

Paired and Altruistic Kidney Donation in the UK: Algorithms and Experimentation

DAVID F. MANLOVE and GREGG O'MALLEY, University of Glasgow

We study the computational problem of identifying optimal sets of kidney exchanges in the UK. We show how to expand an integer programming-based formulation due to Roth et al. [2007] in order to model the criteria that constitute the UK definition of optimality. The software arising from this work has been used by the National Health Service Blood and Transplant to find optimal sets of kidney exchanges for their National Living Donor Kidney Sharing Schemes since July 2008. We report on the characteristics of the solutions that have been obtained in matching runs of the scheme since this time. We then present empirical results arising from experiments on the real datasets that stem from these matching runs, with the aim of establishing the extent to which the particular optimality criteria that are present in the UK influence the structure of the solutions that are ultimately computed. A key observation is that allowing four-way exchanges would be likely to lead to a moderate number of additional transplants.

Categories and Subject Descriptors: G.1.6 [Numerical Analysis]: Optimization—*Integer programming*; G.2.1 [Discrete Mathematics]: Combinatorics—*Combinatorial algorithms*; G.2.2 [Discrete Mathematics]: Graph Theory—*Graph algorithms*; G.2.3 [Discrete Mathematics]: Applications; G.4 [Mathematical Software]: Algorithm design and analysis

General Terms: Algorithms; Design; Experimentation; Measurement

Additional Key Words and Phrases: Domino paired donation chain, integer programming, kidney paired exchange, multi-objective lexicographic optimization problem

ACM Reference Format:

David F. Manlove and Gregg O'Malley. 2014. Paired and altruistic kidney donation in the UK: Algorithms and experimentation. *ACM J. Exp. Algor.* 19, 2, Article 2.6 (December 2014), 21 pages.
DOI: <http://dx.doi.org/10.1145/2670129>

1. INTRODUCTION

Transplantation is the most effective treatment currently known for kidney failure. In the UK alone, as of March 31, 2012, there were 6,633 patients waiting on the

A preliminary version of this article appeared in the Proceedings of SEA 2012: the 11th International Symposium on Experimental Algorithms.

This research was financially supported by the National Health Service Blood and Transplant (NHSBT) between April 1, 2010 and June 30, 2011, and between January 1, 2012 and June 30, 2013 (project 10-11-ODT). We gratefully acknowledge collaboration with Rachel Johnson, Joanne Allen, Iain Harrison and with other staff at NHSBT. The views expressed here are those of the authors and not necessarily those of NHSBT. This research was also financially supported by the Engineering and Physical Sciences Research Council through a Knowledge Transfer Account between July 1, 2011 and December 31, 2011, and through grant EP/K010042/1 between June 6, 2013 and June, 5 2016. We would also like to thank Péter Biró, Rob Irving, Kirstin MacDonald and Ana Viana for valuable input into this work. Finally, we would like to thank the reviewers for their careful reading of the article and valuable comments, which have helped to improve the presentation of the article.

Authors' address: School of Computing Science, University of Glasgow, Glasgow G12 8QQ, UK.

Corresponding author's email address: david.manlove@glasgow.ac.uk.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2014 ACM 1084-6654/2014/12-ART2.6 \$15.00

DOI: <http://dx.doi.org/10.1145/2670129>

transplant list for a donor kidney, with the median waiting time being 1,168 days for an adult and 354 days for a child (based on patient registrations between April 1, 2005 and March 31, 2009). Kidneys used for transplantation can come from both deceased and living donors. In the UK, around 36% of all kidney transplants between April 1, 2011 and March 31, 2012 came from living donors [NHS Blood and Transplant 2012].

It is often the case that a patient requiring a kidney transplant has a willing donor, but due to blood- and/or tissue-type incompatibilities, the transplant cannot take place. However, in the UK, the Human Tissue Act 2004 and the Human Tissue (Scotland) Act 2006 (HTA) introduced, among other things, the legal framework required to allow the transplantation of organs between donors and recipients with no genetic or emotional connection.

With the introduction of the HTA, a patient with an incompatible donor can now “swap” his donor with that of another patient in a similar position via “kidney exchanges” that involve two or more incompatible donor–recipient pairs. For example, a *pairwise (kidney) exchange* involves two incompatible donor–recipient pairs (d_1, r_1) and (d_2, r_2) , where d_1 is compatible with r_2 , and d_2 is compatible with r_1 : d_1 donates a kidney to r_2 in exchange for d_2 donating a kidney to r_1 . *Three-way exchanges* extend this concept to three pairs in a cyclic manner.

In a number of countries, centralized programs (also known as *kidney exchange matching schemes*) have been introduced to help optimize the search for kidney exchanges. These include the United States [Roth et al. 2004, 2005, 2007; Abraham et al. 2007; Alliance for Paired Donation 2012; National Kidney Registry 2013; United Network for Organ Sharing 2013], the Netherlands [Keizer et al. 2005; de Klerk et al. 2005; Glorie and van de Klundert 2012], South Korea [Kwak et al. 1999; Park et al. 1999, 2004; Kim et al. 2007], and Romania [Lucan et al. 2003; Lucan 2007].

Following the introduction of the HTA, in early 2007, the UK established what has now become the National Living Donor Kidney Sharing Schemes (NLDKSS), administered by the National Health Service Blood and Transplant (NHSBT) (formerly UK Transplant) [Johnson et al. 2008]. The purpose of the NLDKSS is twofold: first, to identify those pairs that are compatible with one another and then, subsequently, to optimize the selected set of kidney exchanges subject to certain criteria. It is the responsibility of NHSBT (and in particular its Kidney Advisory Group) to supply the scoring system that is used to measure the benefit of potential transplants and the optimality criteria for the selection of kidney exchanges.

In general, it is seen as logistically challenging to carry out the transplants involved in a kidney exchange when the number of pairs involved in a single such exchange is large. This is mainly because all operations have to be performed simultaneously due to the risk of a donor renegeing on her commitment to donate a kidney after her loved one has received a kidney. As well, longer chains involve more participants and therefore have a higher risk that the whole chain will break down if one of the donors or recipients involved in the cycle becomes ill. Mainly for these reasons, at present, the NLDKSS does not allow exchanges involving more than three pairs.

A kidney exchange matching scheme may also include *altruistic* (or *nondirected*) donors [Roth et al. 2006] who do not have an associated recipient and who are willing to donate a kidney to a stranger. An altruistic donor d_0 can either donate directly to a recipient (without a donor) on the Deceased Donor Waiting List (DDWL) or else can trigger a *Domino Paired Donation chain* (DPD chain) [Montgomery et al. 2006; Roth et al. 2006; Ashlagi et al. 2011] involving one or more incompatible donor–recipient pairs: here d_0 donates to a recipient r_1 in exchange for r_1 's donor donating to the recipient r_2 in the next pair in the chain, and so on, with the final donor donating to the DDWL. According to NHSBT terminology, a DPD chain is *short* (resp. *long*) if it consists of one (resp. two) incompatible donor–recipient pairs. We remark that not all

the transplants on a DPD chain need to be carried out simultaneously because if a donor reneges, the next recipient in the chain should still have his or her willing donor and can participate in the next matching run (this is not necessarily true in the case of a cycle: with nonsimultaneous transplants, if a donor reneges, a recipient may fail to receive a kidney, and, moreover, his or her willing donor may already have had a nephrectomy). At present, the NLDKSS allows short but not long chains. The concept of a DPD chain may be generalized to a *Nonsimultaneous Extended Altruistic Donor (NEAD) chain* [Rees et al. 2009; Ashlagi et al. 2011, 2012; Dickerson et al. 2012b]—these are not currently permitted in the NLDKSS but are discussed in further detail in Section 6.

Within the past decade, kidney exchange has received considerable attention in the computer science and operations research [Abraham et al. 2007; Anderson et al. 2012; Ashlagi et al. 2010; Ashlagi and Roth 2011; Awasthi and Sandholm 2009; Biró et al. 2009; Caragiannis et al. 2011; Constantino et al. 2013; Dickerson et al. 2012a, 2012b, 2013; Pedroso 2014; Toulis and Parkes 2011], economics [Ashlagi et al. 2012; Ashlagi and Roth 2012; Roth et al. 2004, 2005, 2007; Ünver 2010], and medical [Ashlagi et al. 2011; Chen et al. 2012; Gentry et al. 2007; Johnson et al. 2008; Keizer et al. 2005; Kim et al. 2007; de Klerk et al. 2005; Li et al. 2014; Lucan 2007; Montgomery et al. 2006; Park et al. 2004; Rees et al. 2009; Roth et al. 2006; Saidman et al. 2006; Segev et al. 2005] literature, with many of these papers being interdisciplinary in nature. It has been observed that when only pairwise exchanges are permitted, an optimal solution can usually (depending, of course, on the definition of optimality) be found in polynomial time using maximum weight matching in a general graph (see, e.g., Biró et al. [2009] for more details). However, when pairwise and three-way exchanges are allowed, the problem of finding a set of exchanges that maximizes the number of transplants is NP-hard [Abraham et al. 2007] and indeed APX-hard [Biró et al. 2009].

Roth et al. [2007] described two Integer Programming (IP)-based formulations of the problem of finding a maximum weight set of kidney exchanges when both pairwise and three-way exchanges are permitted (here, the weights can measure the benefit of potential transplants). Abraham et al. [2007] showed that, due to scaling issues with the first of these models (the so-called *edge formulation*), the second model (the so-called *cycle formulation*) is the preferred way to model the problem using an IP. They implemented the cycle formulation using a branch-and-price strategy [Barnhart et al. 1998], showing that their approach could clear markets with up to 10,000 donor-recipient pairs.

1.1. Contribution of This Article

In this article, we present an application-driven case study showing how the cycle formulation can be extended to handle kidney exchange in the UK. In particular, we show how to model a complex, hierarchical set of criteria that form the definition of an optimal set of kidney exchanges. These criteria are quite natural yet specific to the requirements of the UK scheme. We have implemented the technique, and it has been used by the NHSBT to find optimal sets of kidney exchanges for the NLDKSS since July 2008. Our contribution in this article is as follows:

- (1) We describe the IP constraints that are required to enforce the NLDKSS optimality definition. The description could help inform decision makers in other countries who are in the early stages of setting up a kidney exchange matching scheme.
- (2) We report on our practical experience over a 4.5-year period of using the technique to find optimal solutions for matching runs of the NLDKSS, which are carried out approximately every quarter. Our implementation involves a branch-and-cut

strategy to compute solutions to the IP models. Extending the software to incorporate a branch-and-price approach is discussed in Section 6.

- (3) We present empirical results arising from experiments conducted using a web application incorporating the IP-based algorithm that is capable of automating the experimental comparison of solutions according to a range of different optimality criteria. Again, these results arise from real datasets and indicate the extent to which the particular optimality criteria present in the UK influence the structure of the solutions that are ultimately computed. A key observation is that allowing four-way exchanges would be likely to lead to a moderate number of additional transplants.

1.2. Structure of This Article

The remainder of this article is organized as follows. Section 2 defines important notation and terminology and lists the hierarchical set of criteria that constitute the UK optimality definition. Then, in Section 3, we describe the IP model, indicating the variables and constraints involved and giving intuition for both. We also discuss two recent papers that have presented alternative IP models for kidney exchange [Glorie and van de Klundert 2012; Constantino et al. 2013]. Section 4 describes the actual matching run results, whereas Section 5 presents the experimental evaluation of alternative optimality criteria. Finally, Section 6 discusses some possible directions for future work.

2. THE NLDKSS OPTIMALITY CRITERIA

The problem of finding an optimal set of kidney exchanges essentially corresponds to computing optimal cycle packings in weighted directed graphs. Suppose we have a set S_P of incompatible donor–recipient pairs involving a set $S_D = \{d_1, \dots, d_{n_D}\}$ of donors of cardinality n_D and a set $S_R = \{r_1, \dots, r_{n_R}\}$ of recipients of cardinality n_R . That is, $S_P \subseteq S_D \times S_R$. Let $n_P = |S_P|$.

Now suppose we have a set $S_A = \{d'_1, d'_2, \dots, d'_{n_A}\}$ of altruistic donors of cardinality n_A . We associate with each altruistic donor $d'_i \in S_A$ a dummy recipient r'_i who is compatible with every donor $d_i \in S_D$ and incompatible with every donor $d'_i \in S_A$. Let

$$S = S_P \cup \{(d'_i, r'_i) : d'_i \in S_A\}.$$

Then, $|S| = n_P + n_A$. Let $n = |S|$.

Enumerate the pairs in S as $\{p_1, \dots, p_n\}$, where, without loss of generality, $S_P = \{p_1, \dots, p_{n_P}\}$. For each $p_k \in S_P$, where $p_k = (d_i, r_j)$, let $d(p_k)$ denote d_i and let $r(p_k)$ denote r_j . Similarly, for each $p_k \in S \setminus S_P$, where $p_k = (d'_i, r'_i)$, let $d(p_k)$ denote d'_i and let $r(p_k)$ denote r'_i .

For each $r_i \in S_R$, define

$$S_P(r_i) = \{p_k \in S_P : r(p_k) = r_i\}.$$

Note that possibly $|S_P(r_i)| > 1$; that is, r_i has multiple willing but incompatible donors. Similarly, for each $d_i \in S_D$, define

$$S_P(d_i) = \{p_k \in S_P : d(p_k) = d_i\}.$$

Also, possibly (but more unusually and forbidden in the NLDKSS), $|S_P(d_i)| > 1$; that is, d_i has multiple recipients for whom he or she is a willing but incompatible donor.

We model the kidney exchange problem by forming a weighted directed graph (digraph) $D = (V, A)$, where $V = \{v_1, \dots, v_n\}$: here, v_k corresponds to p_k ($1 \leq k \leq n$). Moreover $(v_i, v_j) \in A$ if and only if $d(v_i)$ is compatible with $r(v_j)$. In this way, two-cycles and three-cycles in D not involving an altruistic donor correspond to pairwise and three-way exchanges, respectively, whereas two-cycles and three-cycles in D involving

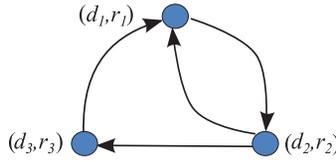


Fig. 1. Example of a three-cycle containing a back-arc and an embedded two-cycle.

an altruistic donor correspond to short and long chains, respectively, where, in practice the final donor in the chain donates a kidney to the DDWL. (Note that our model handles both short and long chains, despite the NLDKSS allowing only short chains at present.)

An arc (v_i, v_j) where $d(p_j) \in S_D$ has a real-valued weight $w(v_i, v_j) > 0$ that arises from a scoring system employed by NHSBT to measure the potential benefit of a transplant from $d(p_i)$ to $r(p_j)$. Factors involved in computing this weight include *waiting time* for $r(p_j)$ (based on the number of previous matching runs that $r(p_j)$ has been unsuccessfully involved in), $r(p_j)$'s *sensitization* (based on calculated HLA antibody reaction frequency), *HLA mismatch* levels between $d(p_i)$ and $r(p_j)$ (which, roughly speaking, corresponds to levels of tissue-type incompatibility), and points relating to the difference in ages between $d(p_i)$ and $d(p_j)$ (see Johnson et al. [2008] for more details).

An arc (v_i, v_j) where $d(p_j) \in S_A$ has weight $w(v_i, v_j) = 0$ (a discussion of this is given toward the end of Section 3.1). The *weight* of a cycle C in D is the sum of the weights of the individual arcs in C .

A *set of exchanges* in D is a permutation π of V such that (i) for each $v_i \in V$, if $\pi(v_i) \neq v_i$ then $(v_i, \pi(v_i)) \in A$, and (ii) no cycle in π has length > 3 . If $\pi(v_i) \neq v_i$ then v_i is said to be *matched*; otherwise, v_i is *unmatched*. A pair $p_i \in P$ is if v_i is matched, and p_i is *unmatched* otherwise.

Suppose some $v_k \in V$ is unmatched. Assume that $p_k \in S_P$. Let $d_i = d(p_k)$ and let $r_j = r(p_k)$. If all pairs in $S_P(d_i)$ (respectively, $S_P(r_j)$) are unmatched, then d_i (respectively, r_j) is said to be *unmatched* and will not participate in a kidney exchange. Now assume that $p_k \in S \setminus S_P$. Then $d(p_k)$, an altruistic donor, is said to be *unmatched*, but in practice he or she will donate directly to the DDWL. For this reason, we define the *size* of π (corresponding to the number of transplants yielded by this set of exchanges) to be the number of vertices matched by π plus the number of unmatched vertices corresponding to altruistic donors.

Given a three-cycle C in D with arcs $(v_i, v_j), (v_j, v_k), (v_k, v_i)$, we say that C contains a *back-arc* if, without loss of generality, $(v_j, v_i) \in A$. In such a case, we say that C contains an *embedded two-cycle* involving arcs $(v_i, v_j), (v_j, v_i)$. A three-cycle with a back-arc and an embedded two-cycle are illustrated in Figure 1. An *effective two-cycle* is either a two-cycle or a three-cycle with at least one back-arc.

A back-arc can be seen as a form of fault-tolerance in a three-cycle. To understand why, consider the three-cycle in Figure 1. If either d_3 or r_3 drops out (for example, due to illness), then the pairwise exchange involving (d_1, r_1) and (d_2, r_2) might still be able to proceed. On the other hand, if either of the pairs (d_1, r_1) or (d_2, r_2) were to withdraw, then this pairwise exchange would have failed anyway. Thus the risk involved with a three-way exchange, due to the greater likelihood (as compared to a pairwise exchange) of the cycle breaking down before transplants can occur, is mitigated with the inclusion of a back-arc.

We now present the definition of an *optimal* set of exchanges for the NLDKSS, as determined by the Kidney Advisory Group of NHSBT.

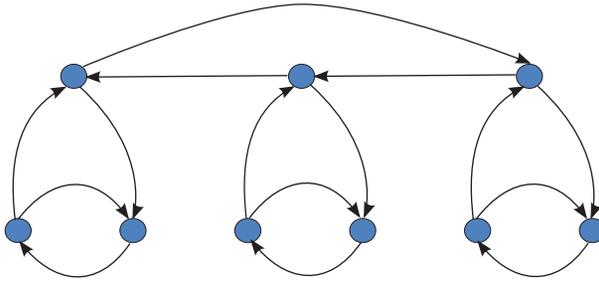


Fig. 2. There are two sets of exchanges, each of size 9 and containing three effective two-cycles: (i) choose the three three-cycles at the bottom, or (ii) choose three two-cycles at the bottom and the three-cycle at the top.

Definition 2.1. A set of exchanges π is *optimal* if:

- (1) The number of effective 2-cycles in π is maximized;
- (2) subject to (1), π has maximum size;
- (3) subject to (1)–(2), the number of three-cycles in π is minimized;
- (4) subject to (1)–(3), the number of back-arcs in the three-cycles in π is maximized;
- (5) subject to (1)–(4), the overall weight of the cycles in π is maximized.

We give some intuition for Definition 2.1 as follows. The first priority is to ensure that there are at least as many two-cycles and three-cycles with embedded two-cycles as there would be in an optimal solution containing only two-cycles. This is to ensure that the introduction of three-way exchanges is not detrimental to the maximum number of pairwise exchanges that could possibly take place. Subject to this, we maximize the total number of transplants (this is the number of unmatched altruistic donors, plus twice the number of pairwise exchanges and short chains, plus three times the number of three-way exchanges¹). Subject to this, we minimize the number of three-way exchanges. Despite Criterion 1, this is still required: for example, an optimal solution could either comprise (i) three three-way exchanges, each with a back-arc, or (ii) three pairwise exchanges and one three-way exchange (each of solutions (i) and (ii) has size 9 and contains three effective two-cycles)—see Figure 2 for an illustration. Solution (ii) is preferred by NHSBT because of the additional logistical difficulty of carrying out three-way exchanges compared to pairwise exchanges (this solution has two fewer three-way exchanges than solution (i)). Next, the number of back-arcs in three-way exchanges is maximized (note that a three-way exchange could contain more than one back-arc). Finally, we maximize the sum of the cycle weights.

3. FINDING AN OPTIMAL SOLUTION

3.1. Our Solution Methodology

In this section, we describe an algorithm that uses a sequence of IP formulations to find an optimal set of kidney exchanges with respect to Definition 2.1. After each run of the IP solver, we use the optimal value calculated at that iteration to enforce a constraint that must be satisfied in subsequent iterations. This ensures that once Criteria 1.. r in Definition 2.1 have been satisfied by an intermediate solution, they continue to hold when we additionally enforce Criterion $r + 1$ ($1 \leq r \leq 4$). At the outset, an IP formulation, called the *basic IP model*, is created. This extends the cycle formulation of Roth et al. [2007] in order to enable unmatched altruistic donors to be quantified (and to handle recipients with multiple donors and donors with multiple recipients).

¹The size should also include take into account long chains if they are permitted.

Let $\mathcal{C} = \{C_1, C_2, \dots, C_{n_C}\}$ denote the set of all possible cycles of lengths 2 and 3 in the digraph D . Without loss of generality, suppose that the two-cycles in \mathcal{C} are C_1, \dots, C_{n_2} , the three-cycles in \mathcal{C} are $C_{n_2+1}, \dots, C_{n_2+n_3}$, and the three-cycles with back-arcs in \mathcal{C} are $C_{n_2+1}, \dots, C_{n_2+n_3^b}$ (so $n_C = n_2 + n_3$).

Let x be an $(n_C + n_A) \times 1$ vector of binary variables $x_1, x_2, \dots, x_{n_C+n_A}$, where for $1 \leq j \leq n_C$, $x_j = 1$ if and only if C_j belongs to an optimal solution, and for $1 \leq j \leq n_A$, $x_{n_C+j} = 1$ if and only if altruistic donor d'_j is unmatched.

The basic IP model then contains the following constraints (in what follows, for a donor or recipient $a \in S_A \cup S_D \cup S_R$, we use the abbreviation $a \in C_j$ as shorthand for the case that $a = d(p_k)$ or $a = r(p_k)$ for some $p_k \in S$ where v_k belongs to C_j):

$$\sum_{C_j \in \mathcal{C}: r_i \in C_j} x_j \leq 1 \quad (1 \leq i \leq n_R) \quad (\text{A})$$

$$\sum_{C_j \in \mathcal{C}: d_i \in C_j} x_j \leq 1 \quad (1 \leq i \leq n_D, |S_P(d_i)| > 1) \quad (\text{B})$$

$$x_{n_C+i} + \sum_{C_j \in \mathcal{C}: d'_i \in C_j} x_j = 1 \quad (1 \leq i \leq n_A) \quad (\text{C})$$

Now, let c be a $1 \times (n_C + n_A)$ vector of values corresponding to the coefficients of current objective criterion; for example, for $1 \leq j \leq n_C$, c_j could be the length of C_j . The objective of the basic IP model is to solve $\max cx$ subject to (A)–(C).

We now provide some intuition for the model. Constraint (A) ensures that each recipient can be involved in at most one cycle in any solution, whereas Constraint (B) does likewise for the donors. Note that these constraints appear separately in order to model the possibility of recipients having multiple donors and donors having multiple recipients. In the case of Constraint (B), we need only range over the donors who have multiple recipients because the corresponding constraint for donors with a single recipient is already enforced by Constraint (A). Constraint (C) states that if no cycle containing a given altruistic donor is selected, then that donor is set to be unmatched.

We now describe the sequence of steps used to compute an optimal set of exchanges in D according to Definition 2.1. Item r in the following list corresponds to the step in the algorithm that enforces Criterion r (together with Criteria $1 \dots r - 1$) in the optimality definition. At each iteration, we indicate the additional constraints that are added to the basic IP model and also the objective function used at each iteration (where appropriate).

- (1) *The number of effective two-cycles is maximized.* Construct an undirected graph $G = (V, E)$ corresponding to the underlying digraph D , where the vertices in G and D are identical, and an edge in G corresponds to a two-cycle in D (i.e., $\{v_i, v_j\} \in E$ if and only if $(v_i, v_j) \in A$ and $(v_j, v_i) \in A$). Compute N_2 , the size of a maximum cardinality matching in G using the Micali-Vazirani implementation of Edmonds's algorithm [Micali and Vazirani 1980].
- (2) *Subject to (1), the size is maximized.* To enforce Condition (1) in Definition 2.1, we add the following constraint to the basic IP model:

$$\sum_{j=1}^{n_2+n_3^b} x_j \geq N_2. \quad (\text{D})$$

To enforce Condition (2), we add the objective $\max cx$, where $c_j = 2$ ($1 \leq j \leq n_2$), $c_j = 3$ ($n_2 + 1 \leq j \leq n_2 + n_3$), and $c_j = 1$ ($n_2 + n_3 + 1 \leq j \leq n_2 + n_3 + n_A$). That is, for $r \in \{2, 3\}$, each variable representing an r -cycle has coefficient r , and each variable representing an altruistic donor has coefficient 1, where the objective is to maximize. Let N denote the optimal value computed for the IP.

- (3) *Subject to (1)–(2), the number of three-cycles is minimized.* To enforce Conditions (1) and (2) in Definition 2.1, we add the following constraint to the basic IP model together with Constraint (D):

$$2 \sum_{j=1}^{n_2} x_j + 3 \sum_{j=n_2+1}^{n_2+n_3} x_j + \sum_{j=n_2+n_3+1}^{n_2+n_3+n_A} x_j \geq N. \quad (\text{E})$$

To enforce Condition (3), we add the objective $\min cx$, where $c_j = 0$ ($1 \leq j \leq n_2$), $c_j = 1$ ($n_2 + 1 \leq j \leq n_2 + n_3$), and $c_j = 0$ ($n_2 + n_3 + 1 \leq j \leq n_2 + n_3 + n_A$). That is, each variable representing a three-cycle has coefficient 1, whereas all others have coefficient 0. Let N_3 denote the optimal value computed for the IP.

- (4) *Subject to (1)–(3), the number of back-arcs in the three-cycles is maximized.* To enforce Conditions (1)–(3) in Definition 2.1, we add the following constraint to the basic IP model together with Constraints (D) and (E):

$$\sum_{j=n_2+1}^{n_2+n_3} x_j \leq N_3. \quad (\text{F})$$

To enforce Condition (4), let b_j be the number of back-arcs in cycle C_j ($n_2 + 1 \leq j \leq n_2 + n_3^b$) and define the objective $\max cx$, where $c_j = 0$ ($1 \leq j \leq n_2$), $c_j = b_j$ ($n_2 + 1 \leq j \leq n_2 + n_3^b$), and $c_j = 0$ ($n_2 + n_3^b + 1 \leq j \leq n_2 + n_3 + n_A$). That is, each variable corresponding to a two-cycle, or to a three-cycle with no back-arc, or to an altruistic donor, has coefficient 0, whereas each variable x_j representing a three-cycle with at least one back-arc has coefficient b_j . Let N_B denote the optimal value computed for the IP.

- (5) *Subject to (1)–(4), the overall weight is maximized.* To enforce Conditions (1)–(4) in Definition 2.1, we add the following constraint to the basic IP model together with Constraints (D)–(F):

$$\sum_{j=n_2+1}^{n_2+n_3^b} b_j x_j \geq N_B. \quad (\text{G})$$

To enforce Condition (5), let w_j denote the weight of cycle C_j ($1 \leq j \leq n_2 + n_3$), and define the objective $\max cx$, where $c_j = w_j$ ($1 \leq j \leq n_2 + n_3$) and $c_j = 0$ ($n_2 + n_3 + 1 \leq j \leq n_2 + n_3 + n_A$). That is, each variable corresponding to a cycle has a coefficient equal to the weight of that cycle, whereas each variable corresponding to an altruistic donor has coefficient 0. A solution to this final IP is an optimal set of exchanges relative to Definition 2.1.

Recall that the weight of an arc (v_i, v_j) , where v_j corresponds to an altruistic donor vertex, is 0. We give a brief discussion of this choice of arc weight. Clearly, it is difficult for the arc weight to depend on the particular choice of recipient from the DDWL that $d(p_i)$ would donate to in practice since the identity of this recipient is unknown when an optimal solution is constructed. If, say, a large positive constant arc weight of W were to be applied to (v_i, v_j) (representing the fact that a recipient from the DDWL who is an excellent match with $d(p_i)$ is very likely to be found) then, for fairness, an unmatched altruistic donor should also contribute a weight of W to the overall solution

(which could be achieved in practice via a self-loop with weight W on the corresponding vertex). But then every optimal solution contains a contribution of $n_A W$ to the weight (i.e., W from each altruistic donor), and hence the particular choice of W is irrelevant. Thus, to avoid the need to include self-loops on vertices corresponding to altruistic donors, we just assume that $W = 0$.

3.2. Alternative Techniques and Models

Considering the overall solution technique, we remark that an alternative to solving a series of IP formulations would be to solve a single IP relative to a weight function that captures the various criteria in the optimality definition (together with their priority levels) by assigning weights of successively decreasing orders of magnitude starting from Criterion 1 downward. This is, however, impractical: Due to the size of the datasets in practice, it would be computationally infeasible to work with such weights. For example, in the case of the October 2012 dataset, weights could be of the order of 10^{19} . Double-precision (64-bit) arithmetic can represent decimals using at most 16 significant figures, and hence numbers requiring a larger number of significant figures will be represented with some loss of precision, meaning that they cannot be compared accurately.

Another approach would be to abandon the lexicographic optimality definition and instead consider the five criteria defined in Definition 2.1 as separate objective functions in a multiobjective optimization problem [Deb 2014]. We return to this issue in Section 6. However, we remark here that the optimality definition given in Definition 2.1 as agreed by the Kidney Advisory Group of NHSBT is primarily driven by clinical considerations. We also point out that the Netherlands [Glorie and van de Klundert 2012] and South Korea [Kim et al. 2007] employ multiobjective lexicographic optimization, as in our case.

Two recent papers have presented different IP models for kidney exchange problems. Glorie and van de Klundert [2012] presented an IP model for the Dutch kidney exchange matching scheme. The criteria comprising the lexicographic objective function involve maximizing the number of transplants, maximizing the number of blood-type identical transplants, minimizing the length of the longest cycle or chain, and matching the recipient with the longest waiting time, among others. The problem is solved by formulating a series of IPs, each dealing with the separate parts of the optimality definition. These IPs essentially extend the cycle formulation of Roth et al. [2007]. The authors employ a branch-and-price approach to solving the IPs. They developed a novel pricing strategy based on the Bellman-Ford algorithm and showed that it performs well when searching for optimal solutions involving cycles of length up to 4 and DPD chains of length up to 6 in pools with up to 1,000 donor–recipient pairs.

Constantino et al. [2013] presented two new IP formulations for kidney exchange, under the assumption that the objective is to maximize the overall weight of the selected transplants (for some appropriate weight function). Each model is *compact*; that is, the number of variables and constraints is polynomial in the size of the problem instance. Note that neither the cycle formulation nor the edge formulation proposed by Roth et al. [2007] is compact: in the former case, the number of variables is exponential in the input size in general, whereas in the latter case the same is true for the number of constraints. The new models of Constantino et al. [2013] developed the edge formulation of Roth et al. [2007]; the authors referred to them as the *extended edge assignment formulation* (EA) and the *extended edge formulation* (EE). A computational analysis, based on branch-and-cut implementations of the models, demonstrated that EE performed better than EA, and, indeed, for dense instances (i.e., with a relatively large proportion of potential donations between donors and recipients), EE was capable of outperforming the cycle formulation.

4. RESULTS OBTAINED FROM THE NLDKSS

Prior to our involvement, NHSBT used an in-house algorithm that identified only pairwise exchanges. With the need to find both pairwise and three-way exchanges, a new software application was developed based on the algorithm outlined in Section 3. At its heart, the application uses the COIN-Cbc IP solver to solve each of the IP problems involved. COIN-Cbc was chosen due to its open licence agreement and the need to deploy the application commercially. Speed improvements obtained using IBM ILOG CPLEX and Gurobi Optimizer were minimal with the current size of the datasets.

The application can be extended via a plugin architecture that allows constraints to be created, added, or removed in a straightforward manner. This added flexibility allows our software to be easily adapted for use in other kidney exchange matching schemes, whether that involves simply changing the order of constraints or adding completely new ones.

The application can either be accessed programatically through a web API or, alternatively, manually via a web interface.² The former version (along with several earlier prototypes) has been used by NHSBT to find an optimal solution for each of the matching runs (occurring at roughly quarterly intervals) since July 2008. Table I summarizes the input to and output from each matching run between July 2008 and October 2012. In each case, an optimal solution³ was returned within a second (on a Linux Centos 5.5 machine with a Pentium 4 3GHz single-core processor with 2Gb RAM) despite a gradually increasing pool of donors.

The table shows key attributes for each matching run organized into three categories concerning: (i) the underlying digraph D , (ii) the solution identified by the algorithm, and (iii) the actual transplants arising from this solution. In Category (i), we indicate the number of vertices and arcs in D , the number of altruistic donors (which contribute toward the vertex count), and the numbers of two- and three-cycles in D . In Category (ii), we break the identified solution down into the numbers of two-cycles, three-cycles, effective two-cycles (denoted by N_2), and short chains, showing the overall size and weight. (Note that altruistic donors were not included in matching runs prior to January 2012.) Finally, in Category (iii), we show the numbers of short chains, pairwise, and three-way exchanges that actually occurred, together with the total number of actual transplants. Note that the total numbers of transplants shown in Categories (ii) and (iii) for the 2012 matching runs include the number of unmatched altruistic donors (as they donate directly to the DDWL).

The numbers of arcs in 2012 increased for a combination of reasons: (a) the introduction of altruistic donors (leading to the automatic inclusion of arcs of the form (v_i, v_j) where $d(v_i) \in S_D$ and $d(v_j) \in S_A$), (b) the relaxation of a previous requirement that donors of blood type O could only donate to recipients of the same blood type, and (c) the fact that all potential arcs are now included, whereas previously arcs not part of any cycle were suppressed. The number of arcs alone does not, however, reveal the complexity of the underlying optimization problem: This is more accurately represented by the numbers of two- and three-cycles (which account for the variables in the underlying IP). The total number of two- and three-cycles tends to fluctuate considerably and was largest in the July 2009 dataset.

Table II shows the minimum, maximum, average, and median values taken over each of the rows of Table I. In the case of the rows corresponding to the numbers of

²See <http://kidney.optimalmatching.com>.

³Note that the optimality criteria were slightly different from July 2008 to July 2009. See the Appendix for a more detailed discussion of this issue.

Table 1. Results Arising from Matching Runs between July 2008 and October 2012

Matching run	2008			2009				2010				2011				2012			
	Jul	Oct		Jan	Apr	Jul	Oct	Jan	Apr	Jun	Oct	Jan	Apr	Jun	Oct	Jan	Apr	Jun	Oct
Properties of D	#vertices	83	123	126	128	141	147	150	158	152	191	202	176	189	197	195	190	187	215
	#alt. donors	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	3	1	4
	#arcs	628	1,406	1,256	1,413	1,926	1,715	1,527	1,635	1,310	1,943	2,366	1,701	2,130	2,007	2,902	2,494	2,190	3,315
	#2-cycles r_2	2	14	17	20	55	4	17	23	4	20	19	9	34	18	115	21	22	35
	#3-cycles r_3	0	116	72	71	166	4	33	77	1	39	145	27	101	73	87	46	33	77
	$r_2 + r_3$	2	130	89	91	221	8	50	100	5	59	164	36	135	91	202	67	55	112
Identified solution	#2-cycles	1	6	5	5	4	0	3	2	3	3	3	0	5	7	1	0	2	6
	#3-cycles	0	3	1	2	7	2	1	6	0	2	10	4	4	5	6	5	2	5
	N_2	1	6	5	5	10	2	4	8	3	5	10	4	9	10	6	6	4	13
	#short chns	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	2	0	4
	size	2	21	13	16	29	6	9	22	6	12	36	12	22	29	24	20	11	35
	weight	6	930	422	618	1,168	300	135	782	261	473	1,328	464	794	1,094	2,882	1,872	1,175	3,599
Actual transplants	#pairwise	1	4	5	2	3	0	2	4	0	3	2	0	2	6	1	1	0	6
	#3-way	0	0	0	0	2	2	0	3	0	1	5	2	4	3	2	4	1	1
	#short chns	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	3
	total	2	8	10	4	12	6	4	17	0	9	19	6	16	21	10	18	4	22

Table II. Statistical Overview of Matching Runs between July 2008 and October 2012

		Minimum		Maximum		Average	Median
		Value	Dataset	Value	Dataset		
Properties of D	#vertices	83	Jul-08	215	Oct-12	164	167
	#alt. donors	1	Jun-12	4	Oct-12	3	3
	#arcs	628	Jul-08	3,315	Oct-12	1,881	1,821
	#2-cycles n_2	2	Jul-08	115	Jan-12	25	20
	#3-cycles n_3	0	Jul-08	166	Jul-09	65	72
	$n_2 + n_3$	2	Jul-08	221	Jul-09	90	90
Identified solution	#2-cycles	0	Apr-12	7	Oct-11	3	3
	#3-cycles	0	Jun-10	10	Jan-11	4	4
	N_2	1	Jul-08	13	Oct-12	6	6
	#short chns	0	Jun-12	4	Oct-12	2	2
	size	2	Jul-08	36	Jan-11	18	18
	weight	6	Jul-08	3,599	Oct-12	1,017	788
Actual transplants	#pairwise	0	Apr-11	6	Oct-11	2	2
	#3-way	0	Jun-10	5	Jan-11	2	2
	#short chns	0	Jan-12	3	Oct-12	1	1
	total	0	Jun-10	22	Oct-12	10	10

altruistic donors, identified short chains, and actual short chains, we consider 2012 data only (because altruistic donors were only introduced into the NLDKSS from January 2012) when computing these values. The maximum number of identified transplants (36) arose in the January 2011 matching run, with the maximum number of actual transplants (22) coming from the October 2012 run. In general, not all transplants identified by the software will lead to operations in practice: One reason is that more detailed cross-matching between each donor and recipient identified for transplant takes place after the matching run, which may lead to new incompatibilities being identified; also, a donor or recipient may become ill between the date of the matching run and the date of the operation.

In total, 325 potential transplants were identified, comprising 56 pairwise and 65 three-way exchanges, eight short chains, and two unmatched altruistic donors. These resulted in 188 *actual* transplants. This total comprised 42 pairwise and 30 three-way exchanges, four short chains, and six unused altruistic donors, implying that roughly 58% of identified potential operations proceeded to transplant.

5. DATA ANALYSIS SOFTWARE AND EMPIRICAL RESULTS

5.1. Introduction

Due to the complex nature of the optimality criteria used by the NLDKSS, it became obvious that there was a need to analyze the effect of each constraint. Furthermore, as the NLDKSS evolves, it is likely that the maximum length of a DPD chain and/or the maximum length of cycle allowed in a solution will increase. In turn, these developments might lead to additional constraints being required. The effect of such changes is often difficult to quantify because carrying out experimental comparisons can be time-consuming due to the significant development work required and the execution of simulations.

To this end, a web application⁴ (referred to as the *toolkit*) was developed that allows NHSBT staff to examine the impact of adding/removing constraints, allowing longer DPD chains, and increasing the maximum cycle length. The output from the

⁴See <http://toolkit.optimalmatching.com>.

application can give information such as the size and weight of an optimal set of exchanges, the number of each type of exchange (i.e., pairwise, three-way, etc.), and the number of DPD chains. This information can be downloaded in the form of a spreadsheet.

In this section, we report on an empirical analysis, undertaken using the toolkit, of the 18 matching runs that have taken place between July 2008 and October 2012. The aim was to determine the effect (in terms of the overall size or weight) of (i) prioritizing pairwise exchanges, (ii) minimizing the number of three-way exchanges and maximizing the number of back-arcs, (iii) allowing four-way exchanges in the optimality definition, and (iv) allowing long DPD chains.

Again, a Linux Centos 5.5 machine with a Pentium 4 3GHz single-core processor with 2Gb RAM was used, and every optimal solution was computed in under 2 seconds.

When undertaking these experiments, we adopted a strategy whereby the dataset for the first matching run (July 2008) was used to find an optimal set of exchanges with respect to some optimality criteria, and then the dataset for each subsequent matching run was constructed by taking the original dataset for that run (that is the dataset used for the matching run when it originally took place) and removing all those donor–recipient pairs that were matched in previous runs. This strategy makes the assumption that those donor–recipient pairs belonging to an optimal set of exchanges never come back into the matching scheme, which in reality is not the case (due to some identified transplants not proceeding). However, a more complex model introduces problems. One such problem is selecting those people who should re-enter the matching scheme when their transplant identified at a previous run fails to proceed. Choosing people who should re-enter has a significant effect on the characteristics of a solution output in a subsequent round and leads to inconsistent results that do not reflect what happens in the real matching scheme. Furthermore, those donor–recipient pairs who were previously matched in a given dataset under the current optimality criteria, but who are unmatched using a new set of optimality criteria, are problematic because we do not possess the data to allow us to calculate the compatibilities between these pairs and any new pairs entering the scheme in subsequent runs. Indeed, for any such newly arriving donor–recipient pairs, we only add in their compatibilities that were known from previous runs for patients who are still involved in the matching scheme. Therefore, to provide a fair comparison when applying a different optimality criterion using this strategy, we also use this same strategy with the current optimality criteria to get base figures with which to compare the new results.

5.2. Prioritizing (Effective) Two-cycles

First, we examined the effect on the size of an optimal set of exchanges π in three cases concerning whether to prioritize two-cycles or effective two-cycles:

- (A) when Definition 2.1 is unchanged;
- (B) when Criterion 1 is omitted from Definition 2.1;
- (C) when Criterion 1 is replaced by “maximize the number of two-cycles.”

Figure 3 shows the size of an optimal solution, and Table III shows the total number of transplants over each of the 18 matching runs. These reveal that when we relax the need to first maximize the number of two-cycles or effective two-cycles (case B from the list) we obtained only a single extra transplant over all 18 matching runs. In contrast, if we require the number of pairwise exchanges alone to be maximized as first priority, then we would see 27 fewer transplants in total. In many cases, obtaining one additional transplant could make it worth changing the criteria; however, in this

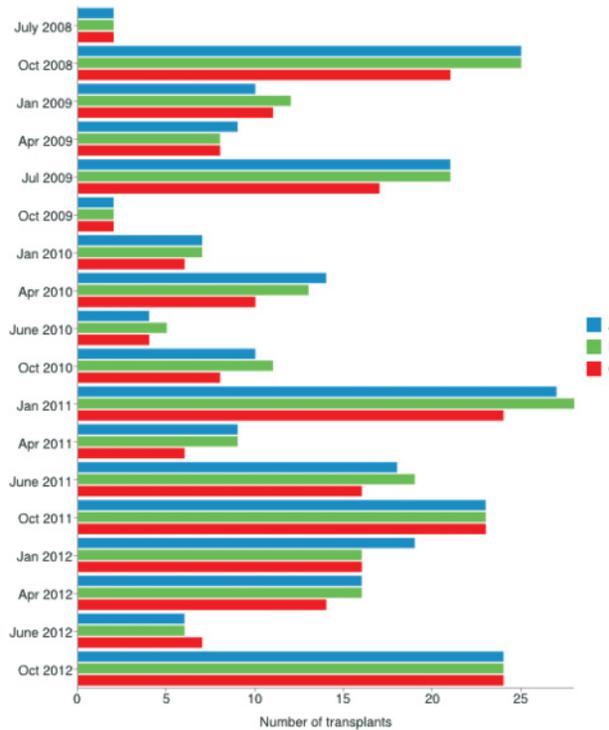


Fig. 3. Effect of prioritizing pairwise exchanges.

case, given the desirable properties of embedded two-cycles, the extra risk involved for a single extra transplant is unlikely to be justified.

5.3. Effect of Criteria 3 and 4

We next analyzed the effect on an optimal solution when we first apply Criteria 1 and 2 from Definition 2.1, then decide whether or not to apply Criteria 3 and 4 (i.e., minimize the number of three-cycles and maximize the number of back-arcs respectively), and subsequently maximize the total weight. This gives four cases that correspond to the combinations of including/excluding Criteria 3 and 4.

It turns out that, in these four cases, the solutions output are different in only a single matching run (January 2012). In the matching run where the solutions are different, only the constraint to maximize the number of back-arcs had any effect; that is, posting the constraint to minimize the number of three-way exchanges had no effect on any solution. It appears that enforcing Criterion 1 (to maximize the number of effective two-cycles) results in a very small set of candidates for a solution that is optimal overall. If we no longer insist that Criterion 1 is enforced, then variations on the weight of an optimal solution are observed in the four cases over many of the matching runs. Prior to January 2012, it would have been natural to ask whether Criteria 3 and 4 should be removed. However, given that Criterion 4 was shown to be discriminating, it seems reasonable to assume that Criterion 3 may also have an impact in the future as the size of the datasets grow. Furthermore, considering that the additional time required to find a solution that satisfies Criteria 1–5 (as opposed to satisfying only Criteria 1, 2, and 5) is minimal (a solution is found in both cases in under 2 seconds for each dataset), it seems sensible to retain both constraints.

Table III. Total Number of Transplants when Prioritizing Pairwise Exchanges

Case	Total Number of Transplants
A	246
B	247
C	219

Table IV. Total Number of Transplants when Increasing the Maximum Cycle Length

Case	Total Number of Transplants
A	192
B	246
C	265

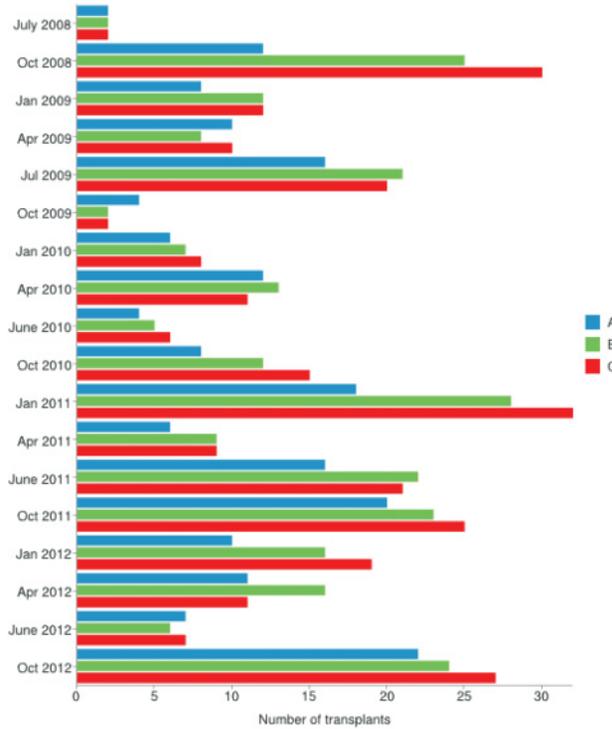


Fig. 4. Effect of increasing the maximum cycle length.

5.4. Increasing the Maximum Cycle Length

We next determined the effect of increasing the maximum cycle length. Initially, the NLDKSS allowed only pairwise exchanges in an optimal solution, but three-way exchanges were permitted from April 2008 (subject to the condition that the number of effective two-cycles is first maximized). Clearly, extending the solution to allow for four-way exchanges ought to increase further the number of transplants, but this must be set against the greater risk of such exchanges not proceeding.

In Figure 4, we show the total number of transplants at each of the 18 matching runs (with the sums of these numbers of transplants shown in Table IV) if an optimal set of exchanges π is defined as follows:

- (A) maximize the size of π , allowing only two-cycles;
- (B) first maximize the number of effective two-cycles, then, subject to that, maximize the total number of transplants, allowing only two-cycles and three-cycles;
- (C) first maximize the number of effective two-cycles, then, subject to this, maximize the size of π , allowing two-, three-, and four-cycles.

Table V. Total Number of Transplants when Altruistic Donors Are Included in the Dataset

Case	Total Number of Transplants
No altruistic donors	74
Allowing only short chains	82
Allowing short and long chains	82

As expected, allowing four-way exchanges leads to an increased number of transplants. Over all of the 18 matching runs, we would have obtained an 19 extra transplants by allowing pairwise, three-way, and four-way exchanges compared to allowing only pairwise and three-way exchanges. This number is smaller than the increase observed when allowing both pairwise and three-way exchanges (compared to allowing only pairwise exchanges), where we would have obtained an extra 54 transplants over all the matching runs to October 2012.

The additional 19 transplants in the former case is likely to be a lower bound for the true increase because the donor–recipient pairs for whom we had limited compatibility information in the experiments may in reality be involved in many additional cycles (this issue does not affect the 246 transplants computed for case (B), where pairwise and three-way exchanges are permitted because this case corresponds to the original optimality criteria; thus, we had full compatibility information for all donor-recipient pairs).

In simulations, Roth et al. [2007] observed that the benefit (in terms of the number of transplants) obtained from allowing four-way exchanges (in addition to pairwise and three-way exchanges) was small. Our results would seem to contradict this; however, one explanation is that the NLDKSS datasets are sparser than those generated by the Saidman dataset generator [Saidman et al. 2006], which was used by Roth et al. [2007] in their simulations. Ashlagi and Roth [2012] observed that, in general, the Saidman generator [Saidman et al. 2006] underestimates the percentage of highly sensitized recipients in the population and thus produces datasets that are denser than those occurring in reality. Ashlagi et al. [2011] argued that, in sparser pools, longer exchanges and cycles can lead to a tangible benefit in terms of the number of transplants, which is consistent with our findings. This issue was also discussed by Glorie and van de Klundert [2012].

5.5. Allowing Long Chains

Finally, we observed the effects of including altruistic donors in the dataset. Altruistic donors were introduced into the NLDKSS in January 2012. Using the data from the four matching runs involving such donors to date, we undertook experiments to allow us to understand both the impact (in terms of increased numbers of transplants) of including altruistic donors and also the benefits of including only short chains or both short and long chains (subject to the optimality criteria in Definition 2.1).

The tests (whose results are presented in Table V) showed that if altruistic donors were not included in the NLDKSS, we would have obtained in total eight fewer transplants over all the matching runs between January 2012 and October 2012—we assume that each unused altruistic donor counts as a single transplant in this case (because they would donate to the DDWL). In contrast, if we allowed both short and long chains, there is in fact no difference in the number of transplants identified. The fact that introducing long chains had no effect is surprising and is likely to be due to the fact that several of the altruistic donors included so far have been incompatible with a significant number of patients. We also note that, in two of the matching runs (April 2012 and June 2012), it was in fact beneficial, in terms of the number of transplants

that could be achieved, to not use one of the included altruistic donors in a short or long chain in each of these runs.

6. FUTURE WORK

Our case study has been driven by a particular practical application, and, as such, the empirical evaluation in Section 5 was based on real datasets (spanning a period of 4.5 years). However, further experiments are required on artificially generated data, which will facilitate both a larger number of trials and bigger datasets. (The generator should of course construct datasets that aim to reflect various properties, such as blood-type prevalence, inherent in the real datasets.) This will provide important information on how far the software, in its current form, is likely to scale. Furthermore, using these datasets may provide greater insight into the effect a particular constraint has on the system.

In what follows, we describe a number of challenges that future work on algorithms for paired and altruistic kidney donation in the UK context may need to address:

- Scalability*: It is important to ensure that the algorithms described in this article can scale up to cope with both larger and denser datasets. Larger datasets will arise as participation in the NLDKSS increases, and denser datasets will follow as a result of increased numbers of altruistic donors and the potential introduction of long chains and four-way exchanges. Although the branch-and-cut strategy described in this article copes comfortably with the current size of NLDKSS datasets, it may not scale to much larger datasets involving, say, 500 donor–recipient pairs. Work is currently under way to extend the algorithm described in this article using branch-and-price techniques along the lines of those described by Abraham et al. [2007], so that we can continue to meet the needs of the NLDKSS in the future.
- NEAD chains* [Rees et al. 2009; Ashlagi et al. 2011, 2012; Dickerson et al. 2012b]: These are similar to DPD chains in that they are triggered by an altruistic donor. They consist of c smaller subchains s_1, \dots, s_c , for some $c > 0$. Each subchain s_r ($1 \leq r \leq c$) is identified during a single matching run, and the operations associated with s_r need not be performed simultaneously. Essentially, s_r is similar to a DPD chain except that, if $r < c$, rather than donating directly to the DDWL, the final donor d in s_r is called a *bridge donor* who triggers subchain s_{r+1} , typically after a period of several months. Thus, d will play the role of an altruistic donor in the matching run in which s_{r+1} is constructed, considering s_{r+1} as a DPD chain. Of course, in the period between d 's recipient receiving a kidney and d donating a kidney themselves, d might become ill or renege. This risk (and indeed the similar risk of other donors in a subchain dropping out) is taken into account when estimating the likely benefit of NEAD chains over DPD chains [Ashlagi et al. 2011]. The final donor in subchain s_c donates to the DDWL. NEAD chains have been used successfully in the United States [Rees et al. 2009], and there are strong arguments for increasing their use in the future [Ashlagi et al. 2012]. Were they to be introduced in the United Kingdom, the algorithm described in this article would need to be adapted if the maximum length of an individual subchain were to exceed that of a long chain.
- Compatible pairs* [Gentry et al. 2007]: Recently, it has been proposed that compatible donor–recipient pairs should participate in kidney exchange matching schemes, such as the NLDKSS, in addition to incompatible donor–recipient pairs. This would enable a recipient r with a willing and compatible donor d to potentially obtain a kidney that is an even better match for him or her than the kidney that r would receive directly from d . Moreover, the inclusion of such pairs would widen the available pool of donors and could enable a recipient with a willing but incompatible donor to participate in

a kidney exchange or a DPD chain that would not otherwise have existed. Of course, if a recipient r with a willing and compatible donor d is not involved in a kidney exchange or DPD chain, then there is always the fall-back position that d can donate to r directly. In practice, as noted by Constantino et al. [2013], this option can be represented in D with the aid of a self-loop involving the vertex corresponding to (d, r) . The introduction of compatible pairs could potentially give rise to a major increase in the size of datasets in future, providing further motivation for the work on scalability as discussed earlier.

- Dynamic kidney exchange* [Awasthi and Sandholm 2009; Ünver 2010; Anderson et al. 2012; Dickerson et al. 2012a]: The optimization problem studied in this article is static in nature: that is, matching runs occur at approximately quarterly intervals, and although new donor–recipient pairs may arrive in between matching runs, they must wait for the next run to determine whether they have been identified for inclusion in a kidney exchange or DPD chain. The length of delay between matching runs is chosen to achieve a balance between, on the one hand, allowing a sufficiently large pool of donors and recipients to build up to ensure that a reasonable number of good-quality kidney exchanges and DPD chains can be identified, and, on the other hand, ensuring that donors and recipients do not have too long to wait before the next run occurs. Another option is to process a new donor–recipient pair (d, r) as soon as they arrive: The mechanism should determine whether to try to match this pair within a chain or exchange immediately or else keep the pair in the pool in order to obtain a greater subsequent benefit. More work is needed to understand the potential benefits of dynamic kidney exchange versus static kidney exchange in the UK context. This could involve measuring attributes such as average waiting time and number and quality (as measured by a weight function) of identified transplants using both methods, on both real and simulated data, over a period of several years.
- Modeling failure probabilities* [Chen et al. 2012; Dickerson et al. 2013; Li et al. 2014; Pedroso 2014]: Recently, a number of papers have focused on modeling the probability of “failure” of vertices and/or arcs in the underlying digraph D for the kidney exchange problem. For example, the failure of a vertex could represent a donor or recipient becoming too ill for nephrectomy or transplant, whereas the failure of an arc could represent the existence of a positive crossmatch that had not been previously determined. Either event would cause a cycle or chain containing the affected vertex or arc to fail. With the aid of such probabilities, a possible goal, then, could be to maximise the expected number of transplants over all possible sets of exchanges. When following this approach [Chen et al. 2012; Dickerson et al. 2013; Li et al. 2014; Pedroso 2014], the expected utility of a cycle can take into account back-arcs and thus for the potential for an embedded two-cycle contained in a three-cycle to proceed if the three-cycle fails. This represents an alternative approach to explicitly prioritizing cycles with back-arcs, as we have done here.
- Multiojective optimization* [Deb 2014]: As discussed in Section 3.2, the current hierarchical optimality definition as given by Definition 2.1 (and the consequent lexicographic optimization strategy) is motivated primarily by clinical requirements. However, clinical priorities evolve over time, and it may be that, in the future, the individual objective functions from Definition 2.1 should no longer be optimized in a lexicographic fashion. This brings us into the realm of multiojective optimization [Deb 2014], where we could seek a Pareto optimal solution by, for example, using a single objective function of the form $g(x) = \sum_{i=1}^5 z_i f_i(x)/opt_i(x)$, where x is a vector of binary variables arising from a problem instance, $f_i(x)$ is the individual objective value corresponding to the i th objective in Definition 2.1, whereas $opt_i(x)$ is the

Table VI. Results Arising from Matching Runs from July 2008 to July 2009

Matching Run		2008		2009		
		Jul	Oct	Jan	Apr	Jul
Identified solution (Table I)	#2-cycles	1	6	5	5	4
	#3-cycles	0	3	1	2	7
	size	2	21	13	16	29
	weight	6	1,033	526	738	1,257
Optimal solution (Definition 2.1)	#2-cycles	1	2	3	2	2
	#3-cycles	0	7	5	5	9
	size	2	25	15	19	31
	weight	6	930	422	618	1,168

optimal value of this objective function (when optimized according to this single objective only) and z_i is some constant. An advantage of such a strategy is that desirable but potentially conflicting objectives (such as maximizing the number of back-arcs and minimizing the number of three-way exchanges) can be traded off against one another. A key problem, of course, is to find suitable z_i values, and future work could address this by means of experimental trials on both real and generated data, measuring gains and losses against objective values obtained from lexicographic optimization and from optimizing objectives individually.

We close by remarking that recent work [Ashlagi et al. 2010; Ashlagi and Roth 2011; Caragiannis et al. 2011; Toulis and Parkes 2011; Ashlagi and Roth 2012] has investigated incentives for a hospital h participating in a kidney exchange matching scheme to withhold its easiest-to-match pairs and deal with them internally while reporting only its hardest-to-match pairs (e.g., involving highly sensitized recipients) to the centralized matching scheme. Patients at other hospitals could be affected because they may lose out on a transplant if h does not truthfully report all of its pairs to the scheme. Although such considerations are, of course, important in general, it is worth mentioning that, at least at present in the UK, there is no legal framework allowing a hospital to undertake operations associated with kidney exchanges or DPD chains internally and outside of the NLDKSS due to the tight regulation of these types of organ transplants by the Human Tissue Authority.

APPENDIX

Between July 2008 and July 2009, the criteria used for an optimal solution in the NLDKSS were different from those in Definition 2.1. Essentially, the previous definition dropped *effective* from Criterion 1 and omitted Criteria 3 and 4 the current definition. The results in the “Identified Solution” rows in Table I are those for the optimality criteria that were in force at the time of the given matching run. However, for comparison purposes, we present in Table VI an indication of the structure, size, and weight of an optimal solution had the present definition been used in each of the matching runs between July 2008 and July 2009. Not surprisingly, the table indicates that a solution with a bigger size and weight could have been identified because three-cycles with embedded two-cycles are now allowed to contribute to the count of the cycles identified in Criterion 1.

REFERENCES

- D. Abraham, A. Blum, and T. Sandholm. 2007. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC'07)*. ACM, 295–304.

- R. Anderson, I. Ashlagi, and D. Gamarnik. 2012. A dynamic model of kidney exchange programs. Extended abstract of poster presented at the 2012 Stochastic Networks Conference. Retrieved from <http://stoch-nets-2012.lids.mit.edu/posters/Anderson.pdf>.
- Alliance for Paired Donation website. <http://www.paireddonation.org> (accessed August 20, 2014).
- I. Ashlagi, F. Fischer, I. Kash, and A. D. Procaccia. 2010. Mix and match. In *Proceedings of the 11th ACM Conference on Electronic Commerce*. ACM, 305–314.
- I. Ashlagi, D. Gamarnik, M. Rees, and A. Roth. 2012. *The Need for (Long) Chains in Kidney Exchange*. Working Paper 18202, National Bureau of Economic Research, Cambridge, MA. Retrieved from <http://www.nber.org/papers/w18202>.
- I. Ashlagi, D. Gilchrist, A. Roth, and M. Rees. 2011. Nonsimultaneous chains and dominos in kidney paired donation – revisited. *American Journal of Transplantation* 11, 5, 984–994.
- I. Ashlagi and A. Roth. 2011. Individual rationality and participation in large scale, multi-hospital kidney exchange. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. ACM, 321–322.
- I. Ashlagi and A. Roth. 2012. New challenges in multihospital kidney exchange. *American Economic Review* 102, 3, 354–359.
- P. Awasthi and T. Sandholm. 2009. Online stochastic optimization in the large: Application to kidney exchange. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, 405–411.
- C. Barnhart, E. Johnson, G. Nemhauser, M. Savelsbergh, and P. Vance. 1998. Branch-and-price: Column generation for solving huge integer programs. *Operations Research* 46, 316–329.
- P. Biró, D. Manlove, and R. Rizzi. 2009. Maximum weight cycle packing in directed graphs, with application to kidney exchange. *Discrete Mathematics, Algorithms and Applications* 1, 4, 499–517.
- I. Caragiannis, A. Filos-Ratsikas, and A. Procaccia. 2011. An improved 2-agent kidney exchange mechanism. In *Proceedings of the 7th International Workshop on Internet and Network Economics*. Lecture Notes in Computer Science Series, vol. 7090. Springer, 37–48.
- Y. Chen, Y. Li, J. Kalbfleisch, Y. Zhou, A. Leichtman, and P. Song. 2012. Graph-based optimization algorithm and software on kidney exchanges. *IEEE Transactions on Biomedical Engineering* 59, 7, 1985–1991.
- M. Constantino, X. Klimentova, A. Viana, and A. Rais. 2013. New insights on integer-programming models for the kidney exchange problem. *European Journal of Operational Research* 231, 1, 57–68.
- K. Deb. 2014. Multi-objective optimization. In *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques* (2nd ed.), E. Burke and G. Kendall, Eds. Springer, 403–449.
- J. Dickerson, A. Procaccia, and T. Sandholm. 2012a. Dynamic matching via weighted myopia with application to kidney exchange. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*. AAAI Press, 1340–1346.
- J. Dickerson, A. Procaccia, and T. Sandholm. 2012b. Optimizing kidney exchange with transplant chains: Theory and reality. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*. Vol. 2. International Foundation for Autonomous Agents and Multiagent Systems, 711–718.
- J. Dickerson, A. Procaccia, and T. Sandholm. 2013. Failure-aware kidney exchange. In *Proceedings of the 14th ACM Conference on Electronic Commerce*. ACM, 323–340.
- S. Gentry, D. Segev, M. Simmerling, and R. Montgomery. 2007. Expanding kidney paired donation through participation by compatible pairs. *American Journal of Transplantation* 7, 2361–2370.
- K. Glorie and J. van de Klundert. 2012. Iterative branch-and-price for large multi-criteria kidney exchange. *Econometric Institute Report* 2012-11, Erasmus University Rotterdam.
- R. Johnson, J. Allen, S. Fuggle, J. Bradley, and C. Rudge. 2008. Early experience of paired living kidney donation in the United Kingdom. *Transplantation* 86, 1672–1677.
- K. Keizer, M. de Klerk, B. Haase-Kromwijk, and W. Weimar. 2005. The Dutch algorithm for allocation in living donor kidney exchange. *Transplantation Proceedings* 37, 589–591.
- M. de Klerk, K. Keizer, F. Claas, M. Witvliet, B. Haase-Kromwijk, and W. Weimar. 2005. The Dutch national living donor kidney exchange program. *American Journal of Transplantation* 5, 2302–2305.
- B. S. Kim, Y. S. Kim, S. I. Kim, M. S. Kim, H. Y. Lee, Y. L. Kim, C. D. Kim, C. W. Yang, B. S. Choi, D. J. Han, Y. S. Kim, S. J. Kim, H. Y. Kim, and D. J. Kim. 2007. Outcome of multipair donor kidney exchange by a web-based algorithm. *Journal of the American Society of Nephrology* 18, 3, 1001–1006.
- J. Kwak, O. Kwon, K. Lee, C. Kang, H. Park, and J. Kim. 1999. Exchange-donor program in kidney transplantation: A single-center experience. *Transplantation Proceedings* 31, 344–345.
- Y. Li, P. Song, Y. Zhou, A. Leichtman, M. Rees, and J. Kalbfleisch. 2014. Optimal decisions for organ exchanges in a kidney paired donation program. *Statistics in Biosciences* 6, 85–104.

- M. Lucan. 2007. Five years of single-center experience with paired kidney exchange transplantation. *Transplantation Proceedings* 39, 1371–1375.
- M. Lucan, P. Rotariu, D. Neculoiu, and G. Iacob. 2003. Kidney exchange program: A viable alternative in countries with low rate of cadaver harvesting. *Transplantation Proceedings* 35, 933–934.
- S. Micali and V. Vazirani. 1980. An $O(\sqrt{|V|} \cdot |E|)$ algorithm for finding maximum matching in general graphs. In *Proceedings of the 21st Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 17–27.
- R. Montgomery, S. Gentry, W. Marks, D. Warren, J. Hiller, J. Houp, A. Zachary, J. Melancon, W. Maley, H. Rabb, C. Simpkins, and D. Segev. 2006. Domino paired kidney donation: A strategy to make best use of live non-directed donation. *The Lancet* 368, 9533, 419–421.
- National Kidney Registry. 2013. Alliance for Paired Donation. Retrieved from <http://www.kidneyregistry.org>.
- NHS Blood and Transplant. 2012. Transplant Activity Report 2011-12. Retrieved from http://www.organdonation.nhs.uk/statistics/transplant_activity_report/archive_activity_reports/pdf/ukt/activity_report_2011_12.pdf.
- K. Park, J. Lee, S. Kim, and Y. Kim. 2004. Exchange living-donor kidney transplantation: Diminution of donor organ shortage. *Transplantation Proceedings* 36, 2949–2951.
- K. Park, J. Moon, S. Kim, and Y. Kim. 1999. Exchange-donor program in kidney transplantation. *Transplantation Proceedings* 31, 356–357.
- J. P. Pedroso. 2014. Maximizing expectation on vertex-disjoint cycle packing. In *Proceedings of the 14th International Conference on Computational Science and Its Applications*. Lecture Notes in Computer Science Series, vol. 8580. Springer, 32–46.
- M. Rees, J. Kopke, R. Pelletier, D. Segev, M. Rutter, A. Fabrega, J. Rogers, O. Pankewycz, J. Hiller, A. Roth, T. Sandholm, M. Ünver, and R. Montgomery. 2009. A nonsimultaneous, extended, altruistic-donor chain. *New England Journal of Medicine* 360, 11, 1096–1101.
- A. Roth, T. Sönmez, and M. Ünver. 2004. Kidney exchange. *Quarterly Journal of Economics* 119, 2, 457–488.
- A. Roth, T. Sönmez, and M. Ünver. 2005. Pairwise kidney exchange. *Journal of Economic Theory* 125, 151–188.
- A. Roth, T. Sönmez, and M. Ünver. 2007. Efficient kidney exchange: Coincidence of wants in a market with compatibility-based preferences. *American Economic Review* 97, 3, 828–851.
- A. Roth, T. Sönmez, M. Ünver, F. Delmonico, and S. Saidman. 2006. Utilizing list exchange and nondirected donation through ‘chain’ paired kidney donations. *American Journal of Transplantation* 6, 11, 2694–2705.
- S. Saidman, A. Roth, T. Sönmez, M. Ünver, and S. Delmonico. 2006. Increasing the opportunity of live kidney donation by matching for two and three way exchanges. *Transplantation* 81, 5, 773–782.
- S. Segev, S. Gentry, D. Warren, B. Reeb, and R. Montgomery. 2005. Kidney paired donation and optimizing the use of live donor organs. *Journal of the American Medical Association* 293, 1883–1890.
- P. Toulis and D. Parkes. 2011. A random graph model of kidney exchanges: efficiency, individual-rationality and incentives. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. ACM, 323–332.
- United Network for Organ Sharing. 2013. Homepage. Retrieved from <http://www.unos.org>.
- M. Ünver. 2010. Dynamic kidney exchange. *Review of Economic Studies* 77, 372–414.

Received November 2012; revised October 2013, August 2014; accepted September 2014