

University of Michigan School of Public Health

The University of Michigan Department of Biostatistics Working
Paper Series

Year 2011

Paper 90

Optimization and Simulation of an Evolving Kidney Paired Donation (KPD) Program

Yijiang Li*

Jack Kalbfleisch[†]

Peter Xuekun Song[‡]

Yan Zhou**

Alan Leichtman^{††}

Michael Rees^{‡‡}

*University of Michigan - Ann Arbor, yijiang@umich.edu

[†]University of Michigan, jdkalbfl@umich.edu

[‡]University of Michigan, Ann Arbor, pxsong@umich.edu

**University of Michigan, Ann Arbor, zhouyan@umich.edu

^{††}University of Michigan, Ann Arbor, aleicht@umich.edu

^{‡‡}University of Toledo, michael.rees2@utoledo.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/umichbiostat/paper90>

Copyright ©2011 by the authors.

Optimization and Simulation of an Evolving Kidney Paired Donation (KPD) Program

Yijiang Li, Jack Kalbfleisch, Peter Xuekun Song, Yan Zhou, Alan Leichtman,
and Michael Rees

Abstract

The old concept of barter exchange has extended to the modern area of living-donor kidney transplantation, where one incompatible donor-candidate pair is matched to another pair with a complementary incompatibility, such that the donor from one pair gives an organ to a compatible candidate in the other pair and vice versa. Kidney paired donation (KPD) programs provide a unique and important platform for living incompatible donor-candidate pairs to exchange organs in order to achieve mutual benefit. We propose a novel approach to organizing kidney exchanges in an evolving KPD program with advantages, including (i) it allows for a more flexible utility-based evaluation of potential kidney transplants; (ii) it takes into consideration stochastic features in managing a KPD program; and (iii) it exploits possible alternative exchanges when the originally planned allocation cannot be fully executed. Another primary contribution of this work is rooted in the development of a comprehensive microsimulation system for simulating and studying various aspects of an evolving KPD program. Various allocations can be obtained using integer programming (IP) techniques and microsimulation models can allow tracking of the evolving KPD over a series of match runs to evaluate different allocation strategies. Simulation studies are provided to illustrate the proposed method.

Optimization and simulation of an evolving kidney paired donation (KPD) program

Yijiang Li^{1,*}, John D. Kalbfleisch¹, Peter X.K. Song¹, Yan Zhou¹, Alan Leichtman², Michael A. Rees³

¹Department of Biostatistics, University of Michigan, Ann Arbor, Michigan 48109, USA

²Department of Internal Medicine, University of Michigan, Ann Arbor, Michigan 48109, USA

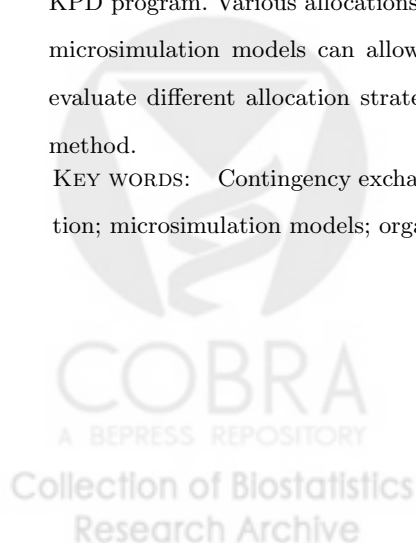
³Department of Urology, University of Toledo Medical Center, Toledo, Ohio 43614, USA

The Alliance for Paired Donation, Maumee, Ohio 43537, USA

**email*: yijiang@umich.edu

SUMMARY: The old concept of barter exchange has extended to the modern area of living-donor kidney transplantation, where one incompatible donor-candidate pair is matched to another pair with a complementary incompatibility, such that the donor from one pair gives an organ to a compatible candidate in the other pair and vice versa. Kidney paired donation (KPD) programs provide a unique and important platform for living incompatible donor-candidate pairs to exchange organs in order to achieve mutual benefit. We propose a novel approach to organizing kidney exchanges in an evolving KPD program with advantages, including (i) it allows for a more flexible utility-based evaluation of potential kidney transplants; (ii) it takes into consideration stochastic features in managing a KPD program; and (iii) it exploits possible alternative exchanges when the originally planned allocation cannot be fully executed. Another primary contribution of this work is rooted in the development of a comprehensive microsimulation system for simulating and studying various aspects of an evolving KPD program. Various allocations can be obtained using integer programming (IP) techniques and microsimulation models can allow tracking of the evolving KPD over a series of match runs to evaluate different allocation strategies. Simulation studies are provided to illustrate the proposed method.

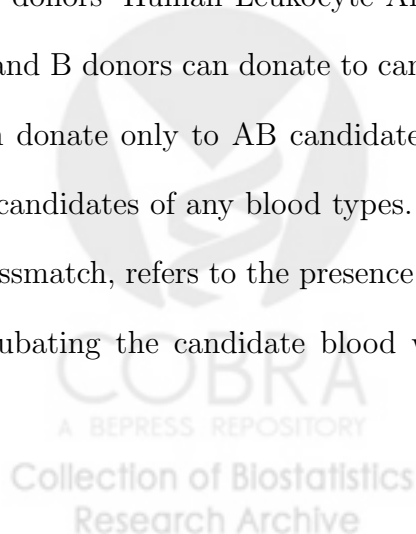
KEY WORDS: Contingency exchanges; expected utility; integer programming; kidney paired donation; microsimulation models; organ exchange.



1. Introduction

For patients with end-stage renal disease, successful renal transplantation provides a considerably better quality of life and improved survival, as compared with the dialysis treatment (Wolfe et al., 1999; Evans et al., 1985; Russell et al., 1992). Cost effectiveness is another advantage of renal transplantation as compared to continuing dialysis (Laupacis et al., 1996). However, due to limited supplies of cadaveric kidneys and a substantial and growing demand for them, many patients who need a transplant have been placed on long waiting lists. According to the Organ Procurement and Transplantation Network (OPTN), as of December 2010, more than 87,000 kidney transplant candidates in the U.S. are on a waiting list; and in 2009, about 34,000 candidates in the U.S. were added to the list, whereas only about 10,000 actually received a kidney transplant from a deceased donor. In one response to this shortage, candidates have increasingly undergone living-donor transplants. Moreover, living-donor transplants have the advantage of a higher graft survival rate, in both the short and the long term, than deceased-donor transplants (Terasaki et al., 1995; Hariharan et al., 2000).

A major issue surrounding the living-donor kidney transplant is the unfortunate fact that willing donors, even related living donors, are often incompatible with their intended candidates, because of ABO blood type incompatibility and/or antibodies against some of the donors' Human Leukocyte Antigens (HLA). With respect to blood type compatibility, A and B donors can donate to candidates of the same blood type or of type AB; AB donors can donate only to AB candidates; and O donors, known as universal donors, can donate to candidates of any blood types. The second form of incompatibility, also called a positive crossmatch, refers to the presence of anti-donor antibodies in the blood of a candidate when incubating the candidate blood with the white blood cells of a prospective donor. Both



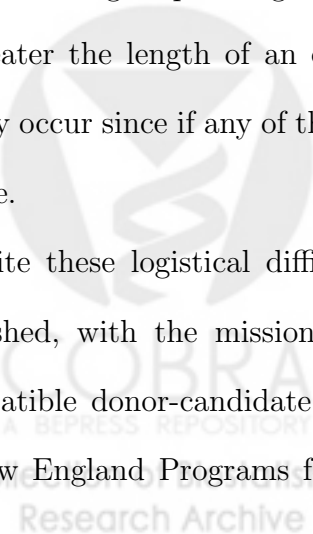
forms of incompatibility can lead to rejection of the transplanted organ and thus prohibit transplant.

Some of these incompatible donors, however, might be completely compatible with other potential candidates. In 1986, the idea of kidney paired donations (KPD) was set forth by a transplant surgeon, Felix Rapaport. Rapaport (1986) and later Ross et al. (1997) envisioned that two candidates with willing, but incompatible, donors could solve each other's problem by simultaneously exchanging the two donors' kidneys; see Figure 1-A for an example of two-way exchanges. More complex exchanges of organs involving three or more pairs are also possible, as schematically illustrated in Figure 1-B.

[Figure 1 about here.]

While three-way or higher exchange cycles increase the chance of finding compatible matches, most transplant centers in practice restrict exchanges to at most three ways mainly for two reasons: (i) All operations on an exchange cycle must be performed simultaneously to avoid the risk that one of the donors may renege. This avoids a situation where a donor withdraws his or her commitment after the other donor has undergone nephrectomy and donated to the candidate associated with this reneging donor. This requirement of performing operations simultaneously creates substantial logistical difficulties of scheduling, say, eight surgeons and eight operating rooms at the same time for a four-way exchange. (ii) In addition, the greater the length of an exchange cycle, the less likely the potential transplants will actually occur since if any of the proposed transplants cannot proceed the whole cycle would collapse.

Despite these logistical difficulties, more and more KPD programs have recently been established, with the mission of promoting mutually beneficial organ exchanges among incompatible donor-candidate pairs. Regional programs in the U.S. include, for example, the New England Programs for Kidney Exchange, the Paired Kidney Exchange Program



at Hopkins, the Alliance for Paired Donation (APD), the National Kidney Registry (NKR) Program, and the University of Michigan Paired Kidney Exchange Program; internationally, KPD programs exist in Netherlands (Keizer et al., 2005) and South Korean (Park et al., 1999). Additional developments include a recent OPTN approval of a national KPD pilot program.

Researchers from several disciplines, collaborating with transplant surgeons, have recently become more involved in organizing and optimizing kidney exchanges in a KPD program. For instance, economists have modeled and analyzed kidney exchanges using game-theoretic approaches. Roth, Sönmez, and Ünver (2004) organized donor-candidate pairs as a “housing market”, a concept first proposed by Shapley and Scarf (1974), and produced an efficient organ exchange mechanism using Gale’s Top Trading Cycles (TTC) algorithm. Roth, Sönmez, and Ünver (2005) and Segev et al. (2005) applied maximum cardinality matching algorithm (Edmonds, 1965) to select exchanges that allow the maximum number of transplants, in the case where only two-way exchanges are considered. Determining optimal exchanges, in general, can be formulated as an integer programming (IP) problem (Roth, Sönmez, and Ünver, 2007). This problem can be efficiently solved by finding a maximum weight perfect matching when no restriction is placed on the cycle length. If, however, only cycles of length up to k are considered, this IP problem is NP-hard when k is larger than two but less than the number of participating pairs (Roth et al., 2007). To address this issue, Abraham, Blum, and Sandholm (2007) have recently proposed an exact algorithm that performs quite satisfactorily in this case. This work has greatly extended these methods to potentially handle large KPD programs.

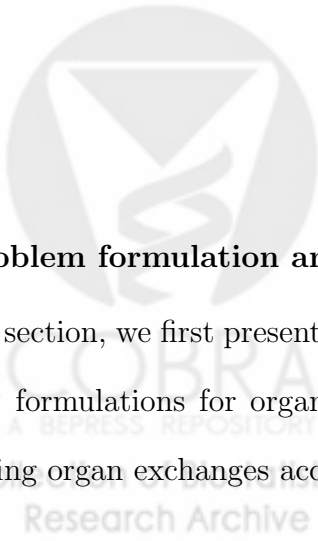
This paper extends and improves upon the research described above to optimize and simulate a KPD program. Our proposed kidney allocation strategy is innovative in several respects. First, it allows a quality-oriented evaluation of a kidney allocation through medical-

outcome-based utilities such as post-transplant graft survival and recipient survival. Second, it explicitly takes into consideration the probability that a predicted compatible transplant will result in an actual transplant operation; this recognition of inherent uncertainty offers substantial improvements over previous approaches. Third, our approach allows for one or more contingency allocations should the originally planned exchanges fail to be executed. Finally, we propose a data-based microsimulation system for simulating an evolving KPD, based on which we can evaluate different kidney allocation strategies, compare their impact on performance outcomes, and assess different choices of utility assignments. The knowledge learned from such a microsimulation should provide invaluable guidance in implementing an actual clinical KPD allocation system.

The rest of the article is organized as follows: in Section 2, we introduce the representation and formulation of a KPD program and present a procedure for arranging kidney exchanges according to the maximum utility cycle-based allocation. In Section 3, we explore some important issues and features in a KPD program that have not been addressed by previous studies and further propose a kidney allocation strategy based on the maximum expected-utility set-based allocation. In Section 4, we present a microsimulation system for simulating an evolving KPD. Section 5 reports simulation results to illustrate the application of the proposed microsimulation system. We conclude with some discussion in Section 6.

2. Problem formulation and the maximum utility cycle-based allocation

In this section, we first present a graph representation of a KPD program, and then describe two IP formulations for organizing kidney exchanges. Also, we introduce a procedure for arranging organ exchanges according to the maximum utility cycle-based allocation.



2.1 Graph representation

We represent a KPD program as a *directed graph*, $G = (V, E)$, where the *vertex set*, $V = \{1, 2, \dots, n\}$, is the set of n incompatible donor-candidate pairs, and the *edge set* E is a binary relation on V , consisting of ordered pairs of vertices in V . An edge from v_1 to v_2 , denoted as (v_1, v_2) , indicates that the donor of pair v_1 is predicted to be compatible with the candidate of pair v_2 . This predicted compatibility is based on a *virtual crossmatch*, which involves computer cross-checking pair characteristics such as blood types and HLA immune information. A negative virtual crossmatch may or may not lead to a negative result in the confirmatory *laboratory crossmatch*, which involves incubating the blood of a candidate with the white blood cells of a prospective donor. It is worth pointing out that prior research on KPD has not taken into consideration this uncertainty, and instead has proceeded as though a negative virtual crossmatch would guarantee, if chosen, a completed transplant. For notational convenience, we denote an edge (v_1, v_2) as e in the case that specifying the donor and candidate is not necessary. Throughout the rest of the paper, we use “a predicted compatible transplant”, “a potential transplant”, “a negative virtual crossmatch”, and “an edge” interchangeably.

In such a directed graph, an *exchange cycle* of length k (or a k -way exchange cycle), $k \geq 2$, is defined as a sequence of vertices, $\langle v_1, v_2, \dots, v_k \rangle$, satisfying (i) v_1, v_2, \dots, v_k are distinct, and (ii) $(v_k, v_1) \in E$ and $(v_{j-1}, v_j) \in E, \forall j = 2, 3, \dots, k$. For an exchange cycle $c = \langle v_1, v_2, \dots, v_k \rangle$, we denote its *vertex set* as $V(c) = \{v_1, v_2, \dots, v_k\}$ and its *edge set* as $E(c) = \{(v_k, v_1), (v_{j-1}, v_j), j = 2, 3, \dots, k\}$. In Figure 2 are shown one two-way exchange cycle, $\langle v_2, v_4 \rangle$, and two three-way exchange cycles, $\langle v_1, v_2, v_3 \rangle$ and $\langle v_2, v_4, v_5 \rangle$. Exchange cycles form a *disjoint* collection if their corresponding vertex sets are disjoint. A *cycle-based allocation* for a KPD program G is defined as a collection of disjoint exchange cycles, and further denoted as $\mathcal{C}(G)$. An alternative *set-based allocation*, denoted as $\mathcal{S}(G)$, will be introduced

and discussed in Section 3. In the context of cycle-based allocations, a fundamental problem in managing a KPD program is to select the “optimal” allocation from among the many mathematically possible ones, and the choice of such an “optimal” allocation depends on how an exchange cycle is evaluated.

[Figure 2 about here.]

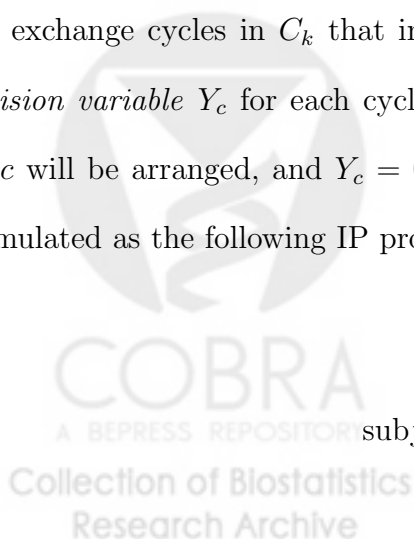
For an edge $(v_1, v_2) \in E$, let $u_{(v_1, v_2)} \geq 0$ denote the *utility* of a kidney transplant from the donor in pair v_1 to the candidate in pair v_2 . The utility of an exchange cycle c , is the sum of its edge utilities, i.e. $U_c = \sum_{e \in E(c)} u_e$, and the utility of a cycle-based allocation $\mathcal{C}(G)$, is the sum of the utilities of its cycles, i.e. $\sum_{c \in \mathcal{C}(G)} U_c$. In this setting, the optimal cycle-based allocation is the one with the maximal utility, denoted as $\mathcal{C}^*(G)$.

Roth et al. (2007) proposed two different IP formulations to determine $\mathcal{C}^*(G)$ that were later also adopted in Abraham et al. (2007). One formulation encodes each exchange cycle as a decision variable, and the other one encodes each edge as a decision variable. We first look at the *cycle formulation* in Section 2.2 and then the *edge formulation* in Section 2.3.

2.2 Cycle formulation

Let C_k be the set of exchange cycles with lengths at most k . For each $i \in V$, let $C_k(i)$ denote the exchange cycles in C_k that involve pair i , i.e. $C_k(i) = \{c \in C_k : i \in V(c)\}$. Define a *decision variable* Y_c for each cycle $c \in C_k$, such that Y_c is 1 if organ exchanges indicated by c will be arranged, and $Y_c = 0$ otherwise. The problem of selecting $\mathcal{C}^*(G)$ can then be formulated as the following IP problem,

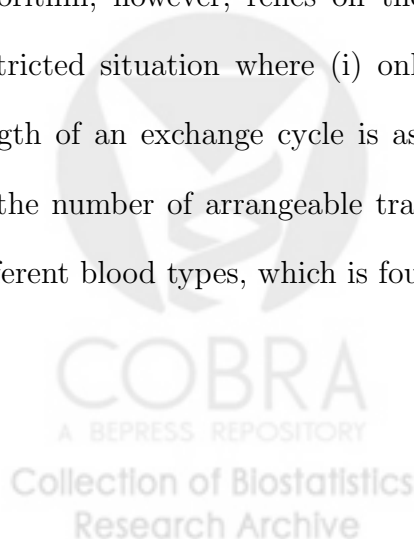
$$\begin{aligned} \max_{\{Y_c\}} \quad & \sum_{c \in C_k} Y_c U_c, \\ \text{subject to} \quad & \sum_{c \in C_k(i)} Y_c \leq 1, \forall i \in V. \end{aligned}$$



Note that the constraint simply codifies the fact that each pair cannot be allowed in more than one exchange cycle simultaneously. Each feasible solution to this IP problem corresponds to one $\mathcal{C}(G)$, and the optimal solution corresponds to $\mathcal{C}^*(G)$.

As a special case in this IP problem, when the utility of an exchange cycle is defined as its length, i.e. $U_c = |V(c)|$, the resulting objective is to maximize the total number of transplants. Most KPD studies have focused on this simplified utility assignment; see more discussion in Section 3.

The cycle formulation usually leads to increased computation time as k grows, though in practice most KPD programs restrict k to be three or smaller due to logistical difficulties. Some particular values of k are worth mentioning. When k is equal to two, the optimization problem could be solved in polynomial time using a maximum weighted matching algorithm, which is an extended version of Edmonds' classical maximal cardinality matching algorithm. When k is equal to $|V|$, i.e. no restriction is placed on the length of an exchange cycle, the *edge formulation*, to be introduced in Section 2.3, yields the optimal solution in polynomial time. As noted earlier, when k is greater than two but less than $|V|$, this optimization problem is NP-hard, which poses with associated computational challenges when the number of incompatible pairs is large. To address this issue, Abraham et al. (2007) developed an exact algorithm that can perform satisfactorily in practice. The efficiency of their proposed algorithm, however, relies on the following property proved in Roth et al. (2007). In a restricted situation where (i) only blood type incompatibility is considered, and (ii) the length of an exchange cycle is assigned as its utility, no improvement would be obtained in the number of arrangeable transplants by allowing k to be greater than the number of different blood types, which is four in this case.



2.3 Edge formulation

Let $Y_{(i,j)}$ be a decision variable for each $(i,j) \in E$, such that $Y_{(i,j)}$ is 1 if edge (i,j) is chosen for a transplant and 0 otherwise. When no restriction is placed on the length of an exchange cycle, we solve the following IP problem:

$$\begin{aligned} & \max_{\{Y_{(i,j)}\}} \sum_{(i,j) \in E} Y_{(i,j)} u_{(i,j)}, \\ & \text{subject to} \quad \sum_{j:(i,j) \in E} Y_{(i,j)} \leq 1, \forall i \in V \\ & \quad \sum_{j:(i,j) \in E} Y_{(i,j)} = \sum_{j:(j,i) \in E} Y_{(j,i)}, \forall i \in V. \end{aligned}$$

In this formulation, the IP problem could be solved efficiently in polynomial time by finding a maximum weight perfect matching. If the exchange cycle length is restricted to be at most k , an additional set of constraints has to be added, i.e.

$$\sum_{(i,j) \in E(c)} Y_{(i,j)} \leq k, \forall c \in C_k.$$

The number of additional constraints, even when $k = 3$, is usually enormously large in a realistic KPD program with several hundred incompatible pairs, which makes it impossible to even store all the constraints in a typical IP solver such as CPLEX or Gurobi. Therefore, the cycle formulation is usually preferred in this case.

2.4 The maximum utility cycle-based allocation

We summarize below the procedure for arranging kidney exchanges according to the maximum utility cycle-based allocation. Given a KPD program $G = (V, E)$,

- (i) Define $u : E \rightarrow R^+$, where u_e is the assigned utility of an edge $e \in E$.
- (ii) Enumerate C_k and find U_c for each $c \in C_k$.
- (iii) Determine $\mathcal{C}^*(G)$ by solving an IP problem with cycle formulation.
- (iv) Arrange kidney exchanges according to $\mathcal{C}^*(G)$.

In (ii), we enumerate C_k by developing an algorithm based on deep-first search (DFS). DFS prioritizes the direction of search to offspring vertices first, and then to sibling vertices. This is the first of three procedures to organize kidney exchanges in a KPD program. Two others are presented and further discussed in Section 3.

3. Optimal kidney allocation in a KPD program

In Section 2, we have introduced a graph representation of a KPD program and presented an allocation strategy based on $\mathcal{C}^*(G)$. In this section, we focus on issues in organizing a KPD program that have not been previously addressed, and explore alternative procedures for arranging kidney exchanges.

3.1 General utilities

Much of the prior work has focused on a simplified edge utility, namely $u_e = 1, \forall e \in E$. As a result, the utility U_c of an exchange cycle c equals its length, and the objective function in the IP problem is the total number of arrangeable transplants. This objective function, however, is restrictive especially when a major interest presumably lies in medical outcomes such as graft survival after transplant. Roth et al. (2007) actually used this simplified utility assignment when they first proposed the IP approach to the KPD optimization problem. And although Abraham et al. (2007) adopted a more general way of assigning edge utilities, or edge weights as referred to in their work, the efficiency of their proposed exact algorithm in solving the IP problem relies partially on using this simplified edge utility, as we have discussed in Section 2.2.

Some recent developments, however, have emerged in assigning more general edge utilities so that they better evaluate potential transplants. In an operational guideline recently posted by the U.S. national KPD pilot program, each potential transplant is initially assigned a base utility of 200 points. Extra points are added as bonuses to edges that, say, have zero antigen

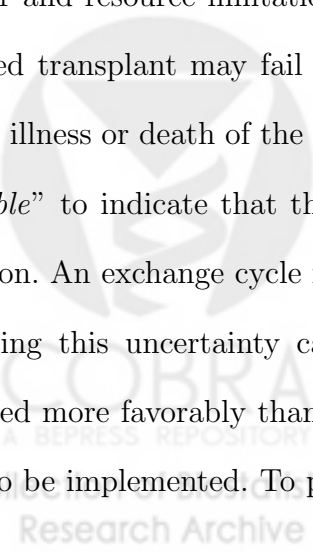
mismatches, or that involve a donor and a candidate in the same transplant center; on the other hand, a certain number of points are deducted, for example, when a donor has one or more of the candidate's other antibody specificities.

We propose to associate with edge (i, j) , a more general utility $u_{(i,j)}$, quantifying the medical outcome from a potential kidney transplant involving the donor in pair i and the candidate in pair j . This outcome could be, for example, graft survival, post-transplant recipient survival, or the incremental years of recipient life that would accrue with a kidney transplant as opposed to remaining on dialysis. Clearly, by incorporating this more general utility, we are able to evaluate and compare competing kidney allocations with a quality-oriented view, and provide kidney transplant candidates with organs that are not only compatible, but that could potentially lead to a good quality of life after transplants.

3.2 *Operational uncertainties in a KPD program*

Prior research on KPD has implicitly assumed that a predicted compatible transplant, if attempted, would yield an actual transplant. In reality, predicted compatible transplants have to be confirmed by laboratory crossmatches, and hence may or may not lead to actual transplant operations. This uncertainty is a necessary ingredient since laboratory crossmatches cannot be undertaken on all possibly compatible donors and candidates due to labor and resource limitations. Further, even if the laboratory crossmatch is negative, a proposed transplant may fail to occur due to other friction including, for example, donor refusal, illness or death of the candidate. Throughout the rest of the paper, we use the term “*is viable*” to indicate that the edge actually results, if chosen, in a completed transplant operation. An exchange cycle is viable if each of its edges is viable.

Ignoring this uncertainty can result in a situation in which long exchange cycles are evaluated more favorably than short ones, despite the fact that longer cycles are much less likely to be implemented. To partially incorporate this uncertainty into the arrangement of



kidney exchanges, the operational guidelines for the national KPD pilot program proposed a deduction of 30 points for a three-way exchange cycle, but not for an exchange cycle that is two-way.

To address this stochastic feature in a KPD program, we associate a probability with each edge corresponding to the chance of that edge being viable. For given $e \in E$, let X_e be a Bernoulli random variable with $X_e = 1$ if e is viable, and $X_e = 0$ otherwise. By letting $p_e = P(X_e = 1)$, a cycle c is viable with probability equal to $P_c = \prod_{e \in E(c)} p_e$, under the assumption that edges in an exchange cycle have an independence relationship. More formally, this assumption is regarded as having independent random variables, $X_e, e \in E(c)$. In the rest of this paper, we assume in general that $\{X_e, e \in E\}$ are independent.

As we should see later in Section 3.4, the incorporation of probabilities into the framework of a KPD program also opens up the opportunity to identify possible alternatives a fall back position that can be implemented if the primary choice does not lead to a completed set of exchanges.

3.3 Maximum expected-utility cycle-based allocation

A natural way to quantify the value of an exchange cycle c is to use its *expected utility*, namely $EU_c = U_c P_c$. This approach recognizes the fact that a longer exchange cycle has a smaller chance of being viable, which counterbalance the fact that such a cycle might potentially contribute greater utilities and allow more transplants. Further, the expected utility of a cycle-based allocation $\mathcal{C}(G)$, is the sum of the expected utilities of its exchange cycles, i.e. $\sum_{c \in \mathcal{C}(G)} EU_c$. Among all cycle-based allocations, the one with the largest expected utility is the *maximum expected-utility cycle-based* allocation, denoted as $\bar{\mathcal{C}}^*(G)$. The following procedure generates this $\bar{\mathcal{C}}^*(G)$ in a KPD program $G = (V, E)$, and arranges kidney exchanges accordingly.

- (i) Define $u : E \rightarrow R^+$, where u_e is the utility of an edge $e \in E$, and $p : E \rightarrow [0, 1]$, where p_e is the probability that e is viable.
- (ii) Enumerate C_k and calculate $EU_c = U_c P_c$ for each $c \in C_k$.
- (iii) Find $\bar{\mathcal{C}}^*(G)$ by an IP-based approach as in Section 2.2, with U_c replaced by EU_c .
- (iv) Arrange kidney exchanges according to $\bar{\mathcal{C}}^*(G)$.

In (iii), the cycle formulation works in a straightforward manner, but the edge formulation cannot be extended, since this formulation represents each edge as a separate decision variable and hence cannot describe EU_c . Notice that the above procedure and the one in Section 2.4 are both *fixed* in that they do not specify how to proceed if an exchange cycle $c \in \bar{\mathcal{C}}^*(G)$ or $c \in \mathcal{C}^*(G)$ is not viable.

3.4 Contingency plans

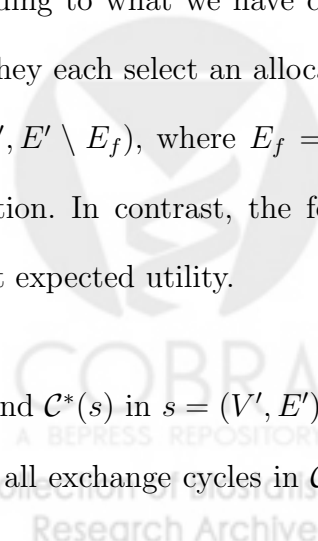
Let us begin with a motivating example. In a small KPD program as represented in Figure 2, there are two three-way exchange cycles, $c_1 = \langle v_1, v_2, v_3 \rangle$ and $c_2 = \langle v_2, v_4, v_5 \rangle$, with expected utilities EU_{c_1} and EU_{c_2} , and one two-way exchange cycle, $c_3 = \langle v_2, v_4 \rangle$, with expected utility EU_{c_3} . We find that $EU_{c_1} > EU_{c_2} > EU_{c_3}$, and thus conclude that $\bar{\mathcal{C}}^*(G)$ is $\{c_1\}$. A further examination, however, reveals that the incompatible pairs in c_3 are part of the pairs in c_2 , i.e. $V(c_3) = \{v_2, v_4\} \subset V(c_2) = \{v_2, v_4, v_5\}$. This observation might suggest that, when c_2 is selected but could not be completed because of problems in either (v_4, v_5) or (v_5, v_2) , the two-way exchange cycle c_3 could still be selected. Therefore, the contribution from this back-up exchange cycle c_3 would add some extra value to the exchange cycle c_2 . Does this extra value make $\{c_2\}$ a preferred allocation to $\{c_1\}$? How should we evaluate $\{c_2\}$ so as to correctly recognize $\{c_3\}$ as a possible back-up allocation? To address these questions in this specific example and address other related issues in general, we propose the following.

First, we give two definitions from graph theory: (i) a graph $G' = (V', E')$ is a *subgraph* of $G = (V, E)$, if $V' \subset V$ and $E' \subset E$, and (ii) a graph $G' = (V', E')$ is an *induced subgraph*

of $G = (V, E)$, if G' is a subgraph of G , and in addition, $E' = \{(u, v) \in E : u, v \in V'\}$. In the context of a KPD program $G = (V, E)$, we then define an *exchange set*, denoted by $s = (V', E')$, as an induced subgraph of G with the extra requirement that it allows at least one exchange cycle of length $|V'|$, where $|V'|$ is also defined as the *size* of this exchange set. In Figure 2 are shown two exchange sets of size 3, $s_1 = (\{v_1, v_2, v_3\}, \{(v_1, v_2), (v_2, v_3), (v_3, v_1)\})$ and $s_2 = (\{v_2, v_4, v_5\}, \{(v_2, v_4), (v_4, v_5), (v_5, v_2), (v_4, v_2)\})$, and one exchange set of size 2, $s_3 = (\{v_2, v_4\}, \{(v_2, v_4), (v_4, v_2)\})$.

By definition, an exchange set would allow one or more cycle-based allocations; if there is more than one and not all of the exchange cycles in the first attempted allocation are viable, we might still have the option to select another allocation, which is called a *contingency*. Therefore, the expected utility generated from arranging exchanges in an exchange set depends on the order in which the possible cycle-based allocations are utilized. Such an ordering defines a *general* procedure, which includes the two previously introduced procedures as special cases. For example, if $s = (V', E')$ and $\mathcal{C}^*(s)$ is adopted as top priority with no allocation assigned as a contingency, this in fact corresponds to the procedure presented in Section 2.4 and in consequence, the generated expected utility is $\sum_{c \in \mathcal{C}^*(s)} EU_c$; similarly, the procedure described in Section 3.3 corresponds to selecting $\bar{\mathcal{C}}^*(s)$ as a first priority but again with no contingency at all, and hence generates an expected utility of $\sum_{c \in \bar{\mathcal{C}}^*(s)} EU_c$. According to what we have discussed, both of these two procedures are fixed in the sense that they each select an allocation only for $s = (V', E')$, but do not specify how to proceed on $(V', E' \setminus E_f)$, where $E_f = \{e \in E' : X_e = 0\}$ is observed when executing that chosen allocation. In contrast, the following “greedy” procedure is *sequential* and generates the largest expected utility.

- (i) Find $\mathcal{C}^*(s)$ in $s = (V', E')$ using the procedure presented in Section 2.4.
- (ii) If all exchange cycles in $\mathcal{C}^*(s)$ are viable, finish with a claimed utility of $\sum_{c \in \mathcal{C}^*(s)} U_c$; if



certain edges, denoted as E_f , fail to be carried forward to actual transplants, then abort the original allocation and go back to (i) with $E' \leftarrow E' \setminus E_f$.

Return to Figure 2, where both the three-way exchange cycle, $c_2 = \langle v_2, v_4, v_5 \rangle$, and the two-way exchange cycle, $c_3 = \langle v_2, v_4 \rangle$, could be selected in the exchange set s_2 . If U_{c_2} is larger than U_{c_3} , c_2 should be chosen as the first priority. When edge (v_4, v_5) and/or edge (v_5, v_2) in c_2 are not viable, c_3 could be tried as a contingency plan. On the other hand, if U_{c_2} is less than U_{c_3} , then c_3 , backed up by c_2 , should be selected as the first priority.

In light of this and similar examples, we quantify the value of an exchange set s using the above greedy procedure, which defines EU_s , the *expected utility* of s . This approach takes into consideration the full-potential contributions from all back-up allocations. The expected utility of an exchange set can be exactly formulated as follows. For an exchange set $s = (V', E')$, let $2^{E'}$ denote the collection of all subsets of E' , and for each $\tilde{E} \in 2^{E'}$, let

$$P(\tilde{E}) = \prod_{e \in \tilde{E}} p_e \prod_{e \in (E' \setminus \tilde{E})} (1 - p_e),$$

which is the probability that the edges in \tilde{E} are viable whereas those in $E' \setminus \tilde{E}$ are not. For $\tilde{s} = (V', \tilde{E})$, let

$$U(\tilde{E}) = \sum_{c \in \mathcal{C}^*(\tilde{s})} U_c.$$

It follows that

$$EU_s = \sum_{\tilde{E} \in 2^{E'}} U(\tilde{E})P(\tilde{E}).$$

In Figure 2, s_2 would be preferred to s_1 if $EU_{s_2} > EU_{s_1}$.

3.5 The maximum expected-utility set-based allocation

In Section 2.1, we introduced a cycle-based allocation as a collection of disjoint exchange cycles. Among all such allocations, the procedure in Section 2.4 arranges kidney exchanges according to the one with the maximum utility while the procedure in Section 3.3 according to the one with the maximum expected utility.

In this subsection, we define a set-based allocation as a collection of disjoint exchange sets, and further denote it as $\mathcal{S}(G)$ for a KPD program G . Following the way in which the expected utility of an exchange set is defined in Section 3.4, the expected utility of a set-based allocation $\mathcal{S}(G)$ is, therefore, $\sum_{s \in \mathcal{S}(G)} EU_s$. Among all set-based allocations, the following procedure arranges kidney exchanges in a KPD program $G = (V, E)$ according to the one with the maximum expected utility, which is denoted as $\overline{\mathcal{S}}^*(G)$.

- (i) Define $u : E \rightarrow R^+$ and $p : E \rightarrow [0, 1]$ as in Section 3.3.
- (ii) Enumerate S_k , the set of all exchange sets of size at most k , where $2 \leq k \leq |V|$.
- (iii) For each $s \in S_k$, calculate its expected utility as $EU_s = \sum_{\tilde{E} \in 2^{E'}} U(\tilde{E})P(\tilde{E})$.
- (iv) Select $\overline{\mathcal{S}}^*(G)$ by forming an IP problem similar to the one discussed in Section 2.2.
- (v) Apply the aforementioned greedy procedure to each $s \in \overline{\mathcal{S}}^*(G)$.

Several remarks on the above procedure are in place. In (ii), a reasonable k , say three or four, is required in practice due to logistical concerns as in Section 1. Enumerating S_k could be accomplished by a DFS-based algorithm similar to the one presented in Section 2.4. In (iii), calculating EU_s involves a summation over $|2^{E'}|$ terms, which poses no computational difficulties in practice for small k .

As we have discussed, arranging kidney exchanges according to $\overline{\mathcal{S}}^*(G)$ allows a more flexible utility-based evaluation of potential transplants, takes into consideration the uncertainties in a KPD, and provides contingency options when possible.

3.6 Estimation of utilities and probabilities

So far in this paper, we have assumed that we are given a utility function $u : E \rightarrow R^+$ and a probability function. In practice, however, these utilities and probabilities are not available and have to be estimated.

In the literature, modeling of outcome-based utility has been considered in deceased-donor kidney transplants by Wolfe et al. (2008). Such models could be adopted comparatively

easily to living-donor kidney transplants. On the other hand, the model for probability can be established through a logistic modeling approach, based on clinical data from multiple KPD programs, including the Scientific Registry of Transplantation Recipients (SRTR), the Alliance for Paired Donation (APD), and the University of Michigan Transplant Center. In this logistic model, some primary predictors include, for example, percentage peak panel reactive antibody (PRA), and cross reactivity of antibody specificities.

Thus, in practice, all three procedures previously discussed for selecting the cycle-based or set-based allocations can be easily adopted by replacing the u and p with estimated \hat{u} and \hat{p} .

3.7 A KPD match run and an evolving KPD program

We consider a *match run* as a series of operations on a collection of incompatible pairs V :

- (i) Form $G = (V, E)$, where E is determined by checking virtual crossmatches on V .
- (ii) Assign each $e \in E$ with an estimated utility \hat{u}_e and an estimated probability \hat{p}_e according to the models discussed in Section 3.6.
- (iii) Arrange kidney exchanges according to $\mathcal{C}^*(G)$, $\bar{\mathcal{C}}^*(G)$, or $\bar{\mathcal{S}}^*(G)$.
- (iv) Recycle any donor-candidate pair that does not proceed to an actual transplant back to the KPD pool awaiting for future matches.

Every KPD program is constantly evolving in that successfully transplanted pairs leave and new incompatible pairs arrive over time. In addition, existing pairs in the pool could withdraw due to factors such as donor or candidate pregnancy, illness, or death. Such an evolving KPD program is managed by repeatedly executing match runs on a regular basis over time. See Figure 3 for an illustration.

4. Microsimulation system

In Sections 2 and 3, we have presented three allocation strategies for managing kidney exchanges, namely, $\mathcal{C}^*(G)$, $\bar{\mathcal{C}}^*(G)$, and $\bar{\mathcal{S}}^*(G)$. A comparison among them and other possible allocations is of high importance with respect to the practical management of a KPD program. Such a comparison, however, usually cannot be performed in the traditional clinical trials due to the nature of kidney transplantation, or any transplantation for that matter. Thus, microsimulation plays an important role for this purpose. In this section, we discuss the key ingredients of such a microsimulation system.

4.1 Generating incompatible pairs

To create an incompatible pair, we generate its donor and candidate separately according to their own population distributions. Candidates are *sampled at random with replacement* from databases of candidates presenting with a willing but incompatible donor. One of such databases is derived from the University of Michigan KPD program, in which candidates are measured by variables of blood type, PRA, candidate antibody specificities (with mean fluorescence intensity (MFI) > 5000), and so on. Currently, the UM database consists of 187 incompatible pairs, and additional databases from other KPD programs (as they become available to us through data-usage agreements) will be incorporated for better variation in candidates. Donors, on the other hand, are generated by separately sampling their blood types and HLA haplotypes. Precisely, blood types are drawn from the U.S. population distribution: O, 44%; A, 42%; B, 10%; and AB, 4% (Stanford Blood Center, 2010). HLA haplotypes are sampled according to their frequencies in the U.S. population, which is derived from an extensive public database on potential bone marrow donors (Maiders et al., 2007).

A simulated donor-candidate pair is regarded as an incompatible pair and hence included in the KPD pool if either their ABO blood types mismatch, or the donor's HLA haplotypes contain any of the candidate's antibody specificities (with MFI > 5000), or both. Note that

only incompatible pairs enter the KPD pool since compatible ones can directly go for a transplant.

4.2 *Simulating a match run and an evolving KPD program*

We simulate a KPD match run on a collection of incompatible pairs V by simulating each of its four steps, as discussed in Section 3.7.

- (i) Determine E according to the simulated blood types and HLA haplotypes.
- (ii) Estimate both u_e and p_e based on generated pair characteristics.
- (iii) Apply $\mathcal{C}^*(G)$, $\bar{\mathcal{C}}^*(G)$, or $\bar{\mathcal{S}}^*(G)$ to arrange kidney exchanges.
- (iv) Since a potential transplant may not be viable, we simulate such uncertainty via a Bernoulli trial with the probability of success equal to that edge probability. The realization of such a Bernoulli trial will indicate if a pair proceeds to an actual transplant and hence leaves the pool, or remains in the pool and awaits future matches.

To address the feature that a KPD program is evolving over time as discussed in Section 3.7, we first generate an initial KPD pool of N incompatible pairs as described in Section 4.1. Further, we assume that the arrival of new incompatible pairs follows a Poisson process with a rate λ . This rate may be governed by a log-linear model with covariates of age, race, and the relationship between donor and candidate, among others. Also, we assume that the withdrawal of existing pairs follows another Poisson process with a rate μ ; and the log-linear model for μ could include other relevant covariates.

5. Simulation results

The proposed microsimulation system in Section 4 enables us to investigate several aspects of KPD, among which comparing different kidney allocation strategies is of special interest. This section presents these results based on simulation studies. We implement the microsimulation

models using the C++ programming language, and the related IP problems are solved by Gurobi Optimizer.¹

5.1 Problem complexity

The number of exchange cycles or exchange sets can be enormous in a reasonably large KPD pool of, say, several hundred pairs, even if the size of exchange cycles or exchange sets is restricted to at most three. Such complexity causes solving related IP problems computationally rather expensive. Table 1 summarizes the averaged numbers (over 200 rounds of simulation) of exchange cycles or exchange sets up to three pairs in a KPD pool whose size varies over 100, 200, 300, 400, and 500, where incompatible pairs are generated according to Section 4.1.

[Table 1 about here.]

5.2 Simulation setup

We perform a total of 600 simulations, in which $\{p_e, e \in E\}$ is generated according to a uniform distribution $U(0.1, 0.5)$. We use two other uniform distributions, $U(10, 20)$ and $U(20, 30)$, to generate $\{u_e, e \in E\}$, each corresponding to 200 rounds of simulation; the utility is fixed at 1 in the remaining 200 simulations.

We start each round of simulation by generating a pool of $N = 200$ incompatible pairs. Additionally generated pairs then enter this pool according to a Poisson process with $\lambda = 10$ pairs per month over a period of $m = 24$ months. For simplicity, we assume no existing pairs drop out of the pool.

We execute a match run at the end of each month on this evolving pool starting with the initial pool. Pairs that arrive during the time of a match run will not participate but wait for the next match run. At each simulation, we make three copies of the evolving

¹Gurobi Optimizer Version 2.0. Houston, Texas: Gurobi Optimization, Inc.

KPD program, and execute each of the three match runs (which arrange kidney exchanges according to $\mathcal{C}^*(G)$, $\bar{\mathcal{C}}^*(G)$, and $\bar{\mathcal{S}}^*(G)$) on each copy such that we can directly compare these three allocation strategies. At the end of each match run, the KPD pool is updated with some pairs leaving or staying. We record several important measures needed for comparison, such as cumulative claimed utilities, cumulative number of transplants, and blood types of candidates receiving transplants.

The above setup provides a simple microsimulation system that enables us to examine and compare different allocation strategies. More realistic models will be further explored by incorporating more comprehensive data on pair characteristics and actual transplant operations, when such data become available to us via data-usage agreements.

5.3 Results

First, we report on the cumulative number of transplants over a period of 24 months across three allocation strategies, $\mathcal{C}^*(G)$, $\bar{\mathcal{C}}^*(G)$, and $\bar{\mathcal{S}}^*(G)$, and under three different models of utilities, i.e. $U(1, 1)$, $U(10, 20)$, and $U(10, 30)$; see Figure 4, which unveils a consistent pattern regardless of utility models that $\bar{\mathcal{S}}^*(G)$ results in the greatest number of transplants, whereas $\mathcal{C}^*(G)$ leads to the fewest number of transplants. Take Figure 4-C as an example. When the edge utility is fixed at 1, allocation strategy $\bar{\mathcal{S}}^*(G)$ gives 50 (median over 200 simulations) transplants after match run 10 (at the end of month 9); in contrast, for exactly the same evolving KPD program, strategy $\mathcal{C}^*(G)$ only leads to a median of 34 completed transplants over the same period of time. Also, allocation strategy $\bar{\mathcal{C}}^*(G)$ allows for a significantly higher number of transplants than $\mathcal{C}^*(G)$ does, though it performs worse than $\bar{\mathcal{S}}^*(G)$.

[Figure 4 about here.]

Figure 5 demonstrates that $\bar{\mathcal{S}}^*(G)$ is advantageous over both $\mathcal{C}^*(G)$ and $\bar{\mathcal{C}}^*(G)$, in the sense that $\bar{\mathcal{S}}^*(G)$ on average achieves the largest cumulative claimed utility. Notice that when edge utilities are fixed at 1, the cumulative claimed utility is the same as the cumulative number of

transplants; in such scenario, the superiority of $\bar{\mathcal{S}}^*(G)$ has been shown in Figure 4, and hence we only compare these allocation strategies under two other utility models, i.e. $U(10, 20)$ and $U(10, 30)$. It is clear in Figure 5 that the strategy of $\bar{\mathcal{S}}^*(G)$ outperforms the other two.

[Figure 5 about here.]

We can also examine via the proposed microsimulation system other aspects of a KPD program. For example, we are interested in exploring how the chance of having a transplant being associated with blood types. In practice, candidates with blood type O are usually at a disadvantageous position due to the limitation that they can only receive kidneys from blood type O donors, who, however, can donate to candidates of any blood type. This phenomenon is clearly observed in Figure 6, where about 60% of the incoming candidates are of blood type O while only about 40% of the performed transplants involve a blood O candidate. As a consequence, candidates of the other blood types (A, B, and AB) are more represented among all candidates receiving a transplant. One possible solution to this difficulty is to assign bonus utility to a potential transplant that involves an O donor giving to an O candidate.

[Figure 6 about here.]

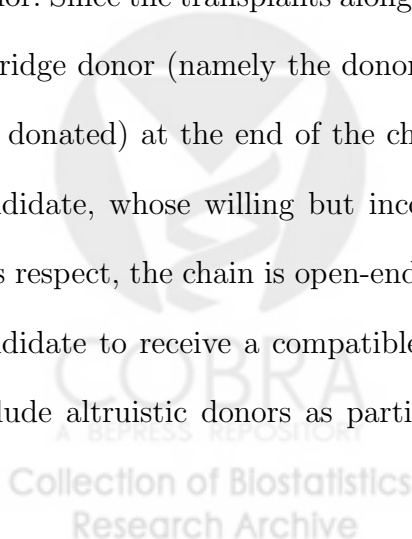
6. Concluding remarks

In this paper, we have proposed a novel approach to arranging kidney exchanges in an evolving KPD program. Our approach identifies the maximum expected-utility set-based allocation that (i) allows for a more flexible utility-based evaluation of potential kidney transplants, (ii) takes into consideration stochastic features in managing a KPD program, and (iii) exploits possible back-up exchanges when the originally planned allocation cannot be fully executed. Another primary contribution is rooted in the development of a comprehensive microsimulation system that enables us to simulate and examine various aspects in an evolving KPD program. This microsimulation system allows us to emulate genetic and

demographic data from existing KPD programs, and to derive statistical models similar to the actual KPD program in practice. In particular, we have suggested (i) models for donors' and candidates' characteristics as well as for their arrival in and withdrawal from a KPD pool; (ii) models for the estimation of the outcome-based utility of a potential transplant; and (iii) models for the prediction of the probability that a planned transplant would indeed occur. Utilizing such a microsimulation system, we are able to quantitatively compare different kidney allocation strategies; and results shed light on decision support in actual KPD programs.

We have illustrated the proposed microsimulation system to compare several kidney allocation strategies. Through simulation studies, we demonstrate advantages of the maximum expected-utility set-based allocation over the other two allocations. Such advantages are attributable to the concept of an exchange set, in which uncertainties in a KPD are properly incorporated and possible contingency allocations are allowed in the case of failure at the original planned exchange. In the future work, we plan to base our simulation on more realistic models that will be developed by incorporating more KPD program source data.

Another possible future work is to consider exchanges initiated by an altruistic donor (Rees et al., 2009). An altruistic donor does not have a designated candidate and donates a kidney voluntarily. Figure 7 gives an illustration on a chain of transplants initiated by an altruistic donor. Since the transplants along the chain are not required to be performed simultaneously, a bridge donor (namely the donor whose incompatible candidate received a kidney but has yet donated) at the end of the chain could make a donation to a future-arrival compatible candidate, whose willing but incompatible donor then becomes the new bridge donor. In this respect, the chain is open-ended and greatly increases the chance for a highly sensitized candidate to receive a compatible kidney. The proposed work could be easily extended to include altruistic donors as participants in a KPD program. More specifically, a chain of



kidney transplants can be viewed as a special exchange cycle, in which the bridge donor “donates” to a phantom candidate associated with the altruistic donor who first initiated that chain, and hence form a hypothetical “edge” (denoted by dashed arrowed lines in Figure 7). Future research on KPD with altruistic donors is certainly of high importance and great interest; we will report our results in a separate publication in future.

[Figure 7 about here.]

ACKNOWLEDGEMENTS

Dr. Yanhua Chen has provided invaluable advice and comment on this work. The authors are grateful for the support from the Michigan Institute for Clinical and Health Research (MICHR), Michigan School of Public Health, Scientific Registry of Transplant Recipients, and National Institute of General Medical Sciences (NIGMS).

REFERENCES

- Abraham, D. J., Blum, A., and Sandholm, T. (2007). Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. *Ec'07: Proceedings of the Eighth Annual Conference on Electronic Commerce* pages 295–304 370.
- Edmonds, J. (1965). Paths, trees, and flowers. *Canadian Journal of Mathematics* **17**, 449–467.
- Evans, R. W., Manninen, D. L., Garrison, L. P., Hart, L. G., Blagg, C. R., Gutman, R. A., Hull, A. R., and Lowrie, E. G. (1985). The quality of life of patients with end-stage renal disease. *New England Journal of Medicine* **312**, 553–559.
- Hariharan, S., Johnson, C. P., Bresnahan, B. A., Taranto, S. E., McIntosh, M. J., and Stablein, D. (2000). Improved graft survival after renal transplantation in the united states, 1988 to 1996. *New England Journal of Medicine* **342**, 605–612.
- Keizer, K. M., de Klerk, M., Haase-Kromwijk, B. J. J. M., and Weimar, W. (2005). The dutch

- algorithm for allocation in living donor kidney exchange. *Transplantation Proceedings* **37**, 589–591.
- Laupacis, A., Keown, P., Pus, N., Krueger, H., Ferguson, B., Wong, C., and Muirhead, N. (1996). A study of the quality of life and cost-utility of renal transplantation. *Kidney International* **50**, 235–242.
- Maiers, M., Gragert, L., and Klitz, W. (2007). High-resolution hla alleles and haplotypes in the united states population. *Human Immunology* **68**, 779–788.
- Park, K., Moon, J. I., Kim, S. I., and Kim, Y. S. (1999). Exchange donor program in kidney transplantation. *Transplantation* **67**, 336–338.
- Rapaport, F. T. (1986). The case for a living emotionally related international kidney donor exchange registry. *Transplantation Proceedings* **18**, 5–9.
- Rees, M. A., Kopke, J. E., Pelletier, R. P., Segev, D. L., Rutter, M. E., Fabrega, A. J., Rogers, J., Pankewycz, O. G., Hiller, J., Roth, A. E., Sandholm, T., Ünver, M. U., and Montgomery, R. A. (2009). A nonsimultaneous, extended, altruistic-donor chain. *New England Journal of Medicine* **360**, 1096–1101.
- Ross, L. F., Rubin, D. T., Siegler, M., Josephson, M. A., Thistlethwaite, J. R., and Woodle, E. S. (1997). Ethics of a paired-kidney-exchange program. *New England Journal of Medicine* **336**, 1752–1755.
- Roth, A. E., Sönmez, T., and Ünver, M. U. (2004). Kidney exchange. *Quarterly Journal of Economics* **119**, 457–488.
- Roth, A. E., Sönmez, T., and Ünver, M. U. (2005). A kidney exchange clearinghouse in new england. *American Economic Review* **95**, 376–380.
- Roth, A. E., Sönmez, T., and Ünver, M. U. (2007). Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *American Economic Review* **97**, 828–851.

- Russell, J. D., Beecroft, M. L., Ludwin, D., and Churchill, D. N. (1992). The quality of life in renal transplantation—a prospective study. *Transplantation* **54**, 656–660.
- Segev, D. L., Gentry, S. E., Warren, D. S., Reel, B., and Montgomery, R. A. (2005). Kidney paired donation and optimizing the use of live donor organs. *Journal of the American Medical Association* **293**, 1883–1890.
- Shapley, L. and Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics* **1**, 23–37.
- Terasaki, P. I., Cecka, J. M., Gjertson, D. W., and Takemoto, S. (1995). High survival rates of kidney-transplants from spousal and living unrelated donors. *New England Journal of Medicine* **333**, 333–336.
- Wolfe, R. A., Ashby, V. B., Milford, E. L., Ojo, A. O., Ettenger, R. E., Agodoa, L. Y. C., Held, P. J., and Port, F. K. (1999). Comparison of mortality in all patients on dialysis, patients on dialysis awaiting transplantation, and recipients of a first cadaveric transplant. *New England Journal of Medicine* **341**, 1725–1730.
- Wolfe, R. A., McCullough, K. P., Schaubel, D. E., Kalbfleisch, J. D., Murray, S., Stegall, M. D., and Leichtman, A. B. (2008). Calculating life years from transplant (lyft): Methods for kidney and kidney-pancreas candidates. *American Journal of Transplantation* **8**, 997–1011.



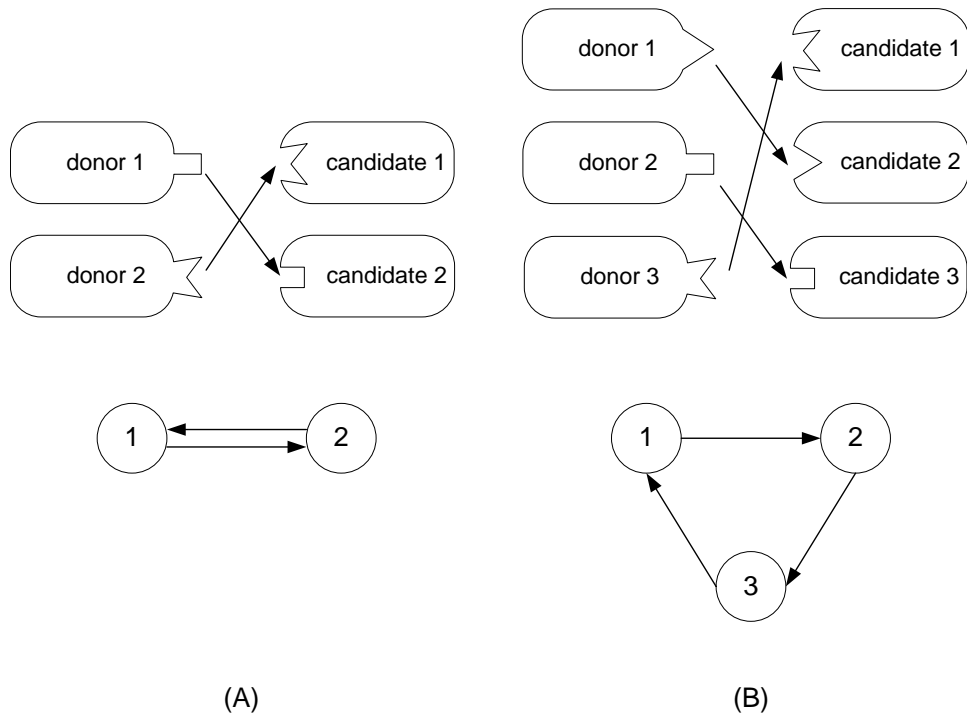


Figure 1. (A): A two-way exchange cycle; (B): A three-way exchange cycle. The two graphs in the bottom are the corresponding graphical representations of the top ones.




Figure 2. One two-way exchange cycle, $\langle v_2, v_4 \rangle$, and two three-way exchange cycles, $\langle v_1, v_2, v_3 \rangle$ and $\langle v_2, v_4, v_5 \rangle$.



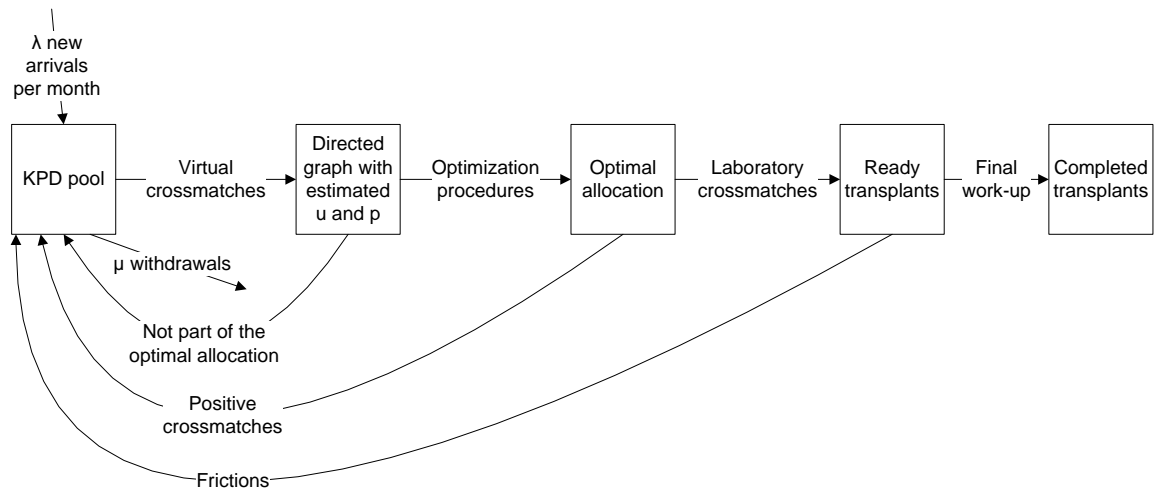


Figure 3. A flow diagram of an evolving KPD program

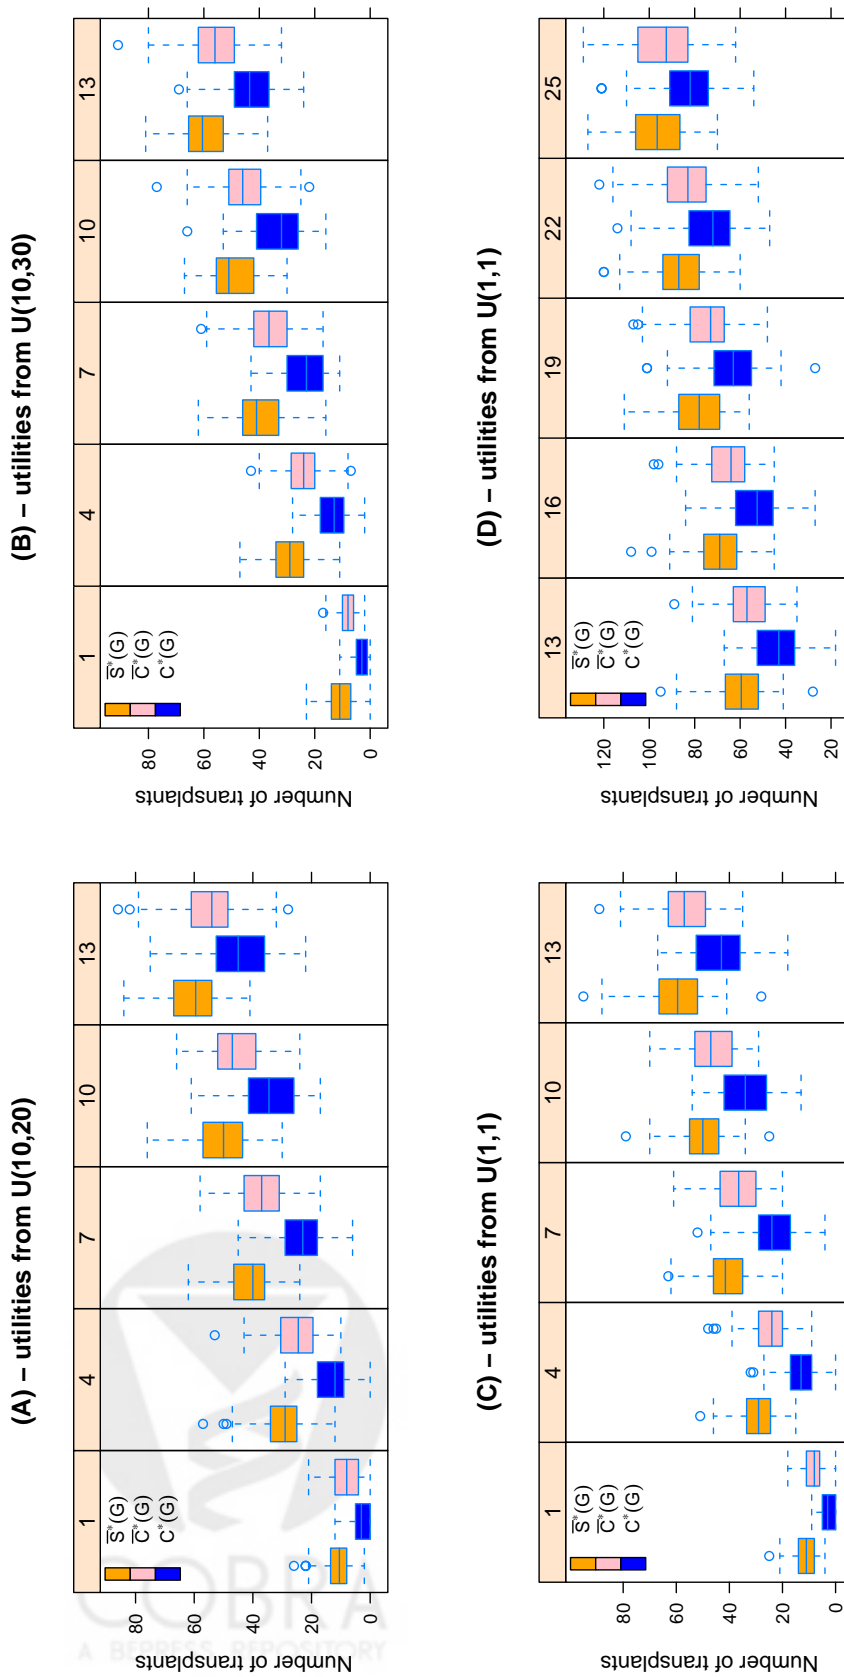


Figure 4. Cumulative number of transplants over 12 or 24 months across three allocation strategies, under three utility generating mechanisms (results are based on 200 simulations). Edge utilities are generated from $U(1, 1)$, $U(10, 20)$, and $U(10, 30)$; and edge probabilities are generated from $U(0.1, 0.5)$. The numbers at the top of each panel indicate match runs.

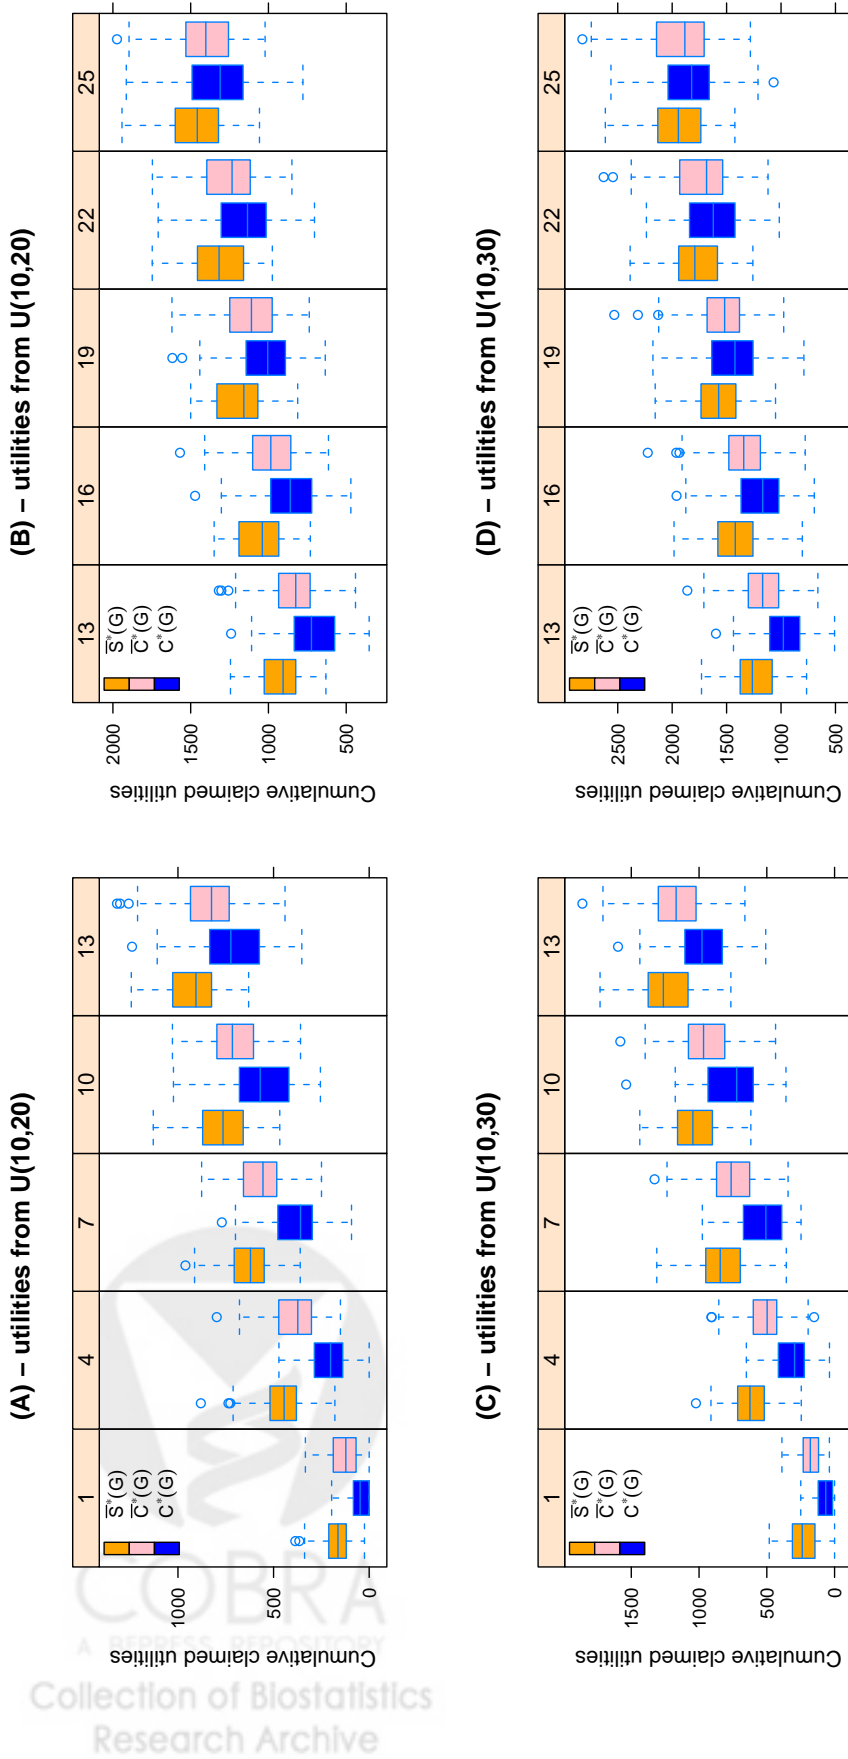


Figure 5. Cumulative claimed utilities of transplants over 24 months across three allocation strategies, under two utility generating mechanisms (results are based on 200 simulations). Edge utilities are generated from $U(10, 20)$, and $U(10, 30)$; and edge probabilities are generated from $U(0.1, 0.5)$. The numbers at the top of each panel indicates match runs.

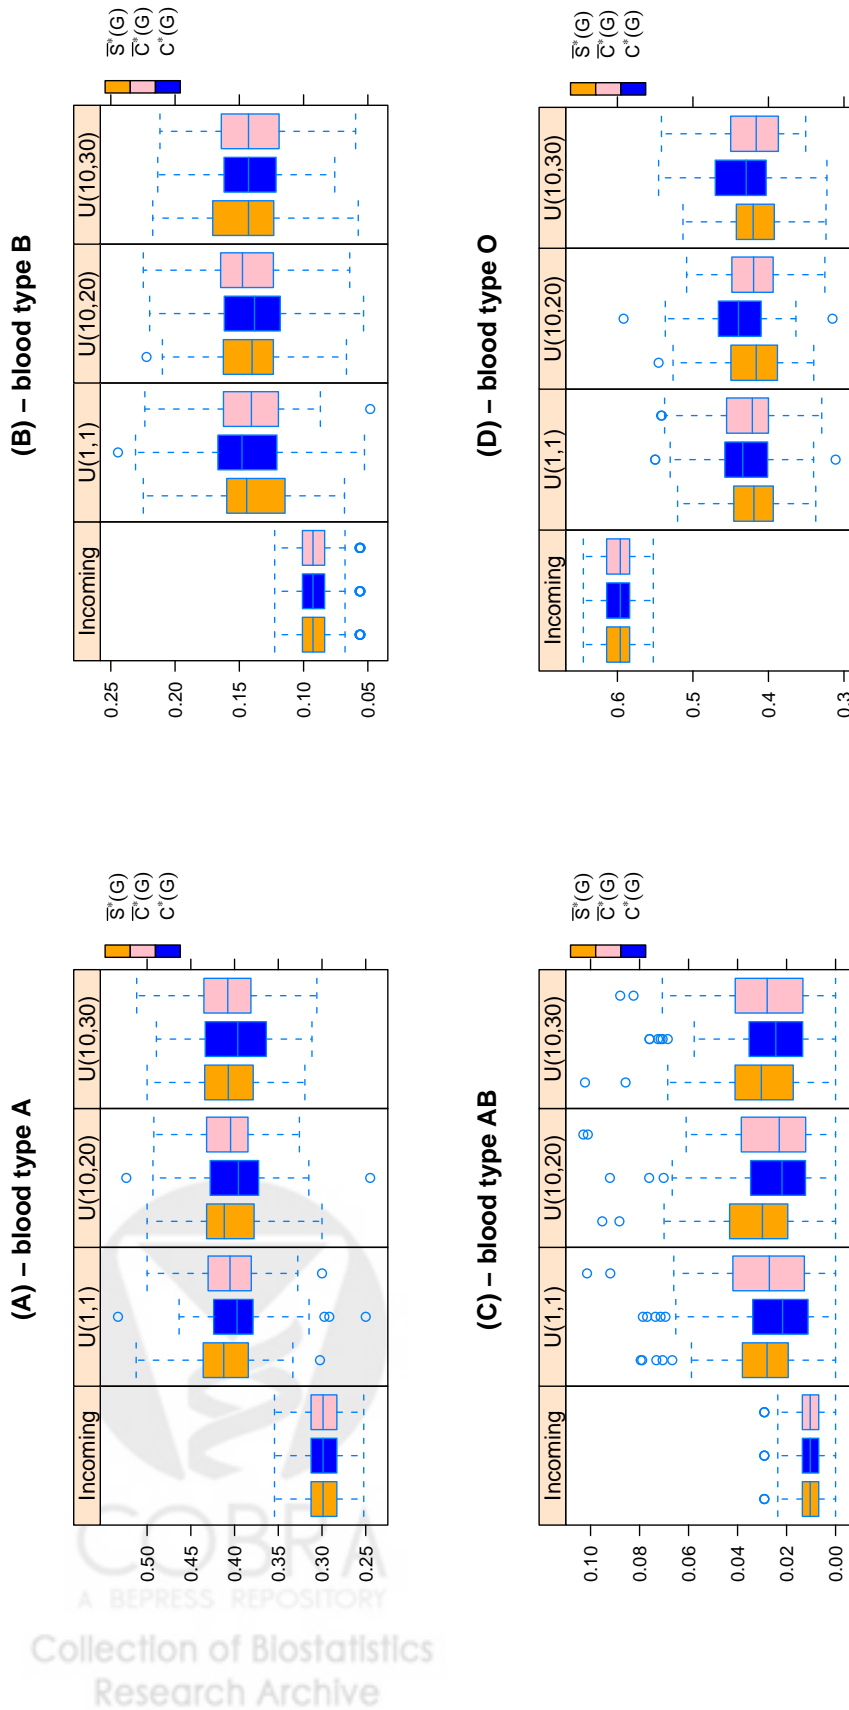


Figure 6. Summaries of the percentage of candidates with certain blood type among all candidates joining the KPD pool and that percentage among all candidates who received a transplant (results are based on 200 simulations). Similar patterns are observed under three utility generating mechanisms, i.e. $U(1, 1)$, $U(10, 20)$, and $U(10, 30)$.

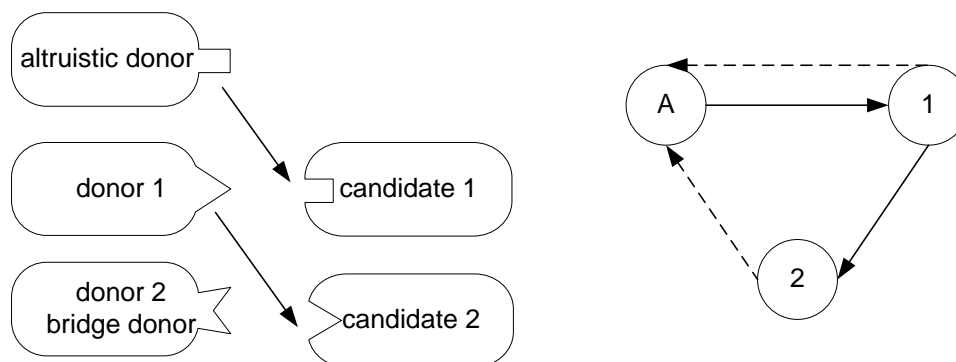


Figure 7. A chain of transplants initiated by an altruistic donor. The dashed arrow indicates a hypothetical “edge”, which represents a “donation” to a phantom candidate associated with the altruistic donor.

Table 1

The averaged numbers of exchange cycles or exchange sets up to three pairs in KPD pools of various sizes; standard deviations are given in the parentheses; the summary is calculated over 200 rounds of simulation.

pool size	exchange cycles (length 2 & 3) mean (standard deviation)	exchange sets (size 2 & 3) mean (standard deviation)
100	388 (237)	383 (229)
200	2659 (998)	2630 (977)
300	9413 (3164)	9305 (3100)
400	21076 (6140)	20829 (5992)
500	40290 (9337)	39815 (9120)

