Inter-Club Kidney Exchange

Gabriele Farina[*] and **John P. Dickerson**[+] and **Tuomas Sandholm**[*]
Carnegie Mellon University[*] University of Maryland[+]
gfarina@cs.cmu.edu, john@cs.umd.edu, sandholm@cs.cmu.edu

Abstract

A kidney exchange is a centrally-administered barter market where patients swap their willing yet incompatible donors. Modern kidney exchanges use 2-cycles, 3-cycles, and chains initiated by non-directed donors (altruists who are willing to give a kidney to anyone) as the means for swapping.

We propose significant generalizations to kidney exchange. We allow more than one donor to donate in exchange for their desired patient receiving a kidney. We also allow for the possibility of a donor willing to donate if any of a number of patients receive kidneys. Furthermore, we combine these notions and generalize them. The generalization is to exchange among organ *clubs*, where a club is willing to donate organs outside the club if and only if the club receives organs from outside the club according to given specifications.

Forms of organ clubs already exist—under an arrangement where one gets to be in the club as a potential recipient if one is willing to donate one's organs to the club upon death. Our approach can be used as an inter-club exchange mechanism that increases systemwide good (and can also be applied to live donation). In this paper we introduce these ideas, present the notion of operation frames that can be used to sequence the operations across batches, and present integer programming formulations for the market clearing problems for these new types of organ exchanges.

Introduction

Chronic kidney disease is a condition that causes a reduction of the kidney function, often with life-threatening consequences. Its societal burden is likened to that of diabetes (Neuen et al. 2013). The National Institute of Diabetes and Digestive and Kidney Diseases estimates that one in 10 American adults, more than 20 million, have some level of chronic kidney disease (NIH 2011).

Kidney transplantation is the most effective treatment for kidney failure. However, the demand for donor kidneys far exceeds the supply. The United Network for Organ Sharing (UNOS) reported that as of October 28th, 2016, the waiting list for kidney transplant had 99,382 patients. Table 1 shows the aggregate number of patients entering and leaving the US waiting list after receiving a kidney (from a deceased or living donor) for the past five years (OPTN 2016).

		Removals for donation	
Year	Additions	Deceased donor	Living donor
2015	35,037	12,236	5585
2014	36,156	11,559	5284
2013	36,393	11,152	5265
2012	34,832	10,850	5088
2011	33,560	11,026	5154

Table 1: Number of patients entering and leaving the US national wait list (due to receiving a kidney).

Roughly two thirds of transplanted kidneys are sourced from cadavers, while the remaining one third come from willing healthy living donors. Patients who are fortunate enough to find a willing living donor must still contend with *compatibility* issues, including blood and tissue type biological compatibility. If a willing donor is incompatible with a patient, the transplantation cannot take place.

This is where *kidney exchange* comes in. A kidney exchange is a centrally-administered barter market where patients swap their willing yet incompatible donors. Modern kidney exchanges use 2-cycles, 3-cycles, and chains initiated by non-directed donors (altruists who are willing to give a kidney to anyone) as the means for swapping.

The idea of kidney exchange was introduced in 1986 (Rapaport 1986), and the first organized kidney exchanges started around 2003-04 (Roth, Sönmez, and Ünver 2004; 2005). Today there are kidney exchanges in the US, Canada, UK, the Netherlands, Australia, and many other countries. In the US, around 10% of live-donor kidney transplants now take place via exchanges.

Kidney exchange started first as matching markets where one donor-patient pair would give to, and receive from, another donor-patient pair. In other words, 2-cycles (Roth, Sönmez, and Ünver 2005) were the only structures used. Then, kidney exchange was generalized to also use 3-cycles (Roth, Sönmez, and Ünver 2007), and then short and finally never-ending chains initiated by non-directed donors (altruists who are willing to give to anyone without needing an organ in return) (Roth et al. 2006; Rees et al. 2009).

Significant work has been invested into scaling the market clearing algorithms, that is, the algorithms that find the optimal combination of non-overlapping (because any one donor can give at most one kidney) cycles and chains (Abraham, Blum, and Sandholm 2007; Constantino et al. 2013; Manlove and O'Malley 2015; Anderson et al. 2015b; Dickerson et al. 2016). There has also been significant work on improving the objective function to take into consideration several aspects of real-world kidney exchange programs, such as failure-awareness (Dickerson, Procaccia, and Sandholm 2013; Anderson 2014; Glorie et al. 2015) and a long-term (rather than batch) approaches of how the market clearing is done (Awasthi and Sandholm 2009; Ünver 2010; Dickerson, Procaccia, and Sandholm 2012a; Dickerson and Sandholm 2015) in the way the donors are assigned to patients.

Summary of Our Contribution

We propose a significantly generalized, more expressive, approach to kidney exchange. We allow more than one donor to donate in exchange for their desired patient receiving a kidney. We also allow for the possibility of a donor willing to donate if any of a number of patients receive kidneys. Furthermore, we combine these notions and generalize them.

Our generalization can be formalized around the concept of exchange among organ *clubs*, where, roughly speaking, a club is willing to donate organs outside the club if and only if the club receives organs from outside the club according to given specifications. More specifically, exchange clubs extend the notion of a donor-pair pair, allowing for a set of healthy donors equally willing to donate one of their kidneys in exchange for an equal (or greater) number of kidneys received by a target set of patients.

Forms of organ clubs already exist—under an arrangement where one gets to be in the club as a potential recipient if one is willing to donate one's organs to the club upon death. For example, there was such a club called *LifeSharers* in the US for several years (Hennessey 2006). It shut down in 2016 amid controversy regarding whether an organ club would actually hurt the nationwide organ allocation. Similarly, there is an organ club in the military "that allows families of active-duty troops to stipulate that their loved ones' organs go to another military patient or family (Kime 2016)." Also, Israel started an organ club where those who have given consent to become organ donors upon death (or whose family members have donated an organ in the past) get priority on the organ waitlist if they need organs; this increased organ donation in Israel by 60% in just one year (Stoler et al. 2016; Ofri 2012). One way to think of the approach that we are proposing is as an inter-club exchange mechanism that increases systemwide good—and can also be applied to live donation.

Our approach is beneficial also in a setting where there are no organ clubs in the traditional sense. We will nevertheless find the notion of a club useful in a technical sense to define the constraints, as we will detail later. We propose a formalization of this new kind of organ exchange, and propose an organ exchange approach where clubs are conceptually the primary agents—whether they are actually clubs, altruists, or donor-patient pairs, or a combination thereof. We support both intra-club and inter-club donations.

We propose a linear integer programming formulation for the optimization problem arising from the new club formalism, and note that it suffers of a synchronization issue, making it unsuitable for real-life implementations. Specifically, the issues are that (1) a club (of which a donor-patient pair is a special case, as is an altruist donor) wants to receive no later than it gives, and (2) there are logistical limits as to how many operations can be conducted simultaneously. In order to fix this problem, we introduce the concept of *operation frames*. Operations frames provide a convenient framework for handling the problem of synchronizing different transplants, by imposing a partial order on them. We then propose a second linear integer formulation overcoming the shortcoming of the first.

The Standard Model

Today's kidney exchanges (and other modern barter exchanges) can be modeled as follows. There is a directed *compatibility graph* $G = (V, E)$, where vertices represent participating parties and edges representing potential transactions (Roth, Sönmez, and Ünver 2007; Abraham, Blum, and Sandholm 2007). In the kidney exchange context, the set of vertices V is partitioned as $V = V_p \cup V_n$, where V_p represents the set of donor-patient pairs, and V_n represents the set of *non-directed* donors (NDDs).

For sake of simplicity, we will consider all non-directed donor vertices as formal donor-patient pairs, where the patient is an artificial object—denoted by $\perp$—that is incompatible with any donor in the system. Vertices u and v are connected by a directed edge $u \to v$ if the donor in u is compatible with the patient in v. The exchange administrator can also define a weight function $w : E \to \mathbb{R}$ representing, for each edge $e = (u, v) \in E$, the underlying quality or priority given to a potential transplant from $u \to v$.

Given the model above, we wish to solve the *clearing problem*, that is, we wish to select some subset of edges with maximum total weight subject to underlying feasibility constraints. For example, a donor d in a donor-patient pair $v = (d, p) \in V_p$ will donate a kidney if and *only if* a kidney is allocated to his or her paired patient p. Non-directed donors have no such constraint. In the model described so far, any solution consists of only two kinds of structure:

- *chains*, that is paths in G initiated by NDDs and then consisting entirely of donor-patient pairs; and
- *cycles*, that is loops in G consisting of vertices in V_p—and not non-directed donors in V_n.

Furthermore, in any feasible solution, these structures cannot share vertices: no donor can give more than one kidney. Figure 1 gives a feasible solution for a small example graph.

In kidney exchange, a length cap L is imposed on cycles for logistical reasons. All transplants in a cycle must be performed simultaneously so that no donor can back out after his patient has received a kidney but before he has donated his kidney. In most fielded exchanges worldwide, $L = 3$, so only 2-cycles and 3-cycles are allowed.

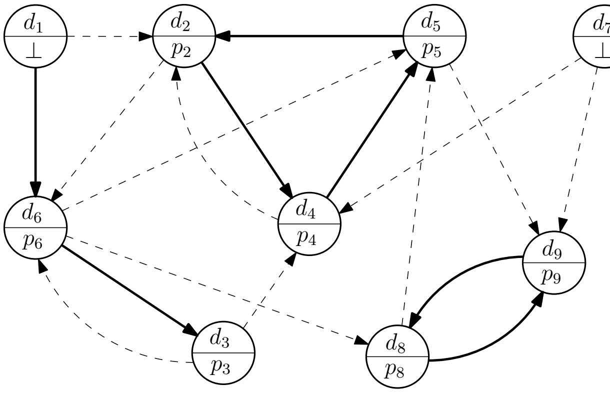

Figure 1: Example of a feasible solution in a pool with 9 donors and 7 patients. Donors d_1 and d_7 are NDDs. Dashed arrows represent compatible edges not selected in the solution. The solution uses one chain, one 2-cycle, and one 3-cycle. For simplicity, edge weights are omitted.

Chains do not need to be constrained in length, because it is not necessary to enforce that all transplants in the chain occur simultaneously. There is a chance that a donor backs out of her commitment to donate, but this event is less catastrophic than the equivalent in cycles. Indeed, a donor backing out in a cycle results in some other patient in the pool losing his donor while not receiving a kidney—that is, a participant in the pool is *strictly* worse off than before—while a donor backing out in a chain simply results in the chain ending. While that latter case is unfortunate, no participant in the pool is strictly worse off than before. In practice, however, a chain length cap is used, in order to make the planned solution more robust to last-minute failures (Dickerson, Procaccia, and Sandholm 2012b; Dickerson et al. 2016).

The problem can be formulated as an integer program to find the optimal solution, and indeed there has been significant work on developing increasingly scalable integer programming algorithms and formulations for this problem (e.g., (Roth, Sönmez, and Ünver 2007; Abraham, Blum, and Sandholm 2007)). The state of the art formulation is called PICEF (Dickerson et al. 2016). Its number of variables is polynomial in chain length cap and exponential in cycle length cap, which is not a problem in practice because the latter cap is small. Furthermore, the LP relaxation is very tight, causing good upper bounding in the search tree and therefore fast run time.

Exchange Clubs as a Modeling Construct

We propose significant generalizations to (kidney) exchange. We allow more than one donor to donate in exchange for their desired patient receiving a kidney. We also allow for the possibility of a donor willing to donate if any of a number of patients receive kidneys. Furthermore, we combine these notions and generalize them. We formalize this by introducing the modeling concept of *exchange clubs*.

Definition 1. (Exchange club) *An* exchange club *c is a tuple $(D_c, P_c, \alpha_c, \gamma_c)$ composed of*

- *a non-empty set of donors D_c;*
- *a (possibly empty) set of patients P_c;*
- *a real $\alpha_c \geq 1$ called "matching multiplier". Intuitively, this means that for each matched patient in P_c, the club is willing to donate (in expectation) α_c kidneys to the pool;*
- *a real $\gamma_c \geq 0$ called "matching debt".*

The idea of exchange clubs is that donors in D_c are willing to donate kidneys only if doing so results in a tangible benefit (that is, kidneys donated) to patients in P_c. More precisely, let $n_d^{\text{ext}}(t)$ be the number of kidneys donated from donors in D_c to clubs other that c by time t, and let $n_p^{\text{ext}}(t)$ be the number of kidneys donated from donors outside of c to patients in P_c; then the following inequality must hold for all time t in order for club c to be willing to participate in the solution:

$$n_d^{\text{ext}}(t) \leq \alpha_c n_p^{\text{ext}}(t) + \gamma_c \tag{1}$$

For now, we ignore parameter γ_c, whose role and motivation will become clear in the following sections.

We can now formalize the uncapped generalized clearing problem as follows.

Definition 2. (Disjoint clubs) *We say that two exchange clubs c and c' are* disjoint *if $P_c \cap P_{c'} = \varnothing$ and $D_c \cap D_{c'} = \varnothing$.*

Problem 1. (Uncapped generalized clearing problem) *Let $\mathcal{C}$ be a set of mutually disjoint exchange clubs; let $\mathcal{D} = \cup_{c \in \mathcal{C}} D_c$ and $\mathcal{P} = \cup_{c \in \mathcal{C}} P_c$ denote the overall set of donors and patients respectively. Furthermore, let $E \subseteq \mathcal{D} \times \mathcal{P}$ be the set of compatibility edges, and let $w : E \to \mathbb{R}$ a weighting function assigning a weight to every compatibility edge. We want to find a set of edges that maximizes the sum of weights and satisfies Inequality 1 assuming all the selected transplants occur simultaneously.*

Matching Debts

We now explain the meaning of γ_c. Suppose a number n_p^{ext} of patients in club c receive kidneys from other clubs, and that the optimal solution of the problem requires that n_d^{ext} donors from club c donate a kidney to other clubs. If $n_d^{\text{ext}} < \alpha_c n_p^{\text{ext}}$, we say that club c owes $\alpha_c n_p^{\text{ext}} - n_d^{\text{ext}}$ kidneys to the system. This is exactly the meaning of the "matching debt" of a club. It reflects the sum of all debts that a club has cumulated in the past. Except for clubs defined by non-directed donors, each club starts with a debt of 0 at the beginning, and potentially increases and decreases its debt to the system over time.

The Standard Model is a Special Case

The (uncapped) standard model is a special case of our new model:

- each non-directed donor defines a club c with no patient, and where he or she is the only donor. Furthermore, the club has $\gamma_c = 1$ (the value of α_c is irrelevant);
- each (d, p) donor-patient pair in the standard models defines a club c, where $D_c = \{d\}$, $P_c = \{p\}$ and $\alpha_c = 1$.

At the same time, our new model allows for some important generalizations. For instance, consider the case where one patient p has a set of two donors both willing to donate a kidney in exchange for only one kidney donated to p. In this case, the two donors and the p form a club with $\alpha_c = 2$.

The introduction of exchange clubs as a modeling construct calls for a different representation of the problem because the traditional donor-patient pairs cannot capture all the new aspects. Therefore, we explicitly represent donors and patients as different types of vertices in the graph. Figure 2 illustrates this under the further assumption that $\alpha_c = 1, \gamma_c = 0$ for all clubs. We represent donor vertices with a square and patient vertices with a circle. Observe that in Figure 2 it is not possible to extend the given solution with an edge from Donor 8 to Patient 7, as doing so would violate Inequality 1: Club D does not receive any kidney from other clubs, and therefore it cannot be asked to donate.

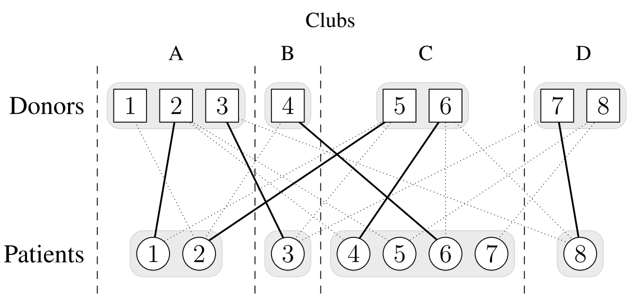

Figure 2: Tiny example problem instance. Vertical dashed lines separate different exchange clubs (so that, for instance, the first club has $D_1 = \{1, 2, 3\}, P_1 = \{1, 2\}$). Solid donor-patient edges exemplify a solution. Dotted donor-patient edges represent compatible pairs not used in the matching. For simplicity, we do not show edge weights in the figure. This figure assumes $\alpha_c = 1$, $\gamma_c = 0$ for all four clubs.

Incentive Issues

We briefly discuss some of the incentive issues in the new model. The issues we discuss arise when considering a club with a matching multiplier $\alpha_c > 1$, that is, the club is willing to donate more than it receives.

First of all, we give a bit more context as to why Inequality 1 requires that $n^{\text{ext}}_{\{d,p\}}$ refers to kidneys donated from and to *other* clubs, i.e. clubs different from c. This is because when $\alpha_c > 1$, we do not want transplants within the club to increase the number of donations the club is willing to make to other clubs. Furthermore, if one were to increase the debt a club owes to outside the club based on the intra-club transplants, that would (further) incentivize the club to not reveal their intra-club transplants. In terms of actual mechanism design, it is already known that even in versions of the standard model, one cannot achieve both efficiency and incentive compatibility (i.e., all centers being motivated to reveal all their pairs) in a single-shot setting (Ashlagi and Roth 2014; Ashlagi et al. 2015), but credit mechanisms hold promise in this regard (Hajaj et al. 2015).

Second, we might be forced to consider that when $\alpha_c > 1$, intra-club donations might be preferable to inter-club donations. Consider the scenario where $\alpha_c > 1$ and club c faces the decision of whether to accept a donation to patient $p \in P_c$ from a different club c', or match the same patient p with a donor $d \in D_c$. Unless the donation from club c' is much better than the one from donor d, club c would have incentive to match internally to avoid accruing a debt of $\alpha_c - 1$ to the system. In order to account for this issue, we can constrain the inter-club donations so that they can only happen if they are better than the alternative intra-club donations (if the latter exist). We ignore this issue in the rest of this paper, but we observe that the task is not difficult, as the constraints can be easily incorporated in Formulations 1 and 2 that we will present later in the paper.

Finally, we don't assume that every donor in D_c be incompatible with all patients in P_c. Indeed, we argue that even if some donor d is compatible with some patient p, we cannot simply conclude that matching d with p is necessarily a good idea, as it depends on the weight of the edge between them. In general, avoiding a greedy intra-club match would result in greater (or equal) value for the system as a whole.

Uncapped Problem Formulation

It is not clear how one could apply an integer program formulation like the state-of-the-art PICEF formulation for the standard kidney exchange problem (Dickerson et al. 2016) in this new setting. The problem here is that the realizability of a particular donation depends on what transplants have already been conducted. It does not seem immediate how such aspects could be encoded in a formulation like PICEF.

We now present an integer linear program formulation, Formulation 1, for finding an optimal solution to the uncapped clearing problem (Problem 1).

$$\max \sum_{(d,p)\in E} w_{dp}\, x_{dp}$$

$$① \quad \sum_{\substack{p\in\mathcal{P}\\(d,p)\in E}} x_{dp} \le 1 \qquad \forall\, d \in \mathcal{D}$$

$$② \quad \sum_{\substack{d\in\mathcal{D}\\(d,p)\in E}} x_{dp} \le 1 \qquad \forall\, p \in \mathcal{P}$$

$$③ \quad \sum_{d\in D_c} \sum_{\substack{p\in\mathcal{P}\setminus P_c\\(d,p)\in E}} x_{dp} \le \gamma_c + \alpha_c \sum_{p\in P_c} \sum_{\substack{d\in\mathcal{D}\setminus D_c\\(d,p)\in E}} x_{dp} \qquad \forall\, c \in \mathcal{C}$$

$$④ \quad x_{dp} \in \{0, 1\} \qquad \forall\, (d,p) \in E$$

Formulation 1: MIP formulation for the uncapped clearing problem.

Here, we let x_{dp} be a binary value (Constraint ④) indicating whether the edge $(d, p) \in E$ is selected in the solution. Constraints ① and ② ensure that each donor donates at most one kidney and that each patient receives at most one kidney, respectively. Constraint ③ encodes the condition that each club c donates at most $\lfloor \gamma_c + \alpha_c n^{\text{ext}}_p \rfloor$ kidneys. Finally, the objective function makes sure that we select a maximum-weight solution.

Operation Frames

While the above modeling approach is promising, we observe that it has a major shortcoming: it might require that a potentially large number of operations happen at the same time so as to honor the condition that the donors not be operated on strictly before patients in their clubs receive kidneys. An example is provided in Figure 3, where we would need patients $\{1, 2, 3, 6, 7\}$ and $\{1, 2, 3, 5, 6\}$ to be operated on at the same time.

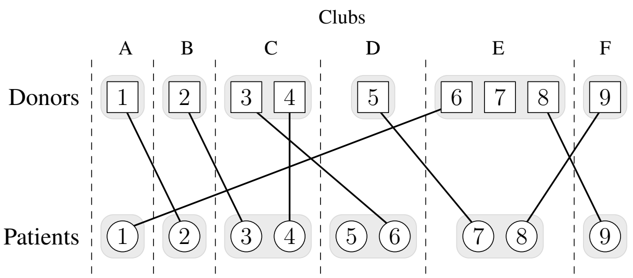

Figure 3: Example of a match requiring 5 simultaneous transplantations. Patients $\{1, 2, 3, 6, 7\}$ and donors $\{1, 2, 3, 5, 6\}$ have to be operated on at the same time. The same holds, at a smaller scale, for patients $\{8, 9\}$ and donors $\{8, 9\}$. The transplantation between donor 4 and patient 4 can be completed any time. Admissible unused edges are not shown; we assumed $\alpha_c = 1, \gamma_c = 0$ for all six clubs.

This is not practically viable for at least two reasons:

- the success probability of all the planned transplants in the structure succeeding in their pre-operation blood type compatibility tests (aka. crossmatch test) and other pre-transplant testing decreases multiplicatively with the number of edges in the planned structure,[1] and
- the logistic (and financial) details are hard to execute—ten people to operate on have to be coordinated, together with the surgeons and staff needed for ten surgeries.

In order to solve this synchronization issue, we introduce the concept of *operation frames*. An operation frame t is an edge set of size up to K_t, representing operations to be performed at the same time. The introduction of operation frames enables us to reason in terms of order in which the operations will be carried out. The chronological order imposed on the operation frames is *partial*. For this reason, we can formalize the set and relationships among operation frames by means of a directed acyclic graph (DAG) $F = (T, B)$, where the set of vertices (i.e., T) coincides with the set of operation frames, while the set of edges $B \subseteq T \times T$ denotes the *happens-strictly-before* chronological (partial) order. We say that operation frame u happens strictly before operation frame v, denoted by $u \rightsquigarrow v$, if there exists a directed path in F from u to v. Figure 4 gives an example.

The introduction of operation frames enables Inequality 1 to be written in terms of *logical* time, that is, substituting the notion of time with the partial happens-strictly-before order.

[1]For further details about pre-transplant test failures, see Dickerson, Procaccia, and Sandholm (2013) and Blum et al. (2015).

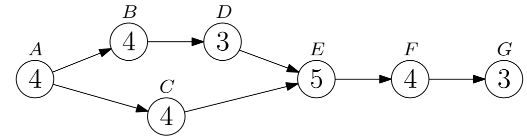

Figure 4: Example of a partial order imposed on the operation frames $\{A, \ldots, G\}$. The number in each operation frame t is the cap size for the operation frame, K_t. Examples of the happens-strictly-before relation include $A \rightsquigarrow D$, $A \rightsquigarrow F$, $C \rightsquigarrow E$. Notice that the pairs (B, C) and (D, C) are not comparable according the partial order "$\rightsquigarrow$".

Thus, for any frame $\tau \in T$, the number of kidneys that were surely (i.e., for any possible linearization of the DAG F) donated to and from club c at the time when τ is executed is

$$n_d(\tau) = \sum_{\tau' \rightsquigarrow \tau} d(c, \tau'), \quad n_p(\tau) = \sum_{\tau' \rightsquigarrow \tau} p(c, \tau'),$$

where $d(c, \tau')$ and $p(c, \tau')$ represent the number of transplant from and to club c scheduled for operation frame τ', respectively.

We argue that operation frames provide a richer problem structure, as it is now possible to assign a (partial) chronological order to the operations we plan to perform. Furthermore, they encode the condition that "no more than K_t people get operated on at the same time" in a very natural way: every operating frame has a parameter K_t. Indeed, operation frames guarantee that not too many surgeries are planned to happen at the same time. Conceptually, they are equivalent to imposing chain and cycle length caps, as it is done in standard model. However, here we are allowing much richer exchange structures (and, as presented so far, there is no way of specifying a different cap on the size of chains versus cycles versus other structures). This is because the operation frame size cap K_t represents an actual limit on the number of simultaneous operations that can be accommodated. Different frames can have different size constraints.

Finally, operation frames allow an integer programming formulation of the problem that (unlike PICEF) uses a number of variables that is polynomial in the maximum size cap $\max_t K_t$. We present that formulation in the next section.

Capped Problem Formulation

The idea of operation frames can be plugged into Formulation 1, leading to the following formulation, Formulation 2.

We let x^t_{dp} be a binary value (Constraint 5) indicating whether the transplant represented by the edge $(d, p) \in E$ is scheduled for operation frame t. Analogous to the uncapped case, Constraints 1 and 2 ensure that each donor donates at most one kidney and that each patient receives at most one kidney, respectively. Constraint 4 ensures that in any operation frame $t \in \mathcal{T}$, no more than K transplants are scheduled. Constraint 3 enforces that at any time t, for each club $c \in \mathcal{C}$, the total number of kidneys donated from club c does not exceed $\lfloor \gamma_c + \alpha_c n_p^{\text{ext}}(t) \rfloor$, where $n_p^{\text{ext}}(t)$ is the total number of kidneys donated to club c from other clubs, before or at operation frame t. The objective function ensures that a maximum-weight solution is found.

$$\max \sum_{(d,p)\in E} \sum_{t\in T} h(t)\, w_{dp}\, x_{dp}^t$$

$$① \quad \sum_{\substack{p\in\mathcal{P}\\(d,p)\in E}} \sum_{t\in T} x_{dp}^t \leq 1 \qquad \forall\, d\in\mathcal{D}$$

$$② \quad \sum_{\substack{d\in\mathcal{D}\\(d,p)\in E}} \sum_{t\in T} x_{dp}^t \leq 1 \qquad \forall\, p\in\mathcal{P}$$

$$③ \quad \sum_{\tau\rightsquigarrow t} \sum_{d\in D_c} \sum_{\substack{p\in\mathcal{P}\setminus P_c\\(d,p)\in E}} x_{dp}^\tau \leq \gamma_c + \alpha_c \sum_{\tau\rightsquigarrow t} \sum_{p\in P_c} \sum_{\substack{d\in\mathcal{D}\setminus D_c\\(d,p)\in E}} x_{dp}^\tau \qquad \forall\, c\in\mathcal{C},\ t\in T$$

$$④ \quad \sum_{(d,p)\in E} x_{dp}^t \leq K \qquad \forall\, t\in T$$

$$⑤ \quad x_{dp}^t \in \{0,1\} \qquad \forall\,(d,p)\in E,\ t\in T$$

Formulation 2: MIP formulation for the capped problem.

One can also model temporal preferences by multiplying the edge weights by discounts $h(t)$, which depend on which operation frame t the surgery is conducted. (This assumes that the time between frames is exogenous—but not necessarily constant—that is, the time between frames does not depend on what transplants the optimizer decides to put in each frame.) This discounting is already include in the objective in Formulation 2.

Operation Frames Make the System Less Myopic

Present-day kidney exchanges operate in a batch setting, potentially planning in a single shot long chains that will, in practice, execute in segments over many months. Solvers for the standard problem (e.g., those based on PICEF) optimize on a batch-by-batch basis, selecting the global optimum solution only inside of a single batch, and not considering future batches. Our approach is more powerful than the standard batch-based one also in the sense that it inherently breaks long structures into shorter ones that execute sequentially. Figure 5 shows an example compatibility graph where our model will return a higher-value solution than the optimal solution in the traditional model.

Figure 2 shows that the capped approach in the standard model cannot consider certain solutions; yet, our capped approach—based on the concept of operations frames—optimizes across all the operation frames at the same time, resulting in less myopic behavior. Under the assumption that the kidney exchange pool is not affected by any exogenous behavior (e.g., compatibility failures, deaths of donors or patients, dynamic insertions and deletions of edges and vertices), our formulation is guaranteed to find a globally optimal allocation of transplants across all operation frames. In contrast, even under these strong assumptions, present-day solvers for the standard model will (by design) fail to find a globally optimal solution across batches.

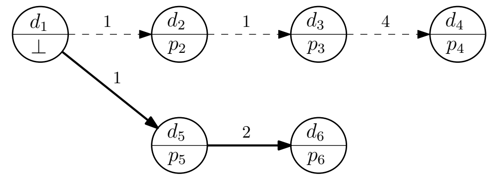

Figure 5: An example graph where Formulation 2 returns a higher-value solution than the standard batch approach, even with only the standard kinds of vertices available. Given a cap $K = 2$, the standard approach will choose the lower 2-chain over the upper 2-chain for utility 3, while our solver will choose to match the upper 3-chain across two frames for greater overall utility of 6.

Conclusions & Future Research

Motivated by the reality of fielded kidney exchanges, in this paper we proposed significant generalizations to kidney exchange—and barter markets more generally. Specifically, we moved the model from individual and independent patient-donor pairs to the modeling concept of multi-donor and multi-patient organ *clubs*, where a club is willing to donate organs outside the club if and only if the club receives organs from outside the club according to expressed preferences. We presented the notion of operation frames that sequence the operations across batches, and gave IP formulations that optimally clear these new types of markets.

Operation frames inherently include a notion of time via the *happens-strictly-before* ordering; yet, this does not capture the full dynamics of kidney exchange, where vertices and edges arrive and disappear over time. Finding an optimal matching policy for fully dynamic kidney exchange is an open problem from both the theoretical (Akbarpour, Li, and Gharan 2014; Anderson et al. 2015a) and computational (Awasthi and Sandholm 2009; Dickerson, Procaccia, and Sandholm 2012a; Dickerson and Sandholm 2015) points of view; perhaps the ordering introduced by operation frames can be used to decrease computational intractability when reasoning in the fully dynamic setting.

More fully exploring incentive issues in this new model is practically interesting. Incentives at the patient or donor level have not been explored thoroughly in the kidney exchange literature beyond "a donor does not have an incentive to donate unless his paired patient receives a kidney." Organ clubs give patients and donors a new, variable amount of bargaining power. Furthermore, in the past, generalizations to the basic kidney exchange model have already allowed mechanism designers to circumvent strong impossibility results (Hajaj et al. 2015; Ashlagi and Roth 2014). The more expressive models presented in this paper could result in similar advances in designing mechanisms with desirable game-theoretic properties.

Acknowledgments. Funded by NSF grants IIS-1617590, IIS-1320620, IIS-1546752, and ARO award W911NF-16-1-0061.

References

Abraham, D.; Blum, A.; and Sandholm, T. 2007. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 295–304.

Akbarpour, M.; Li, S.; and Gharan, S. O. 2014. Dynamic matching market design. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 355.

Anderson, R.; Ashlagi, I.; Gamarnik, D.; and Kanoria, Y. 2015a. A dynamic model of barter exchange. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1925–1933.

Anderson, R.; Ashlagi, I.; Gamarnik, D.; and Roth, A. E. 2015b. Finding long chains in kidney exchange using the traveling salesman problem. *Proceedings of the National Academy of Sciences* 112(3):663–668.

Anderson, R. 2014. *Stochastic models and data driven simulations for healthcare operations*. Ph.D. Dissertation, Massachusetts Institute of Technology.

Ashlagi, I., and Roth, A. E. 2014. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics* 9(3):817–863.

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

Awasthi, P., and Sandholm, T. 2009. Online stochastic optimization in the large: Application to kidney exchange. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 405–411.

Blum, A.; Dickerson, J. P.; Haghtalab, N.; Procaccia, A. D.; Sandholm, T.; and Sharma, A. 2015. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, 325–342.

Constantino, M.; Klimentova, X.; Viana, A.; and Rais, A. 2013. New insights on integer-programming models for the kidney exchange problem. *European Journal of Operational Research* 231(1):57–68.

Dickerson, J. P., and Sandholm, T. 2015. FutureMatch: Combining human value judgments and machine learning to match in dynamic environments. In *AAAI Conference on Artificial Intelligence (AAAI)*, 622–628.

Dickerson, J. P.; Manlove, D.; Plaut, B.; Sandholm, T.; and Trimble, J. 2016. Position-indexed formulations for kidney exchange. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.

Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2012a. Dynamic matching via weighted myopia with application to kidney exchange. In *AAAI Conference on Artificial Intelligence (AAAI)*, 1340–1346.

Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2012b. Optimizing kidney exchange with transplant chains: Theory and reality. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 711–718.

Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2013. Failure-aware kidney exchange. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 323–340.

Glorie, K.; Carvalho, M.; Constantino, M.; Bouman, P.; and Viana, A. 2015. Robust models for the kidney exchange problem. Working paper.

Hajaj, C.; Dickerson, J. P.; Hassidim, A.; Sandholm, T.; and Sarne, D. 2015. Strategy-proof and efficient kidney exchange using a credit mechanism. In *AAAI Conference on Artificial Intelligence (AAAI)*, 921–928.

Hennessey, J. 2006. A members-only club for organ donors. ABC News website. `http://abcnews.go.com/Health/story?id=2071084`.

Kime, P. 2016. Defense Department charts new territory in kidney transplants. Military Times website. `http://www.militarytimes.com/story/military/2016/06/22/defense-department-charts-new-territory-/kidney-transplants/86176348/`.

Manlove, D., and O'Malley, G. 2015. Paired and altruistic kidney donation in the UK: Algorithms and experimentation. *ACM Journal of Experimental Algorithmics* 19(1).

Neuen, B. L.; Taylor, G. E.; Demaio, A. R.; and Perkovic, V. 2013. Global kidney disease. *The Lancet* 382(9900):1243.

NIH. 2011. Kidney disease statistics for the united states. *Bethesda, MD: National Kidney and Urologic Diseases Information Clearinghouse, National Institutes of Health*.

Ofri, D. 2012. In israel, a new approach to organ donation. The New York Times website. `http://well.blogs.nytimes.com/2012/02/16/in-israel-a-new-approach-to-organ-donation/`.

OPTN. 2016. Organ procurement and transplantation network data. U.S. Department of Health and Human Services. `https://optn.transplant.hrsa.gov/data/`.

Rapaport, F. T. 1986. The case for a living emotionally related international kidney donor exchange registry. *Transplantation Proceedings* 18:5–9.

Rees, M.; Kopke, J.; Pelletier, R.; Segev, D.; Rutter, M.; Fabrega, A.; Rogers, J.; Pankewycz, O.; Hiller, J.; Roth, A.; Sandholm, T.; Ünver, U.; and Montgomery, R. 2009. A nonsimultaneous, extended, altruistic-donor chain. *New England Journal of Medicine* 360(11):1096–1101.

Roth, A.; Sönmez, T.; Ünver, U.; Delmonico, F.; and Saidman, S. L. 2006. Utilizing list exchange and nondirected donation through 'chain' paired kidney donations. *American Journal of Transplantation* 6:2694–2705.

Roth, A.; Sönmez, T.; and Ünver, U. 2004. Kidney exchange. *Quarterly Journal of Economics* 119(2):457–488.

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

Roth, A.; Sönmez, T.; and Ünver, U. 2007. Efficient kidney exchange: Coincidence of wants in a market with compatibility-based preferences. *American Economic Review* 97:828–851.

Stoler, A.; Kessler, J.; Ashkenazi, T.; Roth, A.; and Lavee, J. 2016. Incentivizing authorization for deceased organ donation with organ allocation priority: The first 5 years. *American Journal of Transplantation* 16(9):2639–45.

Ünver, U. 2010. Dynamic kidney exchange. *Review of Economic Studies* 77(1):372–414.