

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

David F. Manlove and Gregg O'Malley



School of Computing Science
Sir Alwyn Williams Building
University of Glasgow
Glasgow G12 8QQ
UK

Technical Report
TR-2012-330
March 2012
Revised June 2012

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

David F. Manlove and Gregg O'Malley

School of Computing Science, University of Glasgow, Glasgow, UK.

Email: {david.manlove,gregg.omalley}@glasgow.ac.uk.

Abstract

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 [4, 24] 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 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 4-way exchanges would be likely to lead to a significant number of additional transplants.

1 Introduction

It is understood that transplantation is the most effective treatment that is currently known for kidney failure. In the UK alone, as of 31 March 2012 there were 6612 patients waiting on the transplant list for a donor kidney, with the median waiting time being 1153 days for an adult and 307 days for a child. Kidneys used for transplantation can come from both deceased and living donors. In the UK, around 36% of all kidney transplants between 1 April 2011 and 31 March 2012 came from living donors [19].

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 patients with no genetic or emotional connection.

*Research financially supported by the National Health Service Blood and Transplant (NHSBT) between 1 April 2010 – 30 June 2011, and between 1 January 2012 - 30 June 2013 (project 10-11-ODT). We gratefully acknowledge collaboration with Rachel Johnson, Joanne Allen and with other staff at NHSBT. The views expressed here are those of the authors and not necessarily those of NHSBT. Research also financially supported by a Knowledge Transfer Account from the Engineering and Physical Sciences Research Council between 1 July 2011 - 31 December 2011. Finally, we would like to thank Péter Biró, Rob Irving, Kirstin MacDonald and Ana Viana for valuable input into this work.

With the introduction of the HTA, a patient with an incompatible donor can now “swap” their donor with that of another patient in a similar position, via “kidney exchanges” that involve two or more incompatible patient–donor pairs. For example, a *pairwise (kidney) exchange* involves two incompatible patient–donor pairs (p_1, d_1) and (p_2, d_2) , where d_1 is compatible with p_2 , and d_2 is compatible with p_1 : d_1 donates a kidney to p_2 in exchange for d_2 donating a kidney to p_1 . *3-way exchanges* extend this concept to three pairs in a cyclic manner.

In a number of countries, centralised programmes (also known as *kidney exchange matching schemes*) have been introduced to help optimise the search for kidney exchanges. These include the USA [1, 2, 3], the Netherlands [13, 14], South Korea [15, 21, 20] and Romania [17, 16].

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) [12]. The purpose of the NLDKSS is two-fold: firstly to identify those pairs that are compatible with one another and then subsequently to optimise 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 increases. This is because all operations have to be performed simultaneously due to the risk of a donor renegeing on his/her commitment to donate a kidney after their loved one has received a kidney. Mainly for this reason, at the present time the NLDKSS does not allow exchanges involving more than three pairs.

A kidney exchange matching scheme may also include *altruistic donors*, who do not have an associated patient and who are willing to donate a kidney to a stranger. An altruistic donor d_0 can either donate directly to a patient (without a donor) on the Deceased Donor Waiting List (DDWL), or else trigger a *domino paired chain* (DPC) [6] involving one or more incompatible patient–donor pairs: here d_0 donates to a patient p_1 in exchange for p_1 ’s donor donating to the patient p_2 in the next pair in the chain, with the final donor donating to the DDWL. A DPC is *short* (resp. *long*) if it consists of one (resp. two) incompatible patient–donor pairs). At present the NLDKSS allows short but not long chains.

Kidney exchange has received considerable attention in the computer science, economics and medical literature in recent years [4, 5, 6, 7, 8, 9, 10, 11, 22, 23, 24, 25, 26, 27, 28]. 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., [8] for more details). However when pairwise and 3-way exchanges are allowed, the problem of finding a set of exchanges that maximises the number of transplants is NP-hard [4] and indeed APX-hard [8].

Abraham *et al.* [4], and independently Roth *et al.* [24], described two integer programming (IP)-based formulations of the problem of finding a maximum weight set of kidney exchanges, when both pairwise and 3-way exchanges are permitted (here, the weights can measure the benefit of potential transplants). Abraham *et al.* [4] 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.

In this paper we present an application-driven case study, showing how the cycle formulation can be extended in order to handle kidney exchange in the UK. In particular, we show how to model a complex set of criteria (given in Section 2) that form the definition of an *optimal* set of kidney exchanges. Although most of the criteria have not been explicitly mentioned elsewhere in the literature, they are natural rather than idiosyncratic. We have implemented the technique and it has been used by NHSBT to find optimal sets of kidney exchanges for the NLDKSS since July 2008. Our contribution in this paper is as follows:

1. We describe the IP constraints that are required in order to enforce the NLDKSS optimality definition (Section 3). The description could help to 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 3-year period of using the technique to find optimal solutions for matching runs of the NLDKSS, which are carried out approximately every quarter (Section 4).
3. We present empirical results arising from a web application that is capable of automating the experimental comparison of solutions according to a range of different optimality criteria (Section 5). Again, these results arise from real datasets and indicate 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 4-way exchanges would be likely to lead to a significant number of additional transplants.

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 n incompatible patient–donor pairs $\{(p_i, d_i) : 1 \leq i \leq n\}$ and k altruistic donors $\{d_{n+i} : 1 \leq i \leq k\}$. We associate with each altruistic donor d_{n+i} a *dummy patient* p_{n+i} who is compatible with every donor d_j where $1 \leq j \leq n$.

We model the kidney exchange problem by forming a weighted directed graph $D = (V, A)$, where $V = \{v_1, \dots, v_{n+k}\}$ and v_i corresponds to (p_i, d_i) ($1 \leq i \leq n+k$). Moreover $(v_i, v_j) \in A$ if and only if d_i is compatible with p_j . In this way, 2-cycles and 3-cycles in D not involving an altruistic donor correspond to pairwise and 3-way exchanges respectively, whilst 2-cycles and 3-cycles in D involving an altruistic donor d_{n+i} 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.)

An arc (v_i, v_j) 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_i to p_j . Factors involved in computing this weight include *waiting time* for p_j (based on the number of previous matching runs that p_j has been unsuccessfully involved in), p_j 's *sensitisation* (based on calculated HLA antibody reaction frequency), *HLA mismatch* levels between d_i and p_j (which roughly speaking corresponds to levels of tissue-type incompatibility) and points relating to the difference in ages between d_i




Figure 1: Example of a 3-cycle containing a back-arc and an embedded 2-cycle.

and d_j (see [12] for more details). 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*. Suppose some $v_i \in V$ is unmatched. If $1 \leq i \leq n$, then neither d_i nor p_i will participate in a kidney exchange. However if $i > n$, d_i 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 3-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 2-cycle* involving arcs $(v_i, v_j), (v_j, v_i)$. A 3-cycle with a back-arc and an embedded 2-cycle is illustrated in Figure 1. An *effective 2-cycle* is either a 2-cycle or a 3-cycle with a back-arc.

A back-arc can be seen as a form of fault-tolerance in a 3-cycle. To understand why, consider the 3-cycle in Figure 1. If either p_3 or d_3 drops out (for example due to illness), then the pairwise exchange involving (p_1, d_1) and (p_2, d_2) might still be able to proceed. On the other hand, if either of the pairs (p_1, d_1) or (p_2, d_2) were to withdraw, then this pairwise exchange would have failed anyway. Thus the risk involved with a 3-way exchange, due to the greater likelihood (as compared to a pairwise exchange) of the cycle breaking down before transplants can be scheduled, 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.

Definition 1. A set of exchanges π is optimal if:

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

We give some intuition for Definition 1 as follows. The first priority is to ensure that there are at least as many 2-cycles and embedded 2-cycles as there would be in an optimal solution containing only 2-cycles. This is to ensure that the introduction of 3-way exchanges is not detrimental to the maximum number of pairwise exchanges that could possibly take place. Subject to this we maximise the total number of transplants (this is the number of unmatched altruistic donors, plus

twice the number of pairwise exchanges and short chains, plus 3 times the number of 3-way exchanges and long chains). Subject to this we minimise the number of 3-way exchanges. Despite Criterion 1, this is still required: for example an optimal solution could either comprise three 3-way exchanges, each with a back-arc, or three pairwise exchanges and one 3-way exchange (both solutions have size 9 and contain three effective 2-cycles) – see Appendix A for an illustration. Clearly there is less risk of cycles breaking down with the second solution. Next the number of back-arcs in 3-way exchanges is maximised (note that a 3-way exchange could contain more than one back-arc). Finally we maximise the sum of the cycle weights.

3 Finding an optimal solution

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 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 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 [4, 24] in order to enable unmatched altruistic donors to be quantified. Recall that n is the number of incompatible patient–donor pairs and k is the number of altruistic donors. The basic IP model is then constructed as follows:

1. list all the possible cycles of length 2 and 3 in the directed graph D as C_1, C_2, \dots, C_m , where, without loss of generality, the 2-cycles are C_1, \dots, C_{n_2} , the 3-cycles are $C_{n_2+1}, \dots, C_{n_2+n_3}$, and the 3-cycles with back-arcs are $C_{n_2+1}, \dots, C_{n_2+n_3}$ (so $m = n_2 + n_3$);
2. let x be an $(m + k) \times 1$ vector of binary variables x_1, x_2, \dots, x_{m+k} , where for $1 \leq i \leq m$, $x_i = 1$ if and only if C_i belongs to an optimal solution, and for $1 \leq i \leq k$, $x_{m+i} = 1$ if and only if altruistic donor d_{n+i} is unmatched;
3. let A be an $(n + 2k) \times (m + k)$ $\{-1, 0, 1\}$ -valued matrix, whose entries are all 0 apart from the following:
 - (a) for $1 \leq i \leq n$ and $1 \leq j \leq m$, $A_{i,j} = 1$ if and only if C_j contains d_i ;
 - (b) for each i ($1 \leq i \leq k$), in rows $n + 2i - 1$ and $n + 2i$:
 - i. for $1 \leq j \leq m$, $A_{n+2i-1,j} = 1$ if and only if cycle C_j contains d_{n+i} , and for $1 \leq j \leq k$, $A_{n+2i-1,m+j} = 1$ if and only if $i = j$;
 - ii. for $1 \leq j \leq m$, $A_{n+2i,j} = -1$ if and only if cycle C_j contains d_{n+i} , and for $1 \leq j \leq k$, $A_{n+2i,m+j} = -1$ if and only if $i = j$;
4. let b be an $(n + 2k) \times 1$ vector where:
 - (a) for each i ($1 \leq i \leq n$), $b_i = 1$;
 - (b) for each i ($1 \leq i \leq k$) $b_{n+2i-1} = 1$ and $b_{n+2i} = -1$;
5. let c be a $1 \times (m + k)$ vector of values corresponding to the coefficients of current objective criterion, e.g., c_j could be the length of C_j ;
6. solve $\max cx$ such that $Ax \leq b$.

We now provide some intuition for the model above. Part 3(a) (in combination with 4(a)) ensures that each patient–donor pair is involved in at most one cycle in any solution. Similarly 3(b)(i) (with 4(b)) ensures that each altruistic donor is involved in at most one cycle. 3(b)(i) (with 4(b)) also ensures that if a cycle involving an altruistic donor d_{n+i} is chosen then v_{n+i} must be matched. Similarly, 3(b)(ii) (with 4(b)) ensures that if no cycle involving an altruistic donor d_{n+i} is chosen then v_{n+i} must be unmatched.

We now describe the sequence of steps that is used in order to compute an optimal set of exchanges in D according to Definition 1. Item r in the following list corresponds to the step in the algorithm that enforces Criterion r (together with Criteria 1.. $r - 1$) 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 2-cycles is maximised.*

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 2-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 Edmonds' algorithm [18]. Then add the following constraint:

$$x_1 + x_2 + \dots + x_{n_2+n_3} \geq N_2. \quad (1)$$

2. *Subject to (1), the size is maximised.*

Consider the basic IP model, together with (1), and with the objective $\max cx$, where $c_i = 2$ ($1 \leq i \leq n_2$), $c_i = 3$ ($n_2 + 1 \leq i \leq n_2 + n_3$) and $c_i = 1$ ($n_2 + n_3 + 1 \leq i \leq n_2 + n_3 + k$). 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 maximise. After calculating the optimal value N , add the following constraint:

$$2x_1 + \dots + 2x_{n_2} + 3x_{n_2+1} + \dots + 3x_{n_2+n_3} + x_{n_2+n_3+1} + \dots + x_{n_2+n_3+k} \geq N. \quad (2)$$

3. *Subject to (1)-(2), the number of 3-cycles is minimised.*

Consider the basic IP model, together with (1)-(2), and with the objective $\min cx$, where $c_i = 0$ ($1 \leq i \leq n_2$), $c_i = 1$ ($n_2 + 1 \leq i \leq n_2 + n_3$) and $c_i = 0$ ($n_2 + n_3 + 1 \leq i \leq n_2 + n_3 + k$). That is, each variable representing a 3-cycle has coefficient 1, whilst all others have coefficient 0. After calculating the optimal value N_3 , add the following constraint:

$$x_{n_2+1} + \dots + x_{n_2+n_3} \leq N_3. \quad (3)$$

4. *Subject to (1)-(3), the number of back-arcs in the 3-cycles is maximised.*

Let k_i be the number of back-arcs in cycle C_i ($n_2 + 1 \leq i \leq n_2 + n_3$). Consider the basic IP, together with (1)-(3), and with the objective $\max cx$, where $c_i = 0$ ($1 \leq i \leq n_2$), $c_i = k_i$ ($n_2 + 1 \leq i \leq n_2 + n_3$) and $c_i = 0$ ($n_2 + n_3 + 1 \leq i \leq n_2 + n_3 + k$). That is, each variable corresponding to a 2-cycle or to an altruistic donor has coefficient 0, and each variable x_i representing a 3-cycle has coefficient k_i . Suppose that an optimal solution has value N_B . Add the following constraint:

$$k_{n_2+1}x_{n_2+1} + \dots + k_{n_2+n_3}x_{n_2+n_3} \geq N_B. \quad (4)$$

5. *Subject to (1)-(4), the overall weight is maximised.*

For each i ($1 \leq i \leq n_2 + n_3$), let w_i be the weight of cycle C_i . Consider the basic IP model, together with (1)-(4), and with the objective $\max cx$, where $c_i = w_i$ ($1 \leq i \leq n_2 + n_3$) and $c_i = 0$ ($n_2 + n_3 + 1 \leq i \leq n_2 + n_3 + k$). That is, each variable corresponding to a cycle has coefficient equal to the weight of that cycle, whilst 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 1.

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 downwards. This is however impractical: due to the size of the datasets in practice, it would be computationally infeasible to work with such weights.

Another approach would be to assign smaller weights that somehow prioritise cycles with “good” characteristics, such as 3-cycles with back-arcs. However it is not clear how such weights should be defined, especially as theoretically there is no upper bound on the score of an arc as provided by NHSBT. Any attempt along these lines could never result in a concrete definition of exactly what is being optimised in an optimal solution, as we have obtained here.

4 NLDKSS in Practice

Prior to our involvement, NHSBT used an in-house algorithm that identified only pairwise exchanges. With the need to find both pairwise and 3-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 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 interface¹. The former version (along with several prototypes) has been used by NHSBT to find an optimal solution in each of the matching runs (occurring at roughly quarterly intervals), since July 2008. Table 1 summarises the input to, and output from, each matching run between July 2008 and October 2011. In each case an optimal solution² 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. In total 235 potential transplants were identified (47 pairwise and 47 3-way exchanges), which have re-

¹<http://kidney.optimalmatching.com>

²Note that the optimality criteria were slightly different from July 2008 to July 2009. See Appendix B for a more detailed discussion of this issue.

Matching run		2008		2009				2010				2011			
		Jul	Oct	Jan	Apr	Jul	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
	#arcs	628	1406	1256	1413	1926	1715	1527	1635	1310	1943	2366	1701	2130	2007
	#2-cycles	2	14	17	20	55	4	17	23	4	20	19	9	34	18
	#3-cycles	0	116	72	71	166	4	33	77	1	39	145	27	101	73
Identified solution	#2-cycles	1	6	5	5	4	0	3	2	3	3	3	0	5	7
	#3-cycles	0	3	1	2	7	2	1	6	0	2	10	4	4	5
	size	2	21	13	16	29	6	9	22	6	12	36	12	22	29
Actual transplants	#pairwise	1	4	5	2	3	0	2	4	0	3	2	0	2	6
	#3-way	0	0	0	0	2	2	0	3	0	1	5	2	4	3
	total	2	8	10	4	12	6	4	17	0	9	19	6	16	21

Table 1: Results arising from matching runs between July 2008 and October 2011.

sulted in 134 *actual* transplants³ (34 pairwise and 22 3-way exchanges). Together with the 4 pairwise exchanges that were identified as part of the NLDKSS prior to our involvement, there have been a total of 142 actual transplants to date. Note that altruistic donors were not introduced into the NLDKSS until January 2012, and hence in Table 1, the number of vertices corresponds to the number of patient–donor pairs in each matching run.

The table shows that the matching run in January 2011 had the largest number of vertices and arcs in the underlying digraph, and the largest number of potential transplants of any matching run were identified (36). Even so, the digraph underlying the July 2009 dataset had a larger number of 2-cycles and 3-cycles. It is expected that the digraphs will become much denser once altruistic donors are introduced, and larger as awareness of the scheme grows over time.

5 Data analysis software and empirical results

Due to the complex nature of the optimality criteria used by the NLDKSS, it became obvious that there was a need to analyse the effect of each constraint. Furthermore, as the NLDKSS evolves it is likely that the maximum length of a DPC 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, as 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 altruistic chains, and increasing the maximum cycle length. The output from the application can determine information such as the size and weight of an optimal set of exchanges, the number of each type of exchange (i.e. pairwise, 3-way, etc.), and the number of DPCs. This information can be downloaded in the form of a spreadsheet.

³In 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 patient identified for transplant takes place after the matching run, which may lead to new incompatibilities being identified; also a donor or patient may become ill between the date of the matching run and the date of the operation.

⁴<http://toolkit.optimalmatching.com>

In this section we report on an empirical analysis, using the toolkit, of the 14 matching runs that have taken place between July 2008 and October 2011. The aim is to determine the effect (in terms of the overall size or weight) of (i) prioritising pairwise exchanges, (ii) minimising the number of 3-way exchanges and maximising the number of back-arcs, and (iii) allowing 4-way exchanges in the optimality definition. 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 two seconds.

First we examine the effect on the size of an optimal set of exchanges π in three cases concerning whether to prioritise 2-cycles or effective 2-cycles:

- (A) when Definition 1 is unchanged;
- (B) when Criterion 1 is omitted from Definition 1;
- (C) when Criterion 1 is replaced by “maximise the number of 2-cycles”.

Figure 2 shows the size of an optimal solution in each case, over the 14 matching runs. It reveals that on average if we relax the need to first maximise the number of 2-cycles or effective 2-cycles (case B from the above list) we would obtain only a single extra transplant per matching run. In contrast, if we require the number of pairwise exchanges alone to be maximised as first priority, then we would see a reduction in the number of transplants by, on average, 3 per matching run. In many cases obtaining a single extra transplant could make it worth changing the criteria, however in this case, given the desirable properties of embedded 2-cycles, the extra risk involved for the single extra transplant is unlikely to be justified.

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

It turns out that in each of these four cases, the solution output in each of the 14 matching runs is exactly the same, i.e., posting constraints to minimise the number of 3-ways exchanges or maximise the number of back-arcs has no effect. It appears that enforcing Criterion 1 (maximise the number of effective 2-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. The additional time required to find a solution that satisfies Criteria 3 and 4 (as opposed to satisfying only Criteria 1, 2 and 5) is minimal (a solution is found in both cases in under two seconds for each dataset). Hence Criteria 3 and 4 should be retained as they may well have an impact for larger / denser datasets that are likely to feature in matching runs in the short / medium term.

We next determine the effect of increasing the maximum cycle size. Initially the NLDKSS allowed only pairwise exchanges in an optimal solution, but 3-way exchanges were permitted from April 2008 (subject to the condition that the number of effective 2-cycles is first maximised). Clearly extending the solution to allow for 4-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 3 we show the total number of transplants at each of the 14 matching runs if an optimal set of exchanges π is defined as follows:




Figure 2: Effect of prioritising pairwise exchanges




Figure 3: Effect of increasing the maximum cycle size

- (A) maximise the size of π , allowing only 2-cycles;
- (B) first maximise the number of effective 2-cycles, then subject to that maximise the total number of transplants, allowing only 2-cycles and 3-cycles;
- (C) first maximise the number of effective 2-cycles, then subject to this maximise the number of *effective 3-cycles* (defined to be the number of 3-cycles plus the number of 4-cycles with embedded 3-cycles), then subject to this maximise the size of π , allowing 2-cycles, 3-cycles and 4-cycles.

As expected, allowing 4-way exchanges leads to an increased number of transplants: on average, an additional 4 transplants per matching run (compared to allowing only pairwise and 3-way exchanges). This number is smaller than the increase observed when allowing both pairwise and 3-way exchanges (compared to allowing only pairwise exchanges) where on average there are 7 additional transplants per matching run.

Finally we observe the effects of including altruistic donors in the dataset. Altruistic donors are set to be included in the NLDKSS from January 2012. In order to understand their impact in terms of increased numbers of transplants, the data from the January 2009 matching run was augmented by NHSBT staff with six altruistic donors known at that time. Of particular interest was to determine the benefits of including only short chains or both short and long chains (subject to the optimality criteria in Definition 1).

The test results indicated that, in the absence of altruistic donors, 15 transplants were obtained. When only short chains are permitted, 27 transplants were identified. Finally, if we allow both short and long chains, 31 transplants were identified. This shows that the difference between including only short chains, as opposed to both short and long chains, is of lesser importance than the benefit obtained by allowing short chains, as compared to not including altruistic donors. However, given that

any long chain must have at least one embedded 2-cycle, the risk of including long chains should be seen as minimal.

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 42 months). However further experiments are required on artificially generated data which will facilitate both a larger number of trials and bigger 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.

Future work must also ensure that the algorithms described in this paper can scale as participation in the NLDKSS increases. It is anticipated that column generation techniques, along the lines of those described by Abraham *et al.* [4], will be required to ensure that we can meet the needs of the NLDKSS in the future, given the likelihood of larger datasets and the potential introduction of long chains and 4-way exchanges.

We close by remarking that recent work [5, 7, 9, 28] 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, whilst reporting only its hardest-to-match pairs to the centralised matching scheme. Patients at other hospitals could be affected as 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 DPCs internally and outside of the NLDKSS, due to the tight regulation of these types of organ transplants by the Human Tissue Authority.

References

- [1] <http://www.nepke.org> (New England Program for Kidney Exchange website).
- [2] <http://www.paireddonation.org> (Alliance for Paired Donation website).
- [3] <http://www.paireddonationnetwork.org> (Paired Donation Network website).
- [4] D.J. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets: enabling nationwide kidney exchanges. In *Proceedings of EC '07: the 8th ACM Conference on Electronic Commerce*, pages 295–304. ACM, 2007.
- [5] I. Ashlagi, F.A. Fischer, I.A. Kash, and A. D. Procaccia. Mix and match. In *Proceedings of EC '10: the 11th ACM Conference on Electronic Commerce*, pages 305–314. ACM, 2010.
- [6] I. Ashlagi, D.S. Gilchrist, A.E. Roth, and M.A. Rees. Nonsimultaneous chains and dominos in kidney paired donation – revisited. *American Journal of Transplantation*, 11(5):984–994, 2011.

- [7] I. Ashlagi and A. Roth. Individual rationality and participation in large scale, multi-hospital kidney exchange. In *Proceedings of EC '11: the 12th ACM Conference on Electronic Commerce*, pages 321–322. ACM, 2011.
- [8] P. Biró, D.F. Manlove, and R. Rizzi. Maximum weight cycle packing in directed graphs, with application to kidney exchange. *Discrete Mathematics, Algorithms and Applications*, 1(4):499–517, 2009.
- [9] I. Caragiannis, A. Filos-Ratsikas, and A.D. Procaccia. An improved 2-agent kidney exchange mechanism. In *Proceedings of WINE '11: the 7th International Workshop on Internet and Network Economics*, volume 7090 of *Lecture Notes in Computer Science*, pages 37–48. Springer, 2011.
- [10] Y. Chen, J. Kalbfleisch, Y. Li, P. Song, and Y. Zhou. Computerized platform for optimal organ allocations in kidney exchanges. In *Proceedings of . BIOCAMP 11: the 2011 International Conference on Bioinformatics and Computational Biology*, 2011.
- [11] J. Dickerson, A.D. Procaccia, and T. Sandholm. Optimizing kidney exchange with transplant chains: Theory and reality. In *Proceedings of AAMAS 2012: the 11th International Conference on Autonomous Agents and Multiagent Systems*, to appear. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [12] R.J. Johnson, J.E. Allen, S.V. Fuggle, J.A. Bradley, and C.J. Rudge. Early experience of paired living kidney donation in the United Kingdom. *Transplantation*, 86:1672–1677, 2008.
- [13] K.M. Keizer, M. de Klerk, B.J.J.M. Haase-Kromwijk, and W. Weimar. The Dutch algorithm for allocation in living donor kidney exchange. *Transplantation Proceedings*, 37:589–591, 2005.
- [14] M. de Klerk, K.M. Keizer, F.H.J. Claas, M. Witvliet, B.J.J.M. Haase-Kromwijk, and W. Weimar. The Dutch national living donor kidney exchange program. *American Journal of Transplantation*, 5:2302–2305, 2005.
- [15] J.Y. Kwak, O.J. Kwon, K.S. Lee, C.M. Kang, H.Y. Park, and J.H. Kim. Exchange-donor program in kidney transplantation: A single-center experience. *Transplantation Proceedings*, 31:344–345, 1999.
- [16] M. Lucan. Five years of single-center experience with paired kidney exchange transplantation. *Transplantation Proceedings*, 39:1371–1375, 2007.
- [17] M. Lucan, P. Rotariu, D. Neculoiu, and G. Iacob. Kidney exchange program: a viable alternative in countries with low rate of cadaver harvesting. *Transplantation Proceedings*, 35:933–934, 2003.
- [18] S. Micali and V.V. Vazirani. An $O(\sqrt{|V|} \cdot |E|)$ algorithm for finding maximum matching in general graphs. In *Proceedings of FOCS '80: the 21st Annual IEEE Symposium on Foundations of Computer Science*, pages 17–27. IEEE Computer Society, 1980.

- [19] NHS Blood and Transplant. Weekly Statistics Report (accessed 11 June 2012). http://www.organdonation.nhs.uk/ukt/statistics/downloads/weekly_stats.pdf.
- [20] K. Park, J.H. Lee, S.I. Kim, and Y.S. Kim. Exchange living-donor kidney transplantation: Diminution of donor organ shortage. *Transplantation Proceedings*, 36:2949–2951, 2004.
- [21] K. Park, J.II Moon, S.II Kim, and Y.S. Kim. Exchange-donor program in kidney transplantation. *Transplantation Proceedings*, 31:356–357, 1999.
- [22] A.E. Roth, T. Sönmez, and M.U. Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119(2):457–488, 2004.
- [23] A.E. Roth, T. Sönmez, and M.U. Ünver. Pairwise kidney exchange. *Journal of Economic Theory*, 125:151–188, 2005.
- [24] A.E. Roth, T. Sönmez, and M.U. Ünver. Efficient kidney exchange: Coincidence of wants in a market with compatibility-based preferences. *American Economic Review*, 97(3):828–851, 2007.
- [25] S.L. Saidman, A.E. Roth, T. Sönmez, M.U. Ünver, and S.L. Delmonico. Increasing the opportunity of live kidney donation by matching for two and three way exchanges. *Transplantation*, 81(5):773–782, 2006.
- [26] S.L. Saidman, A.E. Roth, T. Sönmez, M.U. Ünver, and S.L. Delmonico. Utilizing list exchange and undirected good samaritan donation through chain paired kidney exchanges. *American Journal of Transplantation*, 6(11):2694–2705, 2006.
- [27] S.L. Segev, S.E. Gentry, D.S. Warren, B. Reeb, and R.A. Montgomery. Kidney paired donation and optimizing the use of live donor organs. *Journal of the American Medical Association*, 293:1883–1890, 2005.
- [28] P. Toulis and D.C. Parkes. A random graph model of kidney exchanges: efficiency, individual-rationality and incentives. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 323–332. ACM, 2011.

Appendix A




Figure 4: There are two sets of exchanges, each of size 9 and containing 3 effective 2-cycles: (i) choose the three 3-cycles at the bottom, or (ii) choose three 2-cycles at the bottom and the 3-cycle at the top.

Matching run		2008		2009		
		Jul	Oct	Jan	Apr	Jul
Identified solution (Table 1)	#2-cycles	1	6	5	5	4
	#3-cycles	0	3	1	2	7
	size	2	21	13	16	29
Optimal solution (Definition 1)	#2-cycles	1	2	3	2	2
	#3-cycles	0	7	5	5	9
	size	2	25	15	19	31

Table 2: Results arising from matching runs from July 2008 to July 2009.

Appendix B

Between July 2008 and July 2009, the criteria used for an optimal solution in the NLDKSS were different from those in Definition 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 1 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 2 an indication of the structure and size of an optimal solution, had the present definition been used in each of matching run between July 2008 and July 2009. Not surprisingly, the table indicates that more transplants could have been identified, since 3-cycles with embedded 2-cycles are now allowed to contribute to the count of the cycles identified in Criterion 1.