal{E}^i$), while Constraints (1c) imply that shared exchanges $e \in \mathcal{E}^0$ can only be included if the exchange is contained in the proposed solution X .

$$\text{RKEP}(X, i) := \text{maximize} \quad \sum_{e \in \mathcal{E}} w_e^i x_e^i \quad (1a)$$

$$\text{subject to} \quad \sum_{e \in \mathcal{E}: v \in V(e)} x_e^i \leq 1 \quad \forall v \in V \quad (1b)$$

$$x_e^i = 0 \quad \forall e \in \mathcal{E}^0 \setminus X \quad (1c)$$

$$x_e^i \in \{0, 1\} \quad \forall e \in \mathcal{E} \quad (1d)$$

Notice that for any agent $i \in \mathcal{P}$ for which $\text{RKEP}(X, i) \leq w^i(X)$, it is a weakly dominant strategy to accept the proposed solution, provided that the other agents also accept the solution. If all agents accept the solution X , each agent $i \in \mathcal{P}$ scores a value of $w^i(X)$. On the other hand, if for some agent $i \in \mathcal{P}$, $\text{RKEP}(X, i) > w^i(X)$, agent i can achieve a superior outcome by rejecting the proposed solution. An optimal rejection strategy is then encoded by the x_e^i variables. Remark however that a rejection strategy X^i does not guarantee that agent i receives a payout of $w^i(X^i)$, as there might exist a shared exchange $e \in \mathcal{E}^0 \cap X$ not contained in the rejection strategy of another agent j having a vertex involved in e .

2.3 Rejection-proofness

In this work, we will consider mechanisms that propose solutions X , for which accepting is a weakly dominant strategy for every agent, assuming other agents also accept. This gives rise to the definition of *rejection-proofness*.

Definition 2 (Rejection-proof). *A mechanism $f : G \rightarrow \mathcal{X}$ is rejection-proof if, for any compatibility graph $G = (V, A)$ and each agent $i \in \mathcal{P}$, we have*

$$\text{RKEP}(f(G), i) \leq w^i(f(G)) \quad (2)$$

Informally, we will also say that an individual solution X is rejection-proof if and only if $\text{RKEP}(X, i) \leq w^i(X)$ for all agents $i \in \mathcal{P}$. Note that in effect, any proposed solution X gives rise to a game between the agents. X determines the set of rejection strategies for each agent, and $w^i(X)$ defines the value to agent i . Observe that by definition, for a rejection-proof mechanism, it holds that the profile of rejection strategies $(X^1, \dots, X^{|\mathcal{P}|})$ given by $X^i = X \cap \{e \mid V(e) \cap V^i \neq \emptyset\}$, i.e., each agent accepts the solution X , is a pure Nash equilibrium of this game.

3 Experimentally comparing withholding and rejection strategies

In the previous section, we have formally defined both withholding and rejection strategies, as well as the rejection-proofness property. In this section, we experimentally show the benefits agents can achieve through using these strategies, as well as the loss in social value incurred by this strategic behaviour. This will show that incentives for strategic behaviour are present in realistic settings, and that the resulting overall loss is large enough to motivate the development of mechanisms preventing strategic behaviour.

We use the same experimental set-up as Ashlagi and Roth (2014). These authors evaluated the gains agents could achieve through the greedy withholding mechanism. In this section, we extend those simulations to also evaluate rejection strategies.

We consider KEP instances generated through two sets of characteristic distributions for donor-recipient pairs. One set is generated using the traditional Saidman et al. (2006) distributions (also used in Ashlagi and Roth). The second set is constructed based on recent work, which aims to more closely resemble instances of the UK's KEP, as described by Delorme et al. (2022). Specifically, we make use of the SplitPRA-BandXMatch-PRA0 distributions described in this paper. Both sets of instances were generated using the generator found at <https://wpetersson.github.io/kidney-webapp/#/generator>, making use of the pre-built configurations. We consider the same instance sizes as Ashlagi and Roth. Specifically, we create instances with an equal number of pairs for all agents, either 10, 15, 30, 50 or 80. The number of agents is always even and chosen such that there are at most 180 total pairs. We assume agents value every transplant equally, and that there are no non-directed donors.

Results of the experiments are described in Table 1. This table shows the average gain for an individual agent employing the greedy withholding or optimal rejection strategy, assuming all other

agents do not employ such strategies. We note that for both instance types, withholding and rejection strategies provide similar benefits for agents with small pools. For larger agents, with pool sizes of 30 and above, the average gain for the rejection strategy is consistently larger than for the withholding strategy. Furthermore, we note that even though in all but one case, the average gain of withholding was positive, nearly every instance set contains instances where withholding led to transplant losses (only DelBXMPRA0.10x8 and DelBXMPRA0.10x14 did not). By contrast, rejection strategies are risk-free to the agent. If an agent fails to identify a rejection strategy with which the payout can be improved over accepting, the agent can keep the proposed solution. The dual result of higher average profits and reduced risk makes rejection strategies a superior choice for agents.

Instance set	Delorme SplitPRA-BXM-PRA0				Saidman			
	WA	RA	WT	RT	WA	RA	WT	RT
10x10	1.005	1.005	0.943	0.986	1.003	1.034	0.868	0.931
10x12	1.012	1.008	0.931	0.985	1.040	1.037	0.890	0.929
10x14	1.019	1.012	0.944	0.985	1.021	1.050	0.900	0.925
10x16	1.019	1.004	0.938	0.988	1.031	1.046	0.899	0.921
10x2	1.000	1.000	0.935	1.000	0.992	1.042	0.855	0.981
10x4	1.045	1.055	0.882	0.962	1.038	1.064	0.846	0.943
10x6	1.014	1.014	0.926	0.986	1.068	1.065	0.844	0.929
10x8	1.006	1.006	0.927	0.980	1.066	1.066	0.869	0.907
15x10	1.015	1.010	0.905	0.977	1.051	1.060	0.886	0.914
15x2	1.000	1.008	0.939	0.991	1.028	1.077	0.862	0.961
15x4	1.009	1.017	0.873	0.971	1.076	1.100	0.846	0.915
15x6	1.046	1.025	0.877	0.969	1.047	1.087	0.860	0.915
15x8	1.024	1.038	0.907	0.975	1.020	1.047	0.867	0.909
30x2	1.004	1.055	0.854	0.981	1.040	1.080	0.901	0.959
30x4	0.992	1.039	0.813	0.955	1.044	1.077	0.871	0.912
30x6	1.012	1.036	0.839	0.958	1.052	1.088	0.881	0.894
50x2	1.003	1.074	0.820	0.966	1.032	1.055	0.928	0.950
80x2	0.984	1.059	0.807	0.952	1.033	1.058	0.958	0.955

Table 1: Comparison of withholding and rejection strategies. **Instance Set:** The first number of the instance set shows the number of pairs each agent has, the second shows the number of players. **Withhold Agent (WA):** average gain or loss in number of transplants for a specific agent from switching unilaterally from no withholding to greedy withholding, assuming all other agents do not withhold. **Reject Agent (RA):** similar, but with rejection strategies. **Withhold Total (WT):** average gain or loss in the total number of transplants when comparing full cooperation of all agents to greedy withholding strategies for all agents. **Reject Total (RT):** similar, but with rejection strategies for all agents.

The table also shows the loss of transplants overall, if all agents employ the relevant strategy. In our experiments, rejection strategies can lower the total number of transplants by up to 10%. While this is lower than the losses caused by the greedy withholding strategy, they are not negligible. As a result, we conclude that mechanisms should also address the risk of rejection strategies.

4 Maximum rejection-proof kidney exchange solutions

In this section, we will formulate a mechanism that proposes a rejection-proof solution of maximum weight, dubbed the *Maximum Rejection-Proof KEP* mechanism or simply MaxRP-KEP. We also propose an algorithm for solving the resulting bilevel optimization problem.

Definition 3. *Maximum Rejection-Proof KEP (MaxRP-KEP)*

The MaxRP-KEP mechanism selects a kidney exchange solution $X = (x_e)_{e \in \mathcal{E}}$ optimal to the following bilevel program.

$$\text{maximize } \sum_{e \in \mathcal{E}} w_e x_e \quad (3a)$$

$$\text{subject to } \sum_{e \in \mathcal{E}: v \in V(e)} x_e \leq 1 \quad \forall v \in V \quad (3b)$$

$$\sum_{e \in \mathcal{E}} w_e^i x_e \geq \text{RKEP}(X, i) \quad \forall i \in \mathcal{P} \quad (3c)$$

$$x_e \in \{0, 1\} \quad \forall e \in \mathcal{E} \quad (3d)$$

If we restrict ourselves to the objective (3a) and Constraints (3b), we have a standard weighted kidney exchange problem. Constraints (3c) enforce that no agent has a strict incentive to reject the solution. Adding these constraints seriously complicates the problem, as the right-hand is the value to an agent for their optimal rejection strategy. Indeed, Model (3) is actually a bilevel program. We show that for cases in which standard kidney exchange problems are NP-hard, the problem of finding maximum cardinality rejection-proof kidney exchange solutions is complete in the complexity class Σ_2^P (see Woeginger (2021) for more background).

Theorem 4. *Given a set of agents $\mathcal{P}$, a graph $G = (V, A)$ with a partition of V into subsets V^i for each agent $i \in \mathcal{P}$, a fixed maximum cycle length $K \geq 3$, any maximum chain length $L \in \mathbb{N}$ and an integer t , deciding whether there exists a rejection-proof solution $X \in \mathcal{X}$ with value at least t is Σ_2^P -complete.*

The proof of this theorem can be found in the Appendix. Solving Σ_2^P -hard problems is challenging, and requires specialized algorithms. In the remainder of this section, we discuss strategies to solve Model (3). In Section 4.1, we propose a single-level reformulation of the problem with exponentially many constraints. Next, Section 4.2 discusses an iterative constraint generation procedure.

4.1 A single-level reformulation

In this section, we identify a set of constraints which can replace (3c) in Model (3), thus reducing it to a single-level optimization problem, albeit of size exponential in the size of the input. First, for each agent $i \in \mathcal{P}$ and each subset $U^i \subseteq V^i$, we define

$$\beta(U^i) = \max \left\{ \sum_{e \in X} w_e^i \mid X \in \mathcal{X} : \forall e \in X, V(e) \subseteq U^i \right\}. \quad (4)$$

$\beta(U^i)$ denotes the maximum value to agent i of feasible kidney exchange solutions on $G[U^i]$. Note that calculating $\beta(U^i)$ requires solving a standard kidney exchange problem, a NP-hard problem. We now claim the following:

Theorem 5. *A solution $X = (x_e)_{e \in \mathcal{E}}$ to MaxRP-KEP is rejection-proof if and only if for each agent $i \in \mathcal{P}$ and each set $U^i \subseteq V^i$, we have*

$$\sum_{e \in \mathcal{E} : V(e) \cap U^i \neq \emptyset} w_e^i x_e \geq \beta(U^i). \quad (5)$$

Proof. $\Rightarrow$ Suppose there exists some agent $i \in \mathcal{P}$ and a subset $U^i \subseteq V^i$ for which the solution X does not satisfy (5). Then agent i can improve their solution by rejecting all exchanges containing a node $v \in U^i$ and replacing them by a kidney exchange solution on $G[U^i]$ with maximum value to the agent. Hence X is not rejection-proof.

$\Leftarrow$ Suppose the solution $X = (x_e)_{e \in \mathcal{E}}$ is not rejection-proof. Then there exists an agent $i \in \mathcal{P}$ and a solution $X^i = (x_e^i)_{e \in \mathcal{E}} \subseteq X \cup \mathcal{E}^i$ feasible to (1) such that

$$\sum_{e \in \mathcal{E}} w_e^i x_e^i > \sum_{e \in \mathcal{E}} w_e^i x_e. \quad (6)$$

We now construct the set U^i as follows:

$$U^i := \{v \in V^i \mid \exists e \in \mathcal{E} : v \in V(e), x_e^i = 1, x_e = 0\}. \quad (7)$$

Thus, the set U_i is the set of vertices that are part of an exchange in X_i , and not part of an exchange in X . By construction of the set U^i , we have that $x_e^i \leq x_e$ if $V(e) \cap U^i = \emptyset$. Thus, we have:

$$\sum_{e \in \mathcal{E} : V(e) \cap U^i = \emptyset} w_e^i x_e^i \leq \sum_{e \in \mathcal{E} : V(e) \cap U^i = \emptyset} w_e^i x_e. \quad (8)$$

Combining (6) and (8) gives

$$\sum_{e \in \mathcal{E} : V(e) \cap U^i \neq \emptyset} w_e^i x_e^i > \sum_{e \in \mathcal{E} : V(e) \cap U^i \neq \emptyset} w_e^i x_e. \quad (9)$$

Finally, note that by construction of U^i , if for an exchange $e \in X^i$ it holds that $V(e) \cap U^i \neq \emptyset$, then actually $V(e) \subseteq U^i$. These exchanges thus describe a feasible KEP solution on $G[U^i]$. We now have

$$\beta(U^i) \geq \sum_{e \in \mathcal{E}: V(e) \cap U^i \neq \emptyset} w_e^i x_e^i. \quad (10)$$

From combining (9) and (10), we can conclude that (5) is violated for the set U^i . $\square$

We can now replace Constraints (3c) by Constraints (5), for each set $U^i \subseteq V^i$ for each agent $i \in \mathcal{P}$. We refer to these constraints as the *subset rejection constraints*. Violation of any such constraint means the agent has incentive to reject all exchanges making use of vertices in this subset and replacing it with a maximum agent value KEP solution. By this replacement, we obtain the following formulation.

$$\text{maximize } \sum_{e \in \mathcal{E}} w_e x_e \quad (11a)$$

$$\text{subject to } \sum_{e \in \mathcal{E}: v \in V(e)} x_e \leq 1 \quad \forall v \in V \quad (11b)$$

$$\sum_{e \in \mathcal{E}: V(e) \cap U^i \neq \emptyset} w_e^i x_e \geq \beta(U^i) \quad \forall i \in \mathcal{P}, U^i \subseteq V^i \quad (11c)$$

$$x_e \in \{0, 1\} \quad \forall e \in \mathcal{E} \quad (11d)$$

As $\beta(U^i)$ is independent of the value of other variables, this is a single-level reformulation of the problem. However, the formulation is exponential in size due to the number of subsets $U^i \subseteq V^i$. Due to Σ_2^P -completeness, we cannot hope for a subexponential size formulation unless the polynomial hierarchy collapses.

4.2 Solving MaxRP-KEP

In this section, we discuss our solution approach for MaxRP-KEP. The size of this model is exponential in the number of vertices of the graph, since for every agent i , it requires one constraint for each subset $U^i \subseteq V^i$. In our solution method, we start with a relaxed version of this model, dubbed MaxRP-KEP-R, using only a small collection of the subset rejection constraints (5). This gives rise to an intuitive row generation algorithm, in which iteratively MaxRP-KEP-R is solved to optimality and violated subset rejection constraints are identified for each agent. We will discuss the identification of optimal solutions to MaxRP-KEP-R and the separation of violated subset rejection constraints in some more detail.

4.2.1 Identifying optimal solutions to MaxRP-KEP-R.

It has been observed that kidney exchange problems usually have many optimal solutions (Farnadi et al., 2021). This carries over to MaxRP-KEP-R, in particular when the number of rejection constraints is still small, as the problem is then almost similar to a regular KEP. Selecting a solution among the different optima is a delicate matter, as some solutions may be rejection-proof, while others are not. It is desirable to identify a rejection-proof optimal solution if it exists, as then we can restrict the number of iterations of the (costly) procedure of separating violated subset rejection constraints and re-solving of the resulting, strengthened, MaxRP-KEP-R. The following lemma gives us insight in how to compare solutions in terms of rejection-proofness.

Lemma 6. *Consider two solutions $X, Y \in \mathcal{X}$ such that*

- (a) $(Y \cap \mathcal{E}^0) \subseteq (X \cap \mathcal{E}^0)$,
- (b) $w^i(Y) \geq w^i(X)$ for all agents $i \in \mathcal{P}$.

If X is rejection-proof, then Y is rejection-proof.

Proof. Suppose on the contrary that Y is not rejection-proof. Then, for some agent $i \in \mathcal{P}$, there exists a solution Y^i optimal to $\text{RKEP}(Y, i)$ such that $w^i(Y^i) > w^i(Y)$. Note that any solution feasible to $\text{RKEP}(Y, i)$ is also feasible to $\text{RKEP}(X, i)$ because of Condition (a). Finally, $w^i(Y^i) > w^i(Y) \geq w^i(X)$ means the agent has incentive to reject X as well, meaning that X is also not rejection-proof. $\square$

This lemma shows that the number of shared exchanges in a solution and agent value are two factors influencing whether or not a solution is rejection-proof. Indeed, when the number of shared exchanges in a solution X is small, this implies that the number of feasible rejection strategies, i.e. the number of feasible solutions of $\text{RKEP}(X, i)$, is more restricted compared to when there are many shared exchanges. This feasible region is also more restricted if the agent value $w^i(X)$ of i for solution X is larger, which is consistent with Constraints (3c).

We use this knowledge to propose as a tiebreaker the internal preference objective

$$t_1(X) := \sum_{i \in \mathcal{P}} \sum_{e \in \mathcal{E}^i} w_e^i x_e. \quad (12)$$

By counting only the agent value of internal exchanges, it immediately takes into account both conditions of Lemma 6. Clearly, this tiebreaker prioritizes solutions with high agent values, whereas solutions with many shared exchanges are disadvantaged. In fact, we can show that given two solutions as in Lemma 6, the tiebreaking rule will always value the solution more likely to be rejection-proof more highly.

Lemma 7. *Consider two solutions, X and Y , such that $(Y \cap \mathcal{E}^0) \subseteq (X \cap \mathcal{E}^0)$ and $w^i(Y) \geq w^i(X)$ for all agents $i \in \mathcal{P}$. Then, we have $t_1(Y) \geq t_1(X)$, i.e., the tiebreak value of Y is at least as high as the tiebreak value of X .*

Proof. Let $(Y \cap \mathcal{E}^0) \subseteq (X \cap \mathcal{E}^0)$. Then, using the definitions of $w^i(Y)$ and $w^i(X)$, we get

$$\sum_{e \in \mathcal{E}^0} w_e^i y_e \leq \sum_{e \in \mathcal{E}^0} w_e^i x_e \Rightarrow \sum_{e \in \mathcal{E}^i} w_e^i y_e \geq \sum_{e \in \mathcal{E}^i} w_e^i x_e. \quad (13)$$

Now, it follows from (12) and (13) that $t_1(Y) \geq t_1(X)$. $\square$

Notice that there is a strong connection between this tiebreaking objective and solutions obtained from the MaxInt mechanism. In fact, the same two objectives are considered, namely maximizing total number of transplants and maximizing the number of internal transplants, but in a different order. As MaxInt proposes rejection-proof KEP solutions, this supports the use of this tiebreaker.

The tiebreaker is implemented as follows. Let Z be an upper bound to the agent value of internal exchanges in any solution ($Z \geq \max_{X \in \mathcal{X}} \sum_{i \in \mathcal{P}} \sum_{e \in \mathcal{E}^i} w_e^i x_e$). We then replace the objective (3a) of MaxRP-KEP, by

$$\sum_{e \in \mathcal{E}} Z w_e x_e + \sum_{i \in \mathcal{P}} \sum_{e \in \mathcal{E}^i} w_e^i x_e. \quad (14)$$

This objective takes into account both the social value of all transplants ($\sum_{e \in \mathcal{E}} Z w_e x_e$) and the agent value of internal exchanges ($\sum_{e \in \mathcal{E}^i} w_e^i x_e$). Due to the heavy weighting of social value, the agent value will only function as a tiebreaker between solutions with equal social value.

4.2.2 Separating violated subset rejection constraints

Consider a non-rejection-proof solution X . In the proof of Theorem 5, we show that if there exists an agent $i \in \mathcal{P}$ and solution X^i feasible to $\text{RKEP}(X, i)$ with $w^i(X^i) > w^i(X)$, at least one subset rejection constraint must be violated. In the proof, we showed the constraint for agent i and subset

$$U^i = \{v \in V^i \mid \exists e \in \mathcal{E} : v \in V(e), x_e^i = 1, x_e = 0\}.$$

is then violated.

The right-hand side of the violated rejection constraint can be computed as follows.

$$\beta(U^i) = \sum_{e \in \mathcal{E} : V(e) \subseteq U^i} w_e^i x_e^i. \quad (15)$$

5 Maximum Weight Internal Matching

In this section, we describe a first and simple rejection-proof mechanism, dubbed the *Maximum Weight Internal KEP* mechanism, shortened to the MaxInt mechanism or simply MaxInt. This mechanism proposes a solution where the total transplant value is maximized, under the restriction that the value of internal transplants is maximum.

Definition 8. *Maximum Weight Internal Transplants (MaxInt)*

The MaxInt mechanism selects a KEP solution $X \in \mathcal{X}$ optimizing the following integer program (MaxInt-KEP).

$$\begin{aligned}
& \text{maximize} && \sum_{e \in \mathcal{E}} w_e x_e && (16a) \\
& \text{subject to} && \sum_{e \in \mathcal{E}: v \in V(e)} x_e \leq 1 && \forall v \in V && (16b) \\
& && \sum_{e \in \mathcal{E}^i} w_e^i x_e = \beta(V^i) && \forall i \in \mathcal{P} && (16c) \\
& && x_e \in \{0, 1\} && \forall e \in \mathcal{E} && (16d)
\end{aligned}$$

Constraints (16b) again impose that X induces a set of pairwise vertex-disjoint exchanges. Constraints (16c) ensure that for each individual agent $i \in \mathcal{P}$, the exchanges of X induced by V^i forms a maximum weight internal KEP solution. Notice that these constraints are satisfied whenever every agent selects a greedy withholding strategy. Nevertheless, an agent might have many distinct greedy withholding strategies, each of which might have a different impact on the number of transplants that can be realized overall. Therefore, MaxInt can thus be seen as mechanism that coordinates a profile of greedy withholding strategies allowing the maximum overall value of transplants in the shared kidney exchange pool.

Theorem 9. *The MaxInt mechanism is rejection-proof.*

Proof. Let $X \in \mathcal{X}$ be a KEP solution proposed by MaxInt, encoded by binary variables $(x_e)_{e \in \mathcal{E}}$, and $i \in \mathcal{P}$ an arbitrary agent. Then

$$w^i(X) = \sum_{e \in \mathcal{E}^i} w_e^i x_e + \sum_{e \in \mathcal{E}^0} w_e^i x_e.$$

For any optimal solution X^i to RKEP(X, i), encoded by binary variables $(x_e^i)_{e \in \mathcal{E}}$, we have

$$\text{RKEP}(X, i) = \sum_{e \in \mathcal{E}^i} w_e^i x_e^i + \sum_{e \in \mathcal{E}^0} w_e^i x_e^i.$$

As $X \cap \mathcal{E}^i$ is a maximum weight internal set of exchanges on V^i , we have

$$\sum_{e \in \mathcal{E}^i} w_e^i x_e \geq \sum_{e \in \mathcal{E}^i} w_e^i x_e^i. \quad (17)$$

Next, notice that a shared exchange e can only be contained in X^i if it was in the proposed solution X in the first place, hence

$$\sum_{e \in \mathcal{E}^0} w_e^i x_e \geq \sum_{e \in \mathcal{E}^0} w_e^i x_e^i. \quad (18)$$

It follows for adding (17) and (18) that $w^i(X) \geq \text{RKEP}(X, i)$, showing MaxInt is rejection-proof. $\square$

6 Computational Results

This section gives the outcomes of both rejection-proof mechanisms, with regards to total transplants. We consider the same instances and experimental setting as discussed in Section 3. We are now in a position to evaluate the outcomes of the MaxInt and MaxRP-KEP mechanisms, as described in Sections 4 and 5. First, we will compare the number of transplants resulting from the mechanisms against the social optimum, assuming agents cooperate fully. Second, we will look at the incentives for agents. Clearly, when rejection-proof mechanisms are employed, rejection strategies provide no benefit to agents. We will therefore only check whether withholding strategies still provide a benefit to agents.

Table 2 shows that the MaxRP-KEP mechanism achieves rejection-proofness at a small cost on average. For all instance sets, the number of transplants lost, compared to the social optimum is at most 1%. These losses are consistently less than those suffered if only a single agent employed either rejection or withholding strategies (Table 1). The MaxInt mechanism comes at a higher cost, with losses up to

Instance set	Delorme SplitPRA-BXM-PRA0			Saidman		
	MaxRP-KEP	MaxInt	All Withhold	MaxRP-KEP	MaxInt	All Withhold
10x10	0.994	0.944	0.943	0.999	0.885	0.868
10x12	0.996	0.934	0.931	1.000	0.911	0.890
10x14	0.996	0.944	0.944	1.000	0.922	0.900
10x16	0.997	0.940	0.938	0.999	0.920	0.899
10x2	1.000	0.935	0.935	0.998	0.891	0.855
10x4	0.991	0.882	0.882	0.994	0.866	0.846
10x6	0.996	0.931	0.926	0.996	0.868	0.844
10x8	0.994	0.928	0.927	0.995	0.885	0.869
15x10	0.993	0.908	0.905	1.000	0.910	0.886
15x2	0.996	0.947	0.939	0.996	0.888	0.862
15x4	0.990	0.879	0.873	0.996	0.871	0.846
15x6	0.996	0.892	0.877	0.998	0.888	0.860
15x8	0.996	0.910	0.907	0.999	0.891	0.867
30x2	0.992	0.885	0.854	0.997	0.915	0.901
30x4	0.990	0.843	0.813	0.999	0.896	0.871
30x6	0.995	0.866	0.839	1.000	0.912	0.881
50x2	0.992	0.858	0.820	0.998	0.940	0.928
80x2	0.992	0.845	0.807	1.000	0.959	0.958

Table 2: The columns “MaxRP-KEP” and “MaxInt” show the average fraction of the social optimum (maximum number of transplants provided no rejection / withholding strategies are used) achieved by MaxRP-KEP and MaxInt respectively, over all 50 instances per Delorme and Saidman instance set. The column “All Withhold shows” the average fraction of the social optimum whenever each agent uses the greedy withholding strategy.

15% of total transplants. This is still an improvement over the losses suffered in case all agents employ greedy withholding strategies, which are shown in the final column of the table. Finally, Table A.2 in the appendix shows the outcome agents can achieve using greedy withholding strategies if the rejection-proof mechanisms are employed. Contrary to the situation where a maximum weight solution on the common pool is selected (Table 1), agents generally lose transplants by withholding. This is especially the case for agents with large pools, where average losses can be up 8 % for the MaxRP-KEP mechanism.

7 Discussion

In this paper, we address the issue of cooperation in kidney exchange from a new angle. Previous work focused on incentivizing agents to contribute all of their patients and donors to a common pool. The underlying assumption of this work is that agents will strategically choose to reveal or withhold their pairs from a cooperation, so as to maximize the value of transplants to their patients (withholding strategies). One of our contributions is to recognize a second option for agents, the rejection strategies. In Section 3, we showed that such strategies provide better outcomes for the agent than existing withholding strategies with reasonable computational and information requirements. Furthermore, we showed such rejection strategies cause an overall loss of transplants. As a result, we propose that rejection-proofness is a valuable property for cooperation mechanisms.

We provide two different rejection-proof mechanisms, MaxInt and MaxRP-KEP. MaxInt is simple to implement and easy to understand, which may carry benefits in practice. Currently, international cooperations exist (Biró et al., 2020b) using a *consecutive pool* concept. Countries first match internally and send all unmatched pairs to a common pool. This corresponds to agreeing that every country will use the greedy withholding strategy. The MaxInt mechanism is very similar in spirit, but coordinates the greedy withholding strategies so as to maximize the total number of transplants. The second mechanism, MaxRP-KEP, can provide significantly higher numbers of transplants. This performance does come at a cost with regards to implementation and computational difficulty. Further research into algorithms for solving this problem is required to allow for cooperations with many agents with large pools.

Common to both mechanisms is that besides their designed outcome of nullifying rejection strategies, they also both eliminate expected benefits provided by greedy withholding strategies. We do note that they are not strategy-proof, as evidenced by the existence of instances where the greedy withholding strategy does provide a benefit.

As a final remark about the benefits of our mechanisms, we note that they provide a natural way of reconciling differences in transplant valuation in cooperations. These mechanisms make use of agent (w_e^i) and cooperation (w_e) valuations. Through the constraints ensuring rejection-proofness, agents are

assured that their private valuations will be taken into account. This allows participants in a cooperation to set their own priorities.

Looking forward, we believe combinations of rejection-proof and credit mechanisms are interesting to study. Current credit mechanisms do not handle failure or rejection of proposed transplants. By removing incentives to reject this issue is ameliorated. Rejection-proof mechanisms provide guarantees that agents will not lose out by cooperating, but can still allow for unequal division of the benefits of cooperation. Furthermore, in our experiments rejection-proof mechanisms still allow situations where greedy withholding strategies are profitable. Depending on graph structure, there may be situations where they are consistently profitable. Credit mechanisms can reduce these issues in multi-period settings. By combining rejection-proof and credit mechanisms, the risk of strategic behaviour can be further reduced.

Acknowledgements

The authors thank Péter Bíró for helpful discussions. The research of Frits Spieksma is supported by Dutch Research Council (NWO) Gravitation Project NETWORKS [Grant 024.002.003].

References

- Nikhil Agarwal, Itai Ashlagi, Eduardo Azevedo, Clayton R Featherstone, and Ömer Karaduman. Market failure in kidney exchange. *American Economic Review*, 109(11):4026–4070, 2020.
- Sanjeev Arora and Boaz Barak. *Computational complexity: A Modern Approach*. Cambridge University Press, 2009.
- Itai Ashlagi and Alvin E Roth. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics*, 9(3):817–863, 2014.
- Itai Ashlagi, Felix Fischer, Ian A Kash, and Ariel D Procaccia. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*, 91:284–296, 2015.
- Péter Bíró, Bernadette Haase-Kromwijk, Tommy Andersson, Eyjólfur Ingi Ásgeirsson, Tatiana Baltessová, Ioannis Boletis, Catarina Bolotinha, Gregor Bond, Georg Böhmig, Lisa Burnapp, et al. Building kidney exchange programmes in Europe—an overview of exchange practice and activities. *Transplantation*, 103(7):1514, 2019.
- Péter Bíró, Márton Gyetvai, Xenia Klimentova, João Pedro Pedroso, William Pettersson, and Ana Viana. Compensation scheme with shapley value for multi-country kidney exchange programmes. 2020a.
- Péter Bíró, Joris van de Klundert, David Manlove, William Pettersson, Tommy Andersson, Lisa Burnapp, Pavel Chromy, Pablo Delgado, Piotr Dworczak, Bernadette Haase, et al. Modelling and optimisation in European kidney exchange programmes. *European Journal of Operational Research*, 291(2):447–456, 2020b.
- Georg A Böhmig, Jiří Fronek, Antonij Slavcev, Gottfried F Fischer, Gabriela Berlakovich, and Ondrej Viklicky. Czech-austrian kidney paired donation: first european cross-border living donor kidney exchange. *Transplant International*, 30(6):638–639, 2017.
- M Delorme, S Garcia, J Gondzio, J Kalcsics, DF Manlove, W Pettersson, and J Trimble. A new matheuristic and improved instance generation for kidney exchange programmes. *Computers & Operations Research*, 141, 2022.
- Golnoosh Farnadi, William St-Arnaud, Behrouz Babaki, and Margarida Carvalho. Individual fairness in kidney exchange programs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35 (13), pages 11496–11505, 2021.
- Chen Hajaj, John P Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. Strategy-proof and efficient kidney exchange using a credit mechanism. In *Twenty-Ninth AAAI conference on artificial intelligence*, 2015.

- Xenia Klimentova, Ana Viana, João Pedro Pedroso, and Nicolau Santos. Fairness models for multi-agent kidney exchange programmes. *Omega*, page 102333, 2020.
- FT Rapaport. The case for a living emotionally related international kidney donor exchange registry. *Transplantation proceedings*, 18(3) Suppl. 2):5–9, June 1986. ISSN 0041-1345. URL <http://europepmc.org/abstract/MED/11649919>.
- Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Kidney exchange. *The Quarterly journal of economics*, 119(2):457–488, 2004.
- Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005.
- Susan L Saidman, Alvin E Roth, Tayfun Sönmez, M Utku Ünver, and Francis L Delmonico. Increasing the opportunity of live kidney donation by matching for two-and three-way exchanges. *Transplantation*, 81(5):773–782, 2006.
- Bart Smeulders, Danny A M P Blom, and Frits C R Spijksma. Identifying optimal strategies in kidney exchange games is Σ_2^P -complete. *Mathematical Programming*, pages 1–22, 2022.
- Tayfun Sönmez and M Utku Ünver. Market design for kidney exchange. *The Handbook of Market Design*, 93137, 2013.
- María Oliva Valentín, Marta Garcia, Alessandro Nanni Costa, Catarina Bolotinha, Lluís Guirado, Fabio Vistoli, Alberto Breda, Pamela Fiaschetti, and Beatriz Dominguez-Gil. International cooperation for kidney exchange success. *Transplantation*, 103(6):e180–e181, 2019.
- Gerhard J Woeginger. The trouble with the second quantifier. *4OR*, pages 1–25, 2021.

Appendices

A Supporting tables


Instance Set	Total Time (s)	Constraints Added	Iterations	Instances Solved
DelBXMPRA0Score_96_4_1l2m	0.46	5.50	4.77	30
DelBXMPRA0Score_96_4_1l5s	0.17	3.37	3.80	30
DelBXMPRA0Score_96_4_2equal	1.15	9.67	7.63	30
DelBXMPRA0Score_96_4_4equal	0.06	3.13	2.70	30
DelBXMPRA0_96_4_1l2m	0.09	1.60	2.47	30
DelBXMPRA0_96_4_1l5s	0.08	1.07	1.93	30
DelBXMPRA0_96_4_2equal	0.09	1.40	2.17	30
DelBXMPRA0_96_4_4equal	0.04	1.20	1.83	30
SaidmanScore_96_4_1l2m	308.04	36.42	22.83	24
SaidmanScore_96_4_1l5s	303.07	33.26	28.67	27
SaidmanScore_96_4_2equal	732.42	41.83	28.04	24
SaidmanScore_96_4_4equal	518.13	26.64	13.89	28
Saidman_96_4_1l2m	2.97	3.40	3.37	30
Saidman_96_4_1l5s	5.75	7.80	7.40	30
Saidman_96_4_2equal	4.70	2.50	2.77	30
Saidman_96_4_4equal	2.50	9.07	6.87	30

Table A.1: Comparison of computational results for all instance sets. All results were obtained using no initial constraints, the weighted tiebreaking and maximum violation constraints added. Instances solved is the amount of instances solved to optimality in less than 3 hours (10 800s), out of 30. Computation time, constraints added and number of iterations are the means of instances solved to optimality within this time.


Instance Set	Delorme SplitPRA-BXM-PRA0		Saidman	
	WA MaxRP-KEP	WA MaxInt	WA MaxRP-KEP	WA MaxInt
10x10	1.005	1.000	1.009	0.987
10x12	1.000	1.000	1.000	0.979
10x14	1.008	1.004	1.008	0.985
10x16	1.020	1.000	0.997	0.991
10x2	1.000	1.000	0.996	0.985
10x4	0.991	1.011	0.987	0.996
10x6	1.007	1.000	0.976	1.004
10x8	1.006	1.000	0.991	1.000
15x10	1.008	1.000	0.989	0.976
15x2	1.000	1.000	0.972	0.990
15x4	0.986	1.005	0.998	0.993
15x6	0.996	0.980	1.002	0.976
15x8	0.994	0.994	0.979	0.993
30x2	0.970	0.976	0.984	0.991
30x4	0.964	0.992	0.985	0.987
30x6	0.988	0.982	0.987	0.983
50x2	0.945	0.974	0.981	0.991
80x2	0.917	0.971	0.994	1.000

Table A.2: This table shows the number of transplants achieved by an agent employing the greedy withholding strategy (WA: Withhold Agent), divided by the number of transplants without withholding, for both rejection-proof mechanisms and over all 50 instances in the set.


B Supporting figures




(a)



(b)



(c)



(d)

Figure B.1: Comparisons of solution times dependant on use of the tiebreak rule.

C Proof of Theorem 4

This theorem considers the computational complexity of the optimization problem faced by the MaxRP-KEP mechanism. We focus ourselves on the case of unit transplant values, i.e. $w_e = |V(e) \cap (\bigcup_{i \in \mathcal{P}} P^i)|$ and $w_e^i = |V(e) \cap P^i|$ for all exchanges $e \in \mathcal{E}$. We denote by DecRP-KEP the decision variant of this problem, which is described as follows.

DecRP-KEP

Instance: a compatibility graph $G = (V, A)$ and a set of agents $\mathcal{P}$, such that $V = \bigcup_{i \in \mathcal{P}} V_i$ consists of the pools $V^i = P^i \cup N^i$ of donor-recipient pairs and non-directed donors for agent $i \in \mathcal{P}$, and A the set of potential transplants. Furthermore, we are given a maximum cycle length K , maximum chain length L , and a nonnegative integer t .

Question: does there exist a rejection-proof KEP solution consisting of vertex-disjoint cycles of length at most K and chains of length at most L , covering at least t vertices?

Theorem 4 claims that for $K \geq 3$ fixed, and any $L \in \mathbb{N}$, DecRP-KEP belongs to the class Σ_2^P , and furthermore that it is a complete with respect to this class. We prove this statement by providing a reduction from Adversarial (2,2)-SAT, which is proven to be Σ_2^P -complete in Smeulders et al. (2022). This problem is closely related to the foundational Σ_2^P -complete problem, B_2^{CNF} .

Problem: Adversarial (2,2)-SAT

Instance: Boolean CNF formula φ consisting of two sets X, Y of variables, and a set C of disjunctive clauses over X and Y , such that each variable $z \in X \cup Y$ occurs twice in negated ($\neg z$), twice in unnegated form (z) in C .

Question: does there exist a truth setting $\theta_X: X \rightarrow \{\text{True}, \text{False}\}$ such that for any truth setting $\theta_Y: Y \rightarrow \{\text{True}, \text{False}\}$, the formula φ is not satisfied.

We will say that a variable $z \in X \cup Y$ occurs in sense $\psi = t$ or $\psi = f$ in clause $c \in C$, if c contains the literal z or $\neg z$ respectively. Furthermore, we denote by $c(z, i)$ and $c(\neg z, i)$ the clause in which z or $\neg z$ respectively occur for the i -th time, where $i \in \{1, 2\}$. In the following sections, we prove that polynomial reducibility of Adversarial (2,2)-SAT to DecRP-KEP. First, we describe a mapping f of an Adversarial (2,2)-SAT instance φ to an instance $f(\varphi)$ of DecRP-KEP. Next, we lay out specific properties of optimal and feasible solutions of the instances $f(\varphi)$. Finally, we bring all of this together in the polynomial reduction.

Instance mapping

Let us be given an instance $\varphi = \varphi(X, Y, C)$ to Adversarial (2,2)-SAT. We construct an instance $f(\varphi) = (G, \mathcal{P}, K, L, t)$ to DecRP-KEP forming a 2-agent kidney exchange program, i.e., $\mathcal{P} = \{b, g\}$. Informally, we will refer to these agents as the blue and the green agent. The figures of the instance use these same colours to differentiate vertices of the different agents.

X-gadgets: for each variable $x \in X$, we define the vertex set $V_x := V_x^g \cup V_x^b$, with agent's vertex sets V_x^g and V_x^b of agent g and b respectively, as follows.

$$V_x^g := \{x_0\} \cup \{x_{\psi,i} : \psi \in \{t, f\}, i \in \{1, 2\}\}$$

and

$$V_x^b := \{\gamma_{x,\psi,i} : \psi \in \{t, f\}, i \in \{1, 2\}\} \cup \{\alpha_{x,i} : i \in \{1, 2\}\}$$

The arc set $A_x \subseteq \{(i, j) : i, j \in V_x\}$ is defined as

$$A_x := \{(x_0, x_{\psi,1}), (x_{\psi,1}, x_{\psi,2}), (x_{\psi,2}, x_0) : \psi \in \{t, f\}\} \\ \cup \{(x_{\psi,i}, \alpha_{x,i}), (\alpha_{x,i}, \gamma_{x,\psi,i}), (\gamma_{x,\psi,i}, x_{\psi,i}) : \psi \in \{t, f\}, i \in \{1, 2\}\}$$

The graph $G_x = (V_x, A_x)$ will be referred to as the x -*gadget* or *variable gadget* of x , and is depicted within the black square in Figure C.1a.

Y -gadgets: for each variable $y \in Y$, we obtain the vertex set $V_y := V_y^g \cup V_y^b$ of agent b and g respectively. Now, $V_y^g = \emptyset$ and

$$V_y^b := \{y_0\} \cup \left\{ y_{\psi,i} : \psi \in \{t, f\}, i \in \{1, 2\} \right\} \\ \cup \left\{ \gamma_{y,\psi,i} : \psi \in \{t, f\}, i \in \{1, 2\} \right\} \cup \left\{ \alpha_{y,i} : i \in \{1, 2\} \right\}$$

Similar to the arc set A_x for each x -gadget, the arc set $A_y \subseteq \{(i, j) : i, j \in V_y\}$ is defined as

$$A_y := \left\{ (y_0, y_{\psi,1}), (y_{\psi,1}, y_{\psi,2}), (y_{\psi,2}, y_0) : \psi \in \{t, f\} \right\} \\ \cup \left\{ (y_{\psi,i}, \alpha_{y,i}), (\alpha_{y,i}, \gamma_{y,\psi,i}), (\gamma_{y,\psi,i}, y_{\psi,i}) : \psi \in \{t, f\}, i \in \{1, 2\} \right\}$$

The graph $G_y = (V_y, A_y)$ will be referred to as the y -*gadget* or *variable gadget* of y , and is illustrated within the black square in Figure C.1b.

Clause gadgets: for each clause $c \in C$, we define the vertex set $W_c := W_c^g \cup W_c^b$, where the vertex sets W_c^g and W_c^b are respectively defined by

$$W_c^g := \{c_i^g : i \in \{1, 2, 3\}\},$$

and

$$W_c^b := \{\delta_c\} \cup \{c_i^b : i \in \{1, 2, 3, 4\}\}.$$

The arc set A_c within the c -gadget is defined as

$$A_c := \left\{ (c_1^g, c_2^g), (c_2^g, c_3^g), (c_3^g, c_1^g) \right\} \cup \left\{ (\delta_c, c_1^g), (c_1^g, \delta_c) \right\} \\ \left\{ (c_2^g, c_1^b), (c_1^b, c_2^b), (c_2^b, c_2^g) \right\} \cup \left\{ (c_3^g, c_3^b), (c_3^b, c_4^b), (c_4^b, c_3^g) \right\}$$

The graph $G_c = (W_c, A_c)$ will be referred to as the c -*gadget* or the *clause gadget* of c . The c -gadget is illustrated in Figure C.1c.

Arcs connecting variable and clause gadgets: finally, we also add arcs connecting x - and y -gadgets to c -gadgets. In particular, we define these arcs as

$$A_{\gamma,\delta} = \left\{ (\gamma_{z,t,i}, \delta_{c(z,i)}), (\delta_{c(z,i)}, \gamma_{z,t,i}) : z \in X \cup Y, i \in \{1, 2\} \right\} \\ \cup \left\{ (\gamma_{z,f,i}, \delta_{c(\neg z,i)}), (\delta_{c(\neg z,i)}, \gamma_{z,f,i}) : z \in X \cup Y, i \in \{1, 2\} \right\}$$

These arcs are depicted in Figures C.1a and C.1b as dashed arcs. The arcs $A_{\gamma,\delta}$ induce a set of 2-cycles, which we will refer to as (γ, δ) -*cycles*. Summarizing the above, the instance $f(\varphi)$ to MaxRP-KEP consists of a graph $G = (V, A)$ with

$$V = \left(\bigcup_{x \in X} V_x \right) \cup \left(\bigcup_{y \in Y} V_y \right) \cup \left(\bigcup_{c \in C} W_c \right),$$

and

$$A = \left(\bigcup_{x \in X} A_x \right) \cup \left(\bigcup_{y \in Y} A_y \right) \cup \left(\bigcup_{c \in C} A_c \right) \cup A_{\gamma,\delta}.$$

Furthermore, we set the parameters $K = 3$, $L \in \mathbb{N}$ arbitrarily, and $t = 9|X| + 9|Y| + 5|C| + 3$. We proceed by providing the relevant cycle and chain packings, i.e. feasible KEP solutions, for our hardness reduction. Given this characterization of the DecRP-KEP instance $f(\varphi)$, we will now identify some structured packings related to the gadgets and prove that these structures need to appear in optimal solutions.

Notice that we show the result here for maximum cycle length $K = 3$, but all arguments used in the proof below will stay in place for all fixed $K \geq 4$ through the following adaptation of the variable gadgets (for all $x \in X$ and $y \in Y$): subdivide each arc $(\gamma_{z,\psi,i}, z_{\psi,i})$ for $\psi \in \{t, f\}, i \in \{1, 2\}$ into $K - 2$ arcs, where all intermediate nodes belong to the vertex set V_z^b of the blue agent.




Figure C.1: Gadgets for the polynomial reduction defining the transformed instance $f(\varphi)$.

Relevant gadget packings

Consider an x -gadget, and let $\theta_X: X \rightarrow \{\text{True}, \text{False}\}$ be a truth setting for X . We refer to the packing defined by

$$\begin{cases} S_{x,t} = \left\{ \{x_{t,1}, x_{t,2}, x_0\}, \{\gamma_{x,f,1}, x_{f,1}, \alpha_{x,1}\}, \{\gamma_{x,f,2}, x_{f,2}, \alpha_{x,2}\} \right\}, & \text{if } \theta_X(x) = \text{True}, \\ S_{x,f} = \left\{ \{x_{f,1}, x_{f,2}, \gamma_x\}, \{\gamma_{x,t,1}, x_{t,1}, \alpha_{x,1}\}, \{\gamma_{x,t,2}, x_{t,2}, \alpha_{x,2}\} \right\}, & \text{if } \theta_X(x) = \text{False}. \end{cases}$$

as being *consistent* with setting x to **True** or **False** respectively. We can define similar notions of $S_{y,t}$ and $S_{y,f}$ for y -gadgets for $y \in Y$, through replacing X with Y and x with y .

Figure C.2 illustrates the gadget packings consistent with **True** and **False** respectively. We show that for each rejection-proof solution of maximum size, i.e., a solution covering the largest number of vertices, we need to pack each variable gadget consistent with a truth setting.

Lemma 10. *For each maximum solution to MaxRP-KEP on $f(\varphi)$, each variable gadget (i.e., for all $x \in X$ and all $y \in Y$) is packed consistent with either **True** or **False**.*

Proof. We distinguish the proof for x - and y -gadgets, since the x -gadget contains pairs from the green agent, whereas the y -gadget does not.

x-gadget: By using the result of Theorem 5 for the green agent's subset $U^g = \{x_{t,1}, x_{t,2}, x_0\}$, observe that in any rejection-proof solution, the set of cycles and chains containing vertices from U^g should cover at least 3 green vertices. This can only be realized through selecting either the cycle $c_{x,t} = \{x_{t,1}, x_{t,2}, x_0\}$ or the cycle $c_{x,f} = \{x_{f,1}, x_{f,2}, x_0\}$. Without loss of generality, suppose we select $c_{x,t}$. Let now S be a maximum rejection-proof solution, and suppose its sub-packing $S[G_x]$ induced by the x -gadget is not $S_{x,t}$. Again without loss of generality, suppose the cycle $\{x_{f,1}, \alpha_{x,1}, \gamma_{x,f,1}\}$ is not contained in $S[G_x]$. As S is maximum, this means that the cycle $\{\gamma_{x,f,1}, \delta_{c(\neg x,1)}\} \in S$ (if not, then all vertices of the cycle $\{x_{f,1}, \alpha_{x,1}, \gamma_{x,f,1}\}$ are uncovered in S). However, we could then switch from $\{\gamma_{x,f,1}, \delta_{c(\neg x,1)}\}$ to $\{x_{f,1}, \alpha_{x,1}, \gamma_{x,f,1}\}$ without violating subset rejection constraints to obtain rejection-proof solution S' with $|S'| = |S| + 1$, thus contradicting that S is of maximum size.

y-gadget: Notice that any feasible rejection-proof solution S has the following property: for each $y \in Y$, the sub-packing induced by $S[G_y]$ is maximal with respect to G_y . (If not, there exists a 3-cycle $\tilde{c}$ whose vertex set $U^b = V(\tilde{c})$ defines a violated subset rejection constraint). This can only be attained in four ways: the first two correspond to the gadget packings $S_{y,t}$ and $S_{y,f}$ consistent with **True** and **False**. The other alternative maximal packings on G_y are




Figure C.2: Gadget packings for variable $x \in X$ consistent with **True** (C.2a) and **False** (C.2b). Note that the (γ, δ) cycles are not required for a packing to be consistent. These are coloured red to highlight the possibility of including them in a packing, if the δ -vertices are not covered by another cycle already. Consistent packings work similarly for variables $y \in Y$, where the gadget consists of blue vertices only.

$$S_3 = \left\{ \{\gamma_{y,f,1}, y_{f,1}, \alpha_{y,1}\}, \{\gamma_{y,t,2}, y_{t,2}, \alpha_{y,2}\} \right\} \text{ and } S_4 = \left\{ \{\gamma_{y,t,1}, y_{t,1}, \alpha_{y,1}\}, \{\gamma_{y,f,2}, y_{f,2}, \alpha_{y,2}\} \right\}.$$

Without loss of generality, suppose the y -gadget is packed according to S_3 . Then the subset rejection constraint is violated for $U^b = \{y_{t,1}, y_{t,2}, y_0, \alpha_{y,2}, y_{f,2}, \gamma_{y,f,2}\}$, since $\beta(U^b) = 6 > 5$, which is the maximum number of vertices that could be covered in S with cycles having nonempty vertex intersection with U^b (see Figure C.3 for an illustration). This shows that in a feasible rejection-proof solution, all y -gadgets must be packed consistent with either **True** or **False**.

Since both cases cover all of the variable gadgets, the statement follows. $\square$

Notice now that whenever we select a variable gadget packing consistent with a truth setting ψ , then the vertices $\gamma_{z,\psi,i}, i = 1, 2$ are not covered by a cycle internal to the gadget. Nevertheless, we can use the corresponding cycles (γ, δ) -cycles to cover these vertices. In that case, the δ -vertex should not yet be covered by another (γ, δ) -cycle. Observe that using a (γ, δ) -cycle corresponding to a variable $z \in X \cup Y$ and clause $c \in C$ corresponds to satisfying the clause c as the variable z occurs in sense ψ in clause c .

For each clause $c \in C$, there are two possible packings for the clause gadget of c based on whether or not the vertex δ_c is covered by a (γ, δ) -cycle, in other words, whether or not clause c is satisfied. If yes, we select the packing $S_{c,t} = \left\{ \{c_1^g, c_2^g, c_3^g\} \right\}$, given by the unique green 3-cycle. Otherwise, the packing $S_{c,f} = \left\{ \{\delta_c, c_1^g\}, \{c_2^g, c_1^b, c_2^b\}, \{c_3^g, c_3^b, c_4^b\} \right\}$ is selected. We refer to $S_{c,t}$ and $S_{c,f}$ as the *satisfied clause packing* and the *unsatisfied clause packing* respectively. Both packings are depicted in Figure C.4.

Lemma 11. *For each feasible solution to MaxRP-KEP on $f(\varphi)$, each satisfiability gadget is packed with either $S_{c,t}$ or $S_{c,f}$.*

Proof. We again use the result of Theorem 5 to show the statement. Consider the subset $U^g = W_c^g = \{c_1^g, c_2^g, c_3^g\}$ for some clause $c \in C$. Any packing in the clause gadget of c should cover at least all vertices of U^g , and the only two packings with this property are given by $S_{c,t}$ and $S_{c,f}$. $\square$




Figure C.3: When a y -gadget has packing S_3 , the vertices inside the brown parallelogram form a subset U^b with a violated subset rejection constraint.




Figure C.4: Illustration of the satisfied clause packing and the unsatisfied clause packing. As rejection-proofness requires all three green vertices are covered, these are the only possible packings in a feasible solution.

The following lemma illustrates the link between variable and clause gadgets. Whenever the variable gadgets correspond to a truth setting satisfying a clause, then there will be a (γ, δ) cycle selected, linking the clause gadget to a variable gadget. A satisfied clause gadget packing will result.

Lemma 12. *Let the variable gadgets be packed consistent with a truth setting θ_X, θ_Y . In every feasible rejection-proof solution to $f(\varphi)$, for each clause $c \in C$ satisfied by this truth setting, δ_c is covered by a (γ, δ) -cycle.*

Proof. Let $c \in C$ be a clause satisfied with respect to truth settings θ_X, θ_Y . Without loss of generality, let $x \in X$ such that c contains the first occurrence of the literal x and $\theta_X(x) = \text{True}$. Then, consider the blue agent's subset $U^b = \{\gamma_{x,t,1}, \delta_c\}$. If δ_c is not covered by a (γ, δ) -cycle, then U^b is completely uncovered while it induces a 2-cycle internal to blue. Hence, if δ_c is not covered by a (γ, δ) -cycle, then the overall packing would not be rejection-proof, leading to a contradiction. $\square$

Given a feasible solution to MaxRP-KEP on the instance $f(\varphi)$ in which every variable gadget is packed consistent with a truth setting, we define $\theta_X: X \rightarrow \{\text{True}, \text{False}\}$ as the truth setting of X such that the variable gadget of $x \in X$ is packed by $S_{x, \theta_X(x)}$. Furthermore, let $k_{\max}(\theta_X)$ be the maximum

number of clauses in C that can be satisfied given the truth setting θ_X for X , over all truth settings for Y .

Lemma 13. *Let S be a feasible solution to MaxRP-KEP on $f(\varphi)$ such that θ_X is the associated truth setting for X . If $k_{\max}(\theta_X) = |C|$, then the packing of each y -gadget in S is $S_{y, \theta_Y(y)}$ for some truth setting θ_Y for Y that satisfying all clauses.*

Proof. Suppose by contradiction that the packing of y -gadgets in S is not equal to $S_{y, \theta_Y(y)}$ for some satisfying truth setting θ_Y for Y given θ_X , i.e. using θ_X and θ_Y together results in $k < k_{\max}(\theta_X) = |C|$ satisfied clauses. We will show the blue agent has incentive to deviate. We consider the subset $U^b \subset V^b$ of pairs of the blue agent defined by

$$U^b := \left(\bigcup_{y \in Y} V_y^b \right) \cup \{ \delta_c \mid c \in C \} \cup \{ \gamma_{x, \psi, i} \mid x \in X, \psi \in \{t, f\}: \theta_X(x) = \psi \}$$

Currently, it holds for solution S that

$$\sum_{e \in S: V(e) \cap U^b \neq \emptyset} w_e^b = \sum_{e \in S: V(e) \subseteq U^b} w_e^b = 9|Y| + |C| + k.$$

Furthermore, since $k_{\max}(\theta_X) = |C|$, we have that $\beta(U^b) = 9|Y| + 2|C| > 9|Y| + |C| + k$, as each variable gadget allows for an internal packing covering 9 vertices, and satisfiability of all clauses implies that we can pick $|C|$ disjoint (γ, δ) -cycles internal to U^b by construction of the gadgets and the structure of U^b . Hence, the subset rejection constraint for agent b and subset U^b is violated at solution S , which contradicts that S is a rejection-proof solution. $\square$

We use Lemma 10 to 13 to prove our main complexity result, Theorem 4.

Theorem. *DecRP-KEP is Σ_2^P -complete, for fixed maximum cycle length $K \geq 3$ and arbitrary maximum chain length $L \in \mathbb{N}$.*

Proof. We start by showing membership of DecRP-KEP in Σ_2^P . Following the definition of Σ_2^P in Arora and Barak (2009), we consider DecRP-KEP as the language

$$\text{DecRP-KEP} = \{ \langle G, \mathcal{P}, K, L, t \rangle : \exists \text{ rejection-proof KEP solution consisting of cycles of length } \leq K \text{ and chains of length } \leq L \text{ covering at least } t \text{ vertices.} \}$$

Remark that this language is in Σ_2^P since a tuple $\langle G, \mathcal{P}, K, L, t \rangle \in \text{DecRP-KEP}$ if and only if

$\exists X \in \mathcal{X} : X$ covers at least t vertices.

$\forall X' \in \mathcal{X} : X'$ is either (i) not feasible to RKEP(X, i) or (ii) has value $w^i(X') \leq w^i(X)$ for all $i \in \mathcal{P}$.

Remark that there exist compact encodings of X, X' , i.e. polynomial in number of vertices and arcs of the graph G , such as arc-based formulations of kidney exchange programs. Furthermore, it can be checked in polynomial time whether such encodings indeed imply feasible cycle-chain packings. This also holds for checking if X' is feasible to RKEP(X, i) and checking its corresponding objective value $w^i(X')$. Therefore, given a rejection-proof solutions X and some alternative vector X' , we can check in polynomial time that X' does not prove violation of rejection-proofness of X , hence the language DecRP-KEP is in Σ_2^P .

In order to prove Σ_2^P -hardness, we use the results of Lemma 10 to 13. The reduction is based on proving the following statement: an instance φ to Adversarial (2,2)-SAT returns YES if and only if the transformed instance $f(\varphi)$ to DecRP-KEP returns YES. Recall that we use the parameter settings $\mathcal{P} = \{b, g\}$, $K = 3$, $L \in \mathbb{N}$, and $t = 9|X| + 9|Y| + 5|C| + 3$.

$\Rightarrow$) Let us be given a YES-instance $\varphi = \varphi(X, Y, C)$ to Adversarial (2,2)-SAT, i.e. there exists a truth setting $\theta_X : X \rightarrow \{\text{True}, \text{False}\}$ such that for any truth setting $\theta_Y : Y \rightarrow \{\text{True}, \text{False}\}$, there exists some clause $c \in C$ that is left unsatisfied. Let us now fix such a truth setting θ_X for X . Let $k_{\max}(\theta_X) \leq |C| - 1$ be the maximum number of clauses of C that can be satisfied given this θ_X , and let θ_Y be a truth setting attaining $k_{\max}(\theta_X)$ satisfied clauses. We construct a solution S to the corresponding DecRP-KEP instance $f(\varphi)$ as follows:

- for each variable $x \in X$, we add to S the packing $S_{x, \theta_X(x)}$ of the x -gadget consistent with $\theta_X(x)$.

- similarly, for each variable $y \in Y$, we add to S the packing $S_{y, \theta_Y(y)}$ of the y -gadget consistent with $\theta_Y(y)$
- for each clause $c \in C$, we add to S the satisfied clause packing $S_{c,t}$ if c is satisfied by the truth settings θ_X and θ_Y , and the unsatisfied clause packing $S_{c,f}$ otherwise.
- for each (γ, δ) -cycle, add it to S if both its vertices are not covered yet by S .

Within each X - and Y -gadget, the solution S covers nine vertices, which gives us $9|X| + 9|Y|$ vertices already. Furthermore, since the corresponding truth settings θ_X and θ_Y for X and Y respectively ensure that $k_{\max}(\theta_X)$ clauses were satisfied, this gives us another $k_{\max}(\theta_X)$ cycles of length 2, so an additional $2k_{\max}(\theta_X)$ vertices are covered. Finally, if the c -vertex was not in a 2-cycle already, we cover eight vertices internally in the satisfiability gadget, opposed to 3 otherwise. In total, this gives us

$$9|X| + 9|Y| + 2k_{\max}(\theta_X) + 8(|C| - k_{\max}(\theta_X)) + 3k_{\max}(\theta_X) = 9|X| + 9|Y| + 8|C| - 3k_{\max}(\theta_X)$$

vertices covered by S . This is indeed at least $9|X| + 9|Y| + 5|C| + 3$, since $k_{\max}(\theta_X) \leq |C| - 1$. It remains to show that this solution is indeed rejection-proof. The transformed instance $f(\varphi)$ of DecRP-KEP contains a green and blue agent. Observe that S covers all green vertices, so we can restrict ourselves to proving that the blue agent has no incentive to reject S . First, notice that the blue agent cannot use an internal exchange to cover the vertices $c_i^b, i \in \{1, 2, 3, 4\}$ in the clause gadgets. Therefore, such vertices can only be covered if the proposed solution includes its associated shared cycle. Furthermore, proposed shared cycles within each gadget of variable $z \in X \cup Y$ of the form $\{z_{\psi,i}, \alpha_{z,i}, \gamma_{z,\psi,i}\}$, where $\psi \in \{t, f\}$ and $i \in \{1, 2\}$, will also not be rejected by the blue agent. Notice that the only reason to reject such a cycle is to be able to pick the (γ, δ) -cycle incident with $\gamma_{z,\psi,i}$. However, there is no strict benefit to the blue agent to make this switch. Furthermore, this blocks the δ -vertex from forming a 2-cycle with another blue γ -vertex from another variable gadget. Therefore, accepting the cycle always leads to weakly preferred outcomes to the blue agent. The remaining set of the blue agent's vertices is U^b as defined in the proof of Lemma 13. Notice that in S , all δ_c -vertices for $c \in C$ are covered. Furthermore, we have maximum size packings within x - and y -gadgets that allow for the largest number, i.e. $k_{\max}(\theta_X)$, of (γ, δ) -cycles, with respect to the set of solutions that packs all x - and y -gadgets consistent with a truth setting. This means that a total number of $9|Y| + |C| + k_{\max}(\theta_X)$ vertices in U^b are covered by a cycle in S . Hence, a certificate solution S' showing that the blue agent has incentive to reject S needs to contain more (γ, δ) -cycles than S . Recall that S corresponds to a truth setting for Y that maximizes the number of satisfied clauses. Thus, there must exist a variable $y \in Y$ such that S' contains a (γ, δ) -cycle associated with occurrences of both the unnegated and negated literal of y . Whenever this happens, the packing of the y -gadget is not consistent anymore and a 3-cycle within the y -gadget needs to be given up, since this would mean a γ_y -vertex is covered with two distinct cycles. Therefore, a solution containing $k - k_{\max}(\theta_X)$ extra (γ, δ) -cycles comes at the cost of giving up equally many 3-cycles within the y -gadgets, hence covering fewer vertices of U^b in total compared to S . This shows that S is rejection-proof and satisfies the target t , i.e. S is a certificate solution proving that $f(\varphi)$ is a YES-instance to DecRP-KEP.

$\Leftarrow$) Let us consider a feasible solution S to DecRP-KEP for the transformed instance that covers at least $9|X| + 9|Y| + 5|C| + 3$ vertices, meaning that S is in fact of maximum size. Remark that there must exist a clause $c \in C$ for which the c -gadget is packed with $S_{c,f}$, since otherwise a feasible DecRP-KEP solution can have at most $9|X| + 9|Y| + 2|C| + 3|C| = 9|X| + 9|Y| + 5|C|$ covered vertices otherwise. Let the truth setting $\theta_X: X \rightarrow \{\text{True}, \text{False}\}$ such that the variable gadgets for each $x \in X$ are packed in S by $S_{x, \theta_X(x)}$. This follows from Lemma 10, which assures us that each variable gadget packing is consistent with a truth value. Since at least one unsatisfied clause packing is selected and S is rejection-proof, Lemma 13 implies that there cannot exist a truth setting to Y that satisfies all clauses given θ_X . Therefore, θ_X is a valid certificate for $\varphi = \varphi(X, Y, C)$ being a YES-instance to Adversarial (2,2)-SAT. $\square$

D Computation times

D.1 Difficulty of instance sets

While the basic characteristics of the instances such as number of pairs and NDDs were kept constant, we note massive differences in computation time required for the different instance sets. These differences

were most pronounced for the different distributions used for instance generation (Saidman vs. SplitPRA-BandXMatch-PRA0), with the Saidman instances considerably more difficult. Agent preferences also had a large impact, as instances with unit agent valuations for both the overall objective and agent preferences were significantly easier to solve. Finally, more minor differences exist depending on the distribution of pairs and NDDs over the agents.

- **Saidman vs. Delorme.** It is clear that there exist huge differences based on the instance generation procedure used. All instance sets generated using the Delorme distributions are solved completely, with average computation times only exceeding 1 second for one set of instances. Instances generated using the Saidman distributions on the other hand are much more difficult, with not all instances solved within the time limit. For the solved instances, the average time to optimality is also much longer, up 12 minutes for one instance set.

A likely factor for this difference is the far greater arc density in Saidman instances (see Delorme et al. (2022)) For one, this increases the computational cost of solving the integer programs. Second, this gives the agents more freedom and they are thus more likely to reject proposed solutions, as reflected in the number of iterations required to find the optimal rejection-proof solution.

- **Unit vs. Score** Instances in which the agents have non-unit values for the transplants are much more difficult. This is very apparent in the Saidman instances, where all instances with unit valuations are solved within the time limit, with the most difficult instances set taking less than 6 seconds on average. For the non-unit values, instances in each set were unsolved, and average times for solved instances range from 5 to 12 minutes.

The reason for the additional difficulty is reflected in the number of iterations and number of constraints. In case of unit values, the overall value and agent values of a given solution are closely aligned. For non-unit values however, solutions with many low value transplants are often rejected by the agents, adding more constraints to the master problem. This leads to more, and more costly, iterations as the additional constraints add difficulty to the master problem.

- **Distribution over agents** We focus on the difficult Saidman Score instances, and note instances with larger agents are more difficult to solve. The “4equal” set has most instances solved, and while the average time per instance is second highest, this is due to the final solved instances taking just under 3 hours. Excluding this instance, the average stands at only 163 seconds. For all other instances, the largest agent has 50 pairs and NDDs, and we see that more instances are solved, with lower average time, for the instances where the remaining pairs are divided over more agents.

Since larger agents have more opportunities to identify profitable rejections, the increased difficulty is as expected. This is also reflected in the number of iterations required and constraints added. In instance sets with larger agents, more constraints were needed to identify rejection-proof solutions.

The instance sets in this study cannot capture the variety of possibilities possible in real life, and the feasibility of using MaxRP-KEP in practice will depend on the specific properties of the application. This said, we believe the above results are positive. MaxRP-KEP is solvable on most medium-size instances even for the dense Saidman instances, and trivial for the more realistic Delorme instances. The increased difficulty as agent size increases is a concern, especially in cooperations between large countries. Still, size 50 agents are manageable, a size which is only exceeded by 3 European Kidney Exchanges (UK, Spain and The Netherlands (Biró et al., 2019)).

D.2 Tiebreaker

We now discuss the tiebreaker, and its effect on computation time for the challenging SaidmanScore instances. Its effect depends on instance characteristics. Figure B.1 shows the number of solved instances at each point in time and gives the clearest comparison. For the instances with one large and 2 medium or 5 small agents, the tiebreaker generally speeds up computation for easier instances, but the gap disappears for harder instances. For both instance sets, not using the tiebreaker in fact allows for one more instance to be solved within the time limit. For the 2equal and 4equal instance sets, little difference can be seen. We do note that using the tiebreaker strongly reduces the number of iterations required to identify optimal solutions, as can be seen in Table D.3. As before, this is offset by the increased difficulty of solving the master problem, now due to the harder objective function, leading to similar overall results.

Instance Set	tiebreak	Total Time (s)	Master Time (s)	Iterations	Instances Solved
112m	Used	308.04	298.82	22.83	24
	Not Used	942.85	911.45	44.48	25
115s	Used	303.07	290.88	28.67	27
	Not Used	726.00	672.74	43.71	28
2equal	Used	732.42	720.24	28.04	24
	Not Used	1176.52	1156.51	39.17	24
4equal	Used	518.13	516.44	13.89	28
	Not Used	378.73	376.41	18.71	28

Table D.3: Comparison of SaidmanScore instances based on whether or not tiebreaks are used. Results were obtained using no initial constraints and maximum violated constraints added. Instances solved is the number of instances solved to optimality within 3 hours (10800s), out of 30. Reported computation and master times, and numbers of iterations are averages for instances solved to optimality within this time.