



Improved instance generation for kidney exchange programmes

Maxence Delorme^a, Sergio García^c, Jacek Gondzio^c, Jörg Kalcsics^c, David Manlove^b,
William Pettersson^{b,*}, James Trimble^b

^a Department of Econometrics and Operations Research, Tilburg University, The Netherlands

^b School of Computing Science, University of Glasgow, United Kingdom

^c School of Mathematics, University of Edinburgh, United Kingdom

ARTICLE INFO

Keywords:

kidney exchange
Saidman generator
Matheuristic

ABSTRACT

Kidney exchange programmes increase the rate of living donor kidney transplants, and operations research techniques are vital to such programmes. These techniques, as well as changes to policy regarding kidney exchange programmes, are often tested using random instances created by a Saidman generator. We show that instances produced by such a generator differ from real-world instances across a number of important parameters, including the average number of recipients that are compatible with a certain donor. We exploit these differences to devise powerful upper and lower bounds and we demonstrate their effectiveness by optimally solving a benchmark set of Saidman instances in seconds; this set could not be solved in under thirty minutes with previous algorithms. We then present new techniques for generating random kidney exchange instances that are far more consistent with real-world instances from the UK kidney exchange programme. This new process for generating random instances provides a more accurate base for comparisons of algorithms and models, and gives policy-makers a better understanding of potential changes to policy leading to an improved decision-making process.

1. Introduction

Kidney disease affects millions of people worldwide, and the two known treatment options for end-stage kidney disease are dialysis and kidney transplantation. Of these options, kidney transplantation offers a better quality of life, better life expectancy, and is cheaper (Axelrod et al., 2018), but donor kidneys are in short supply. These kidneys are donated by either deceased or living donors, with better outcomes for the recipient correlated with living donors (Hart et al., 2017; Wolfe et al., 2010). However, it is difficult for a recipient to find a compatible and willing living donor on their own. A Kidney Exchange Programme (KEP) alleviates the need for recipients to be compatible with their own donors (Roth et al., 2004). Instead, recipients join KEPs with a willing paired donor (or multiple willing paired donors), and at certain times the KEP performs a matching run, determining an optimal set of exchanges to perform such that a donor donates a kidney if their paired recipient receives a kidney (and at most one donor paired with any recipient donates a kidney). These exchanges involving recipients with paired donors form cycles: the length of these cycles is often limited for logistical reasons, but limiting this maximum length can reduce the efficiency of a KEP. Additionally, individuals can join a KEP to donate a kidney without expecting any reciprocal donation: such individuals are referred to as non-directed (or altruistic) donors, and donors paired

with a recipient are called directed donors. Non-directed donors can initiate a chain of transplants, by donating to a recipient whose donor can then donate to a further recipient and so-on. We will use the term *exchange* to refer to either a cycle or a chain.

Identifying the potential transplants is a vital part of the operation of a KEP, and numerous operations research techniques and algorithms have been devised for finding such exchanges. The problem of finding a largest set of transplants is polynomial-time solvable in very restricted settings (such as limiting cycle and chain lengths to 2), but in general is NP-hard (Abraham et al., 2007), and so integer programming (IP) is often used to find solutions. The first IP models created a graph of potential kidney exchanges (Abraham et al., 2007; Roth et al., 2007), and used one variable for either each potential cycle or each edge, but there have been numerous substantial improvements in this area, including compact formulations as well as column generation, branch-and-price, branch-and-price-and-cut, and failure-aware modelling (Alvelos et al., 2019; Anderson et al., 2015; Constantino et al., 2013; Delorme et al., 2020; Dickerson et al., 2016, 2012; Lam and Mak-Hau, 2020; Manlove and O'Malley, 2015; McElfresh et al., 2019). A summary of models and techniques currently in use by real-world KEPs is also available (Biró et al., 2019b).

* Corresponding author.

E-mail address: william.pettersson@glasgow.ac.uk (W. Pettersson).

<https://doi.org/10.1016/j.cor.2022.105707>

Received 10 May 2021; Received in revised form 8 November 2021; Accepted 9 January 2022

Available online 21 January 2022

0305-0548/© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

In this paper we use the UK Living Kidney Sharing Scheme (UK-LKSS) as our real-world KEP example (Manlove and O'Malley, 2015). The UK-LKSS is run by NHS Blood and Transplant, a special health authority of the National Health Service (NHS) within the UK. The UK-LKSS has been operating since 2007, and is currently the largest KEP (by both pool size and transplant count) in Europe (Biró et al., 2019a).

We use the term *Kidney Exchange problem* (KE) to refer to the problem of finding a set of cycles and chains that yield the maximum number of transplants, given a dataset containing information about donors, recipients and compatibilities between them. The effectiveness of the techniques used to find optimal sets of exchanges is often compared using randomly-generated instances that are created to mimic real-world instances. Such random instances are also useful for policy-makers, as they allow changes to policies (such as prioritisation of highly-sensitised patients, or the introduction of a wider range of exchange structures) to be tested. In the literature, a generator due to Saidman et al. (2006), henceforth referred to as the *Saidman generator*, is commonly used to generate such instances. The Saidman generator draws recipients and donors with various properties (such as blood groups) based on the random distribution of such parameters in populations. A more thorough introduction to the Saidman generator is given in Section 4.2. We also highlight that PrefLib (Mattei and Walsch, 2020), an online repository of matching problems with preferences, includes a set of instances that have been produced using the Saidman generator (Dickerson et al., 2012). We will refer to these as the PrefLib instances. The best known approach in the literature for solving these particular instances has been demonstrated by Lam and Mak-Hau (2020), who use a branch-and-price-and-cut algorithm to solve all 30 non-trivial instances when cycles are limited to length 3, and 19 of 30 non-trivial instances when cycles are limited to length 4. However, these PrefLib instances differ from real-world instances in a range of parameters (such as edge density of the graph of potential kidney exchanges), meaning that algorithms that can solve the PrefLib instances quickly may not be as applicable to real-world problems and vice-versa.

1.1. Related work

Many different approaches to improving the performance of KEPs have been investigated in the literature. We outline here two particular developments in the field, and highlight that for both of these, and indeed for almost any improvement to KEPs, the simulation of any proposed changes depends on the use of suitable random or benchmark instances for a baseline comparison.

The creation of transnational KEPs, either as new programmes or through coordination between or amalgamation of existing national programmes, allows KEPs to share their pools, thus increasing their overall effectiveness through more efficient allocation of resources. However, collaboration between countries when human organs are involved is a sensitive issue, both ethically and legally. Individual rationality has been proposed as one requirement for such collaboration (Ashlagi and Roth, 2011, 2012), and further research has suggested equality measures such as the Shapley value (Biró et al., 2020) as well as IP models that allow for, and have simulated, different constraints within different countries (Mincu et al., 2020). We note that KEPs within some countries (such as the USA) may exhibit similar characteristics without being transnational, as transplant centres (or groups of transplant centres) either collaborate or compete for kidney transplants (Li et al., 2018).

Another important development in the field of kidney exchange is the introduction of robust, or fault-tolerant, KEPs. A potential transplant may be selected by an algorithm or policy for transplantation, but may fail due to any one of a number of reasons (e.g., donor or recipient illness, unexpected incompatibilities). Robust KEP algorithms aim to assign a likelihood of progression (alternatively, a likelihood of failure) to each potential transplant, and then select a set of exchanges

that maximises the expected number of transplants. The difficulty of determining the likelihood of whether a potential transplant will proceed in practice is well known, and depends on cPRA (calculated Panel Reactive Antibodies, a score between zero and one hundred assigned to a recipient indicating the likelihood that a random blood-group compatible donor will be tissue-type incompatible, see Section 2.1 for more background), but cannot be explained by cPRA alone (Glorie, 2012). Section 4 in this paper provides a number of new insights into this relationship between cPRA and transplant compatibility. Recent models for KEPs have assigned a success probability to each potential transplant, and used these to find a set of exchanges that maximises the number of expected transplants (Dickerson et al., 2012), and further work has also included various recourse options that model fall-back strategies which may be used when transplants cannot proceed (Carvalho et al., 2020; Klimentova et al., 2016).

1.2. Our contribution

We develop three upper bounds for the potential number of transplants in KE instances by utilising the available blood-group information of donors and recipients. We empirically show that the best of these three bounds is always equal to the optimal solution value for the tested PrefLib instances, which implies that these instances can be optimally solved by using our new upper bound together with any lower bounding technique that provides a solution matching it. For the sake of completeness, we also introduce the general framework of one such lower bounding technique, an ad-hoc 7-step matheuristic, whose complete description is available in the appendix. Our approach is able to solve the tested PrefLib instances significantly faster than existing techniques.

We then study the differences between PrefLib instances (and more generally, any instance created by a Saidman generator) and real-world instances we observed in the UK-LKSS, and we highlight distinctions relative to a number of relevant parameters including edge density and maximum number of potential transplants. This allows us to devise new methods for generating random KE instances, based on the Saidman generator but incorporating improvements that can be justified both statistically and by real-world scenarios. These improvements all involve the calculated Panel Reactive Antibody value (cPRA) of a recipient, which estimates what proportion of donors in the KEP will be tissue-type incompatible with that recipient. The improvements arise from examining:

- the distribution of cPRA amongst recipients when grouping recipients by their blood group and the blood group of their donor,
- the relationship between cPRA and transplant compatibility, and
- the transplant compatibility of recipients with a cPRA of zero.

We show that our improved generator creates random KE instances that more closely match real-world instances from the UK than the traditional Saidman generator across a wide variety of relevant parameters. We also confirm that our new upper bounds are not tight on these instances, meaning that our matheuristic does not provide a guarantee of optimality anymore, as is usually required in a KEP.

The rest of this paper is organised as follows. Section 2 gives a background on and defines terminology and notation used in kidney exchange programmes. Section 3 details our new upper bound methods and the framework of our 7-step matheuristic together with some empirical proofs of their effectiveness on the PrefLib instances. Section 4 then introduces our new techniques for generating random KE instances, and we close with some conclusions in Section 5.

2. Background

We present background on compatibility in KEPs in Section 2.1, and then in Section 2.2 we introduce representations of KE instances as well as terminology used to describe particular structures within such instances.

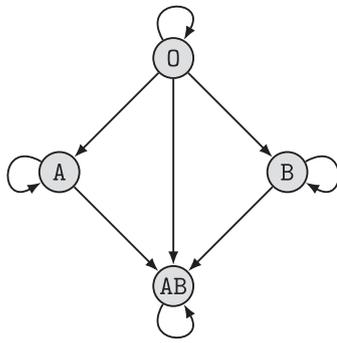


Fig. 1. Blood-group compatibility.

2.1. Compatibility

The compatibility of a transplant between a donor and a recipient is determined by a number of different factors. We will describe two in detail as they are relevant to our work, and will briefly mention some others.

Possibly the most well-known factor is blood-group (or ABO) compatibility. A donor with blood group O (also called universal donor) can potentially give their kidney to a recipient with any blood group. A donor with blood group A (respectively B) can potentially give their kidney to a recipient with blood group A or AB (respectively B or AB). An AB donor can only give their kidney to an AB recipient (also called universal receiver). This compatibility is often presented as a graph such as Fig. 1

Another factor for compatibility is tissue-type compatibility, determined by the presence of human leukocyte antigens (HLA) in a donor and potential presence of HLA antibodies in a recipient. As such, this is also known as HLA compatibility. Tissue-type compatibility is best determined with a lab-based crossmatch test, in which samples from a donor and recipient are combined and examined for reactions (Mulley and Kanellis, 2011). However, such tests are expensive and time-consuming, so it is often not feasible to test every single donor against every single recipient. Instead, some alternatives have been developed. A Panel Reactive Antibody (PRA) test determines a recipient's PRA value by performing crossmatch tests against a given panel of samples. The resulting PRA value of a recipient gives the proportion of donors that a recipient is likely to be incompatible with a percentage rounded to a whole integer (e.g., a recipient with a PRA of 0 is likely to be tissue-type compatible with all donors, while a recipient with a PRA of 95 is likely to be tissue-type compatible with only 5% of donors). A better understanding of HLA incompatibility has allowed for the testing of the presence or absence of individual antigens or antibodies, resulting in the calculated PRA (cPRA) parameter. This is calculated based on not only the presence of antigens in recipients and antibodies in donors, but also a better understanding of which specific antigen and antibody combinations are more likely to result in a failed transplant. Again, the actual cPRA value for a recipient aims to represent the proportion of donors with which the recipient will not be compatible with, and is also rounded to an integer percentage. We highlight as well that HLA incompatibility does come in degrees, and that under certain circumstances, HLA-incompatible transplants can be successfully performed using desensitisation techniques (Jordan et al., 2016). This is of particular importance as we do observe transplants proceeding even when a recipient has a cPRA of 100, which might otherwise seem implausible.

Many other factors influence compatibility, including but not limited to, the ages and general health of the donor and the recipient, the size of the donor kidney, and even potentially the number of veins or arteries present on the donor kidney.

2.2. Notation and terminology

We formally define a Kidney Exchange problem (KE) instance as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ (also called compatibility graph) in which vertex set \mathcal{V} contains every recipient and arc set $\mathcal{A} = \{(i, j) : i, j \in \mathcal{V}\}$ contains the potential transplants: (i, j) belongs to \mathcal{A} if the donor associated to recipient i is compatible with recipient j . Note that compatible recipient/donor pairs (represented by an arc $(i, i) \in \mathcal{A}$) are allowed in some KEPs. A feasible exchange is a cycle in the directed graph: a sequence of distinct vertices (v_1, \dots, v_k) such that there is an arc (v_i, v_{i+1}) for each $i \in \{1, \dots, k-1\}$, and an arc (v_k, v_1) . In an exchange, a donor paired with a recipient only donates a kidney if their paired recipient receives a kidney. For logistical purposes, there is often an upper bound on how many vertices can be part of any exchange (named L in the following).

Some KEPs also include non-directed donors (sometimes referred to as altruistic donors) who are willing to donate a kidney without being paired with a recipient. In the compatibility graph, non-directed donors are represented by a node in \mathcal{V} with no incoming arcs and one outgoing arc to each of the recipients they are compatible with. In the presence of non-directed donors, an exchange can also be a chain, which is a sequence of distinct vertices (v_1, \dots, v_k) where v_1 corresponds to a non-directed donor, and for each $i \in \{1, \dots, k-1\}$ there is an arc (v_i, v_{i+1}) . Usually, the donor associated to recipient v_k gives a kidney to a deceased-donor waiting list.

Some KEPs also allow a recipient to be paired with multiple donors. In Section 3, we focus on the PrefLib instances, in which: (i) every recipient is paired with exactly one donor and every donor is paired with exactly one recipient, (ii) there are no compatible recipient/donor pairs (i.e., no arcs $(i, i) \in \mathcal{A}$), and (iii) there are no non-directed donors. In Section 4 we explicitly allow recipients to be paired with multiple donors, which more closely matches the real-world instances from the UKLSS.

3. Upper bounds and 7-step matheuristic for PrefLib instances

In this section, we describe three upper bounding procedures for KE instances with no non-directed donors and the general framework of our 7-step matheuristic. The first bound U_1 , which is valid for any instance, is derived from a relaxation of the cycle size limit constraints and can be computed in polynomial time. The two other bounds U_2 and U_3 use the donor and recipient blood groups to identify a subset of recipient/donor pairs that cannot participate in a maximum-size exchange, and perform particularly well on PrefLib instances. The 7-step matheuristic takes advantage of the large number of recipient/donor pairs and of the high density in the compatibility graph to divide an instance into several sub-problems where it tries to match recipient/donor pairs of specific blood groups.

3.1. Unbounded cycle lengths upper bound

It is well known that determining the maximum number of transplants in a kidney exchange instance for general values of L is \mathcal{NP} -hard. Abraham et al. (2007) outlined two cases where the problem can be solved in polynomial time:

- When $L = 2$, the compatibility graph can be transformed into an undirected graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ in which edge $\{i, j\}$ belongs to \mathcal{E}' if and only if both arcs (i, j) and (j, i) belong to \mathcal{A} in the original graph. The problem now is to find a matching of maximum size and this can be done in polynomial time using Edmonds' algorithm (Edmonds, 1965).
- When $L = \infty$, the compatibility graph can be transformed into a bipartite graph $\mathcal{G}'' = (\mathcal{L}'', \mathcal{R}'', \mathcal{E}'')$ where \mathcal{L}'' is the set of donors and \mathcal{R}'' is the set of recipients. There is an edge $\{i, j\}$ in \mathcal{E}'' with cost 1 if donor i is compatible with recipient j and an edge

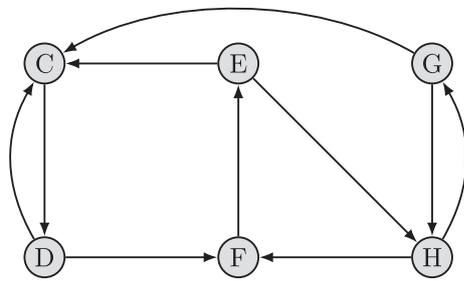


Fig. 2. Compatibility graph for Example 1.

$\{i, j\}$ in \mathcal{E}'' with cost 0 if i and j are in the same recipient/donor pair. The problem now is to find a perfect matching of maximum weight, and again, this can be done in polynomial time using the Hungarian algorithm (Kuhn, 1955).

For a given kidney exchange instance with cycle size limit L , the maximum number of transplants cannot decrease as L increases. As a result, solving the problem with $L = \infty$ gives an upper bound on the maximum number of pairs selected in an exchange for any value of $L \in \mathbb{Z}$. We call this upper bound U_1 .

Example 1. Let us consider the kidney exchange instance with 6 recipients and compatibility graph as represented by Fig. 2.

- If $L \geq 6$, an optimal solution will match the six recipient/donor pairs in either one or two cycles: $\{[G, C, D, F, E, H]\}$ or $\{[C, D, F, E], [H, G]\}$;
- If $L = 4$ or 5 , an optimal solution will match the six recipient/donor pairs in two cycles: $[C, D, F, E]$ and $[H, G]$;
- If $L = 3$, an optimal solution will match five recipient/donor pairs in two cycles: $[E, H, F]$ and $[D, C]$;
- If $L = 2$, an optimal solution will match four recipient/donor pairs in two cycles: $[H, G]$ and $[D, C]$.

3.2. Data-driven upper bound

Our second upper bound U_2 uses blood-group compatibility (see Section 2) to calculate the maximum number of transplants that can possibly take place. Note that blood-group compatibility is a necessary but not sufficient condition for transplant compatibility. In the following, we call:

- A $*/*$ pair a recipient/donor pair in which the recipient and the donor have any blood group, and we denote by n the number of $*/*$ pairs.
- An $X/*$ pair a recipient/donor pair in which the recipient has blood group X and the donor has any blood group, and we denote by $n_{X/*}$ the number of $X/*$ pairs.
- A $*/Y$ pair a recipient/donor pair in which the recipient has any blood group, and the donor has blood group Y and we denote by $n_{*/Y}$ the number of $*/Y$ pairs.
- An X/Y pair a recipient/donor pair in which the recipient has blood group X and the donor has blood group Y , and we denote by $n_{X/Y}$ the number of X/Y pairs.

Upper bound U_2 assumes first that each of the recipient/donor pairs will participate in the exchange unless stated otherwise. It then uses the two following rules to adjust its value:

- Since only cycles are allowed, if an $O/*$ pair p cannot be matched with any $*/O$ pair, then p cannot participate in the exchange.
- Similarly, if a $*/AB$ pair p cannot be matched with any $AB/*$ pair then p cannot participate in the exchange.

Table 1
Blood-group distribution for Example 2.

| Rec./Don. | ./O | ./A | ./B | ./AB | ./* |
|-----------|-----|-----|-----|------|-----|
| O/. | 66 | 134 | 70 | 13 | 283 |
| A/. | 51 | 34 | 59 | 11 | 155 |
| B/. | 10 | 53 | 4 | 1 | 68 |
| AB/. | 3 | 1 | 1 | 1 | 6 |
| */. | 130 | 222 | 134 | 26 | 512 |

Even though it might be difficult to detect with precision which $O/*$ and which $*/AB$ pairs will be excluded from the matching, we know that if $n_{O/*} > n_{*/O}$ (i.e., if there are more $O/*$ than $*/O$ pairs), then at least $n_{O/*} - n_{*/O}$ pairs cannot participate in the exchange. Similarly, we know that if $n_{*/AB} > n_{AB/*}$, then at least $n_{*/AB} - n_{AB/*}$ pairs cannot participate in the exchange. If we now define $\ell_1 = \max\{0, n_{O/*} - n_{*/O}\}$ as the minimum number of $O/*$ pairs that cannot participate in the exchange and $\ell_2 = \max\{0, n_{*/AB} - n_{AB/*}\}$ as the minimum number of $*/AB$ pairs that cannot participate in the exchange, then we can formally introduce the bound $U_2 = n - \ell_1 - \ell_2 + \min\{n_{O/AB}, \ell_1, \ell_2\}$, where the last term identifies the number of O/AB pairs that were counted both in ℓ_1 and in ℓ_2 . Note that a solution with value U_2 is only reachable if the following rules are met: (i) if $\ell_1 > 0$ (which is the case for every PrefLib instance), then every O donor is matched with an O recipient (otherwise more than ℓ_1 pairs of the form $O/*$ would not be able to participate in the exchange), and (ii) if $n_{O/AB}$ is less than or equal to both ℓ_1 and ℓ_2 (which is also the case for every PrefLib instance), then no O/AB pairs should participate in the exchange (otherwise, such a pair should be counted both in ℓ_1 and in ℓ_2). These rules will be used henceforth.

Example 2. Let us consider the kidney exchange instance from the PrefLib repository “MD-00001-00000191”, an instance with $n = 512$ and the blood-group distribution shown in Table 1. Out of the 512 recipient/donor pairs, we know that 153 ($\ell_1 = 283 - 130$) $O/*$ pairs cannot participate in the matching because there are not enough $*/O$ pairs to accommodate them. We also know that 20 ($\ell_2 = 26 - 6$) $*/AB$ pairs cannot participate in the matching because there are not enough $AB/*$ pairs to accommodate them. Out of the 173 removed pairs, we know that up to 13 O/AB pairs could have potentially been counted twice (i.e., both in ℓ_1 and ℓ_2). This gives an upper bound U_1 equal to 352 ($512 - 153 - 20 + 13$). An optimal solution to this instance has size 351 if $L = 3$, and size 352 if $L \geq 4$, showing that, for this particular instance, the bound is very close to the optimal solution value.

3.3. ILP-based upper bound

After some preliminary analysis, we found that the difference observed between upper bounds U_2 and the optimal solution values of the tested PrefLib instances always came from some $*/AB$ pairs that we assumed could be part of the exchange while in reality they could not (in other words, we underestimated ℓ_2). This can be explained by the fact that AB is the least common blood group, so there are fewer opportunities to match $*/AB$ pairs with $AB/*$ pairs. In order to correct that inaccuracy, we designed a tailored ILP model that maximises the number of X/AB pairs that can possibly be included in the exchange ($X \in \{A, B, AB\}$) to determine a refined upper bound U_3 . Note that the other rules adopted from U_2 , namely (i) that any $*/O$ pair must be followed by an $O/*$ pair in the exchange, and (ii) that no O/AB pairs should be included in the exchange, are still in place. Every time one of these two rules is broken, the exchange size is reduced by one unit. In the following, we first describe the ILP model details for $L = 3$ and then we explain how to modify it for $L = 4$. The model needs a list of every cycle solely composed of $*/AB$ pairs. We distinguish:

- C_2^2 , the set of 2-cycles with two $*/AB$ pairs – these can only be in the form $[AB/AB - AB/AB]$.

- C_3^3 , the set of 3-cycles with three */AB pairs – these can only be in the form [AB/AB - AB/AB - AB/AB].

For instance, if all the pairs in Example 1 were AB/AB pairs, then there would be two cycles in C_2^2 : [C, D] and [G, H], and one cycle in C_3^3 : [E, H, F]. The model also needs a list of every cycle partly composed of */AB pairs. We distinguish:

- P_2^1 , is the set of *partial* cycles comprising two pairs, among which exactly one is a */AB pair – these are in the form [X/AB - AB/Y], where $X \in \{A, B, AB\}$ and $Y \in \{O, A, B\}$. These cycles do not necessarily need to be complete, so we make sure that either the donor of the second pair is compatible with the recipient of the first pair, or that there exists at least one additional pair that can complete the cycle;
- C_3^2 , the set of 3-cycles with exactly two */AB pairs – these are in the form [X/AB - AB/AB - AB/Y], where $X \in \{A, B, AB\}$ and $Y \in \{A, B\}$;

By introducing binary decision variables x_c taking value 1 if cycle c is selected and 0 otherwise ($c \in C$, where $C = \{C_2^2 \cup C_3^3 \cup P_2^1 \cup C_3^2\}$), the maximum number of X/AB pairs that can possibly be included in the exchange ($X \in \{A, B, AB\}$) denoted as z can be obtained through the following ILP model:

$$\max \quad z = \sum_{c \in C_2^2} 2x_c + \sum_{c \in C_3^3} 3x_c + \sum_{c \in P_2^1} x_c + \sum_{c \in C_3^2} 2x_c \quad (1)$$

$$\text{s.t.} \quad \sum_{c \in C: i \in V(c)} x_c \leq 1, \quad \forall i \in \mathcal{V}, \quad (2)$$

$$x_c \in \{0, 1\}, \quad c \in C, \quad (3)$$

where $V(c)$ indicates the set of recipient/donor pairs that belong to cycle c . Upper bound U_3 can be defined as $U_3 = n - \ell_1 - \ell_2'$ where $\ell_2' = (n_{*/AB} - z)$. Note that it is not necessary to consider cycles [O/AB - AB/Y] in P_2^1 because selecting such cycle would increase z (and thus, U_3) by one unit at best, but would also break one of our two rules, decreasing by one unit the exchange size. Similarly, it is not necessary to consider cycles [O/AB - AB/AB - AB/O] in C_3^2 because selecting such cycle would only increase U_3 by one unit (two units in z with a one unit penalty due to the broken rule), which is as effective as selecting the truncated [AB/AB - AB/O] cycle that already exists in P_2^1 .

The model can be extended to take into account a cycle size limit of four. In this case, we need:

- An additional set of cycles C_4^4 with four */AB pairs – these can only be in the form [AB/AB - AB/AB - AB/AB - AB/AB].
- An updated set of partial cycles P_2^1 with two pairs having exactly one */AB pair in which we additionally make sure that either the donor of the second pair is compatible with the recipient of the first pair, or that there exists at least one pair or a couple of pairs that can complete the cycle.
- An updated set of partial cycles P_3^2 with three pairs having exactly two */AB pair – these are in the form [X/AB - AB/AB - AB/Y] where $X \in \{A, B, AB\}$ and $Y \in \{O, A, B\}$. These cycles do not necessarily need to be complete, so we make sure that either the donor of the third pair is compatible with the recipient of the first pair, or that there exists at least one pair that can complete the cycle.
- An additional set of 4-cycles C_4^3 with exactly three */AB pairs [X/AB - AB/AB - AB/AB - AB/Y], where $X \in \{A, B, AB\}$ and $Y \in \{A, B\}$.

3.4. Upper bound performance on PrefLib instances

We tested each of the bounds on a set of 30 PrefLib instances without non-directed donors, which is the subset that was used by Lam and Mak-Hau (2020) for testing their branch-and-price-and-cut algorithm. The results are available in Table 2. The ‘‘Instance’’ column gives the

name of the instance, column n gives the number of recipient/donor pairs in the instance, the four following columns give the upper bounds and optimal value for $L = 3$, and the last four columns give the same information for $L = 4$. The optimal values were either obtained from Lam and Mak-Hau (2020) or were deduced from a lower bound obtained in Lam and Mak-Hau (2020) combined with one of the upper bounds we introduced.

We observe that bound U_1 displays a very good quality, as it is equal to the optimal solution value in 27 instances out of 30 for $L = 3$ and on all instances for $L = 4$. In fact, by simply using U_1 with the solution obtained by the algorithm of Lam and Mak-Hau (2020) for $L = 3$, we could solve the 11 benchmark instances with $L = 4$ that they left open in their paper. Upper bound U_2 is also very promising, finding the optimal objective value in 39 instances out of 60. It is not surprising that the values of U_1 and U_2 do not change when the cycle size limit increases as L is not taken into account when the bounds are computed. We also notice that upper bound U_3 displays outstanding performance on the PrefLib instances as it is always equal to the optimal solution value both for $L = 3$ and $L = 4$. It is worth mentioning that each bound was calculated in less than one second of computing time. The experiments were run on an Intel Xeon E5-2680W v3, 2.50 GHz with 192 GB of memory, running under Scientific Linux 7.5, and using Gurobi 7.5.2 to solve the ILP models.

3.5. 7-Step matheuristic framework for solving PrefLib instances

While the tested PrefLib instances could theoretically be solved by using any lower bounding technique whose solution matches U_3 for every instance, to the best of our knowledge, there is no such ad-hoc heuristic available in the literature. To show that such a lower bounding technique does not necessarily take a long computation time, we propose in the following the framework of a 7-step matheuristic tailored for solving PrefLib instances and show its efficiency. When used with U_3 , it is able to solve all the tested instances in a few seconds, which is much faster than the branch-and-price algorithm of Lam and Mak-Hau (2020), who also tested their approach on PrefLib instances.

Matheuristic is a neologism originating from the contraction of ‘‘math’’ and ‘‘heuristic’’ and is often used as an alternative to mathematical models when the latter become too large to solve difficult combinatorial optimisation problems. Matheuristics usually obtain good quality solutions but do not have a proof of optimality. A common approach in matheuristics is to solve a mathematical model with a reduced set of variables once or several times in a row. It is the path we follow with our 7-step matheuristic as we split every instance into several sub-problems in which we try to match recipient/donor pairs of specific blood groups as described in Algorithm 1.

Algorithm 1: 7-step matheuristic.

- 1: Identify the types of cycles and recipient/donor pairs used in the exchange and process the AB/* and */AB pairs
 - 2: Process the B/B pairs
 - 3: Process the A/B and B/A pairs
 - 4: Process the A/A pairs
 - 5: Process the O/O pairs
 - 6: Process the B/O and O/B pairs
 - 7: Process the A/O and O/A pairs
-

The matheuristic starts by determining the types of cycles that should be used if we want to reach a solution of value U_3 (e.g., x_1 cycles of the form [A/B-B/A], x_2 cycles of the form [A/B-B/O-O/A], ...). Then, for each combination of recipient/donor pairs, it uses the compatibility graph to enumerate every feasible cycle or partial cycle of the given types and uses an ILP formulation to select those that are necessary to reach the desired solution. At each step, we define:

- An *objective*, which is the goal of the ILP model that is solved at the current step.

Table 2
 U_1, U_2 , and U_3 values for PrefLib instances with $L = 3$ and $L = 4$.

| Instance | n | $L = 3$ | | | | $L = 4$ | | | |
|-------------------|------|---------|-------|-------|------|---------|-------|-------|------|
| | | U_1 | U_2 | U_3 | OPT | U_1 | U_2 | U_3 | OPT |
| MD-00001-00000191 | 512 | 352 | 352 | 351 | 351 | 352 | 352 | 352 | 352 |
| MD-00001-00000192 | 512 | 337 | 337 | 337 | 337 | 337 | 337 | 337 | 337 |
| MD-00001-00000193 | 512 | 300 | 301 | 300 | 300 | 300 | 301 | 300 | 300 |
| MD-00001-00000194 | 512 | 313 | 313 | 313 | 313 | 313 | 313 | 313 | 313 |
| MD-00001-00000195 | 512 | 322 | 322 | 322 | 322 | 322 | 322 | 322 | 322 |
| MD-00001-00000196 | 512 | 313 | 313 | 313 | 313 | 313 | 313 | 313 | 313 |
| MD-00001-00000197 | 512 | 334 | 337 | 334 | 334 | 334 | 337 | 334 | 334 |
| MD-00001-00000198 | 512 | 332 | 333 | 332 | 332 | 332 | 333 | 332 | 332 |
| MD-00001-00000199 | 512 | 313 | 314 | 313 | 313 | 313 | 314 | 313 | 313 |
| MD-00001-00000200 | 512 | 312 | 313 | 312 | 312 | 312 | 313 | 312 | 312 |
| MD-00001-00000231 | 1024 | 659 | 659 | 658 | 658 | 659 | 659 | 659 | 659 |
| MD-00001-00000232 | 1024 | 662 | 662 | 662 | 662 | 662 | 662 | 662 | 662 |
| MD-00001-00000233 | 1024 | 635 | 636 | 635 | 635 | 635 | 636 | 635 | 635 |
| MD-00001-00000234 | 1024 | 641 | 642 | 641 | 641 | 641 | 642 | 641 | 641 |
| MD-00001-00000235 | 1024 | 660 | 660 | 660 | 660 | 660 | 660 | 660 | 660 |
| MD-00001-00000236 | 1024 | 655 | 655 | 655 | 655 | 655 | 655 | 655 | 655 |
| MD-00001-00000237 | 1024 | 597 | 597 | 597 | 597 | 597 | 597 | 597 | 597 |
| MD-00001-00000238 | 1024 | 672 | 672 | 671 | 671 | 672 | 672 | 672 | 672 |
| MD-00001-00000239 | 1024 | 673 | 674 | 673 | 673 | 673 | 674 | 673 | 673 |
| MD-00001-00000240 | 1024 | 639 | 639 | 639 | 639 | 639 | 639 | 639 | 639 |
| MD-00001-00000271 | 2048 | 1284 | 1284 | 1284 | 1284 | 1284 | 1284 | 1284 | 1284 |
| MD-00001-00000272 | 2048 | 1249 | 1249 | 1249 | 1249 | 1249 | 1249 | 1249 | 1249 |
| MD-00001-00000273 | 2048 | 1232 | 1232 | 1232 | 1232 | 1232 | 1232 | 1232 | 1232 |
| MD-00001-00000274 | 2048 | 1242 | 1242 | 1242 | 1242 | 1242 | 1242 | 1242 | 1242 |
| MD-00001-00000275 | 2048 | 1301 | 1301 | 1301 | 1301 | 1301 | 1301 | 1301 | 1301 |
| MD-00001-00000276 | 2048 | 1311 | 1311 | 1311 | 1311 | 1311 | 1311 | 1311 | 1311 |
| MD-00001-00000277 | 2048 | 1316 | 1316 | 1316 | 1316 | 1316 | 1316 | 1316 | 1316 |
| MD-00001-00000278 | 2048 | 1268 | 1268 | 1268 | 1268 | 1268 | 1268 | 1268 | 1268 |
| MD-00001-00000279 | 2048 | 1271 | 1272 | 1271 | 1271 | 1271 | 1272 | 1271 | 1271 |
| MD-00001-00000280 | 2048 | 1295 | 1295 | 1295 | 1295 | 1295 | 1295 | 1295 | 1295 |

- A list of *cycles to complete*, which contains the incomplete cycles determined at a previous step that need to be completed at the current step if we want to reach a solution of value U_3 because it is the last opportunity to do so.
- A list of *pairs to use*, which contains the pairs that need to be assigned to a cycle at the current step if we want to reach a solution of value U_3 because it is the last opportunity to do so.
- a list of *types of cycles allowed*, which contains the types of cycles that are enumerated in the ILP formulation at the current step;
- A list of *types of incomplete cycles allowed*, which contains the types of incomplete cycles (i.e., with no compatibility requirements between the donor of the second pair and the recipient of the first pair) that are enumerated in the ILP formulation at the current step.

The model is inspired by the cycle formulation (see Roth et al., 2007) and associates a binary variable to every cycle and every incomplete cycle allowed. It maximises the defined “objective” under the constraints that every “cycle to complete” must be completed and every “pair to use” must be assigned to a cycle. Note that the cycles selected at a given step are assumed to be fixed at each subsequent step while the partial cycles selected at a given step must be completed (but cannot be broken) in a subsequent step. A summary of every step of the matheuristic is presented in Table 3 for the case where the cycle size limit L is set to 3. Additional details about each of the matheuristic steps, some justification regarding the step ordering, some technical aspects about the absence of certain cycle types, an explanation as to how the matheuristic can be extended to handle a larger cycle size limit $L = 4$, and an example are available in the Appendix.

3.6. Matheuristic performance on PrefLib instances and discussion

We measured the performance of our matheuristic on the 30 PrefLib instances used by Lam and Mak-Hau (2020) for testing their branch-and-price-and-cut algorithm and the results are displayed in Table 4.

Column “ L ” gives the cycle size limit, column “ n ” gives the number of recipient/donor pairs in the instance, the three following columns indicate the number of optimal solutions found by the matheuristic (out of 10), the average lower bound (which was always equal to U_3 in all tested instances), and the computing time in seconds (using the same computational environment as the one described in Section 3). The six following columns report the number of optimal solutions found and the computing time required by two ILP formulations and the branch-and-price-and-cut as reported in the experiments of Lam and Mak-Hau (2020).

We observe that the matheuristic can solve every instance in a few seconds. In particular, it is able to solve the 11 instances left open by Lam and Mak-Hau (2020). Even though these results are very good, a word of caution is in order: the matheuristic is tailored for large-size PrefLib instances. These are characterised by a very large graph density (around 25%) and a relatively large number of recipient/donor pairs, which makes it very likely that the cycles predicted in the first step can be constructed in subsequent steps. Small-size PrefLib instances could not be solved by the matheuristic but were qualified as “trivial” by Lam and Mak-Hau (2020) (i.e., solvable in seconds by an ILP formulation). Real-world instances from the UKLKSS seem to indicate that graph densities are usually smaller (around 10%).

The main objective of this section is to show that benchmark instances, which are important for developing new algorithms, can easily favour one kind of algorithm. PrefLib instances have a high graph density, which penalises the methods that enumerate every feasible cycle (such as the cycle formulation and PICEF) and benefits methods using lower and upper bounding techniques (like our matheuristic). In the following, we suggest some improvements for the Saidman generator (that was used for generating the PrefLib instances) in order to make it produce instances that match more closely the real-world instances of the UKLKSS.

Table 3
7-step matheuristic overview for $L = 3$.

| Step | Focus | Input | | Output | |
|------|--------------|------------------------------|--------------|---|--|
| | | Cycles to complete | Pairs to use | Cycles allowed | Incomplete cycles allowed |
| 1 | * /AB, AB/ * | - | - | [X/AB - AB/AB - AB/Y] where $X \in \{A, B, AB\}$ and $Y \in \{A, B\}$ | [X/AB - AB/Y] where $X \in \{A, B, AB\}$ and $Y \in \{O, A, B\}$ |
| 2 | B/B | [B/AB - AB/B] | - | [B/AB - AB/B], [B/AB - AB/B - B/B], [B/B - B/B], [B/B - B/B - B/B] | - |
| 3 | A/B, B/A | [A/AB - AB/B], [B/AB - AB/A] | [B/B - B/B] | [A/AB - AB/B - B/A], [B/AB - AB/A - A/B], [A/B - B/A], [A/B - B/B - B/A] | [A/B - B/O], [B/A - A/O], [A/B - B/A] |
| 4 | A/A | [A/AB - AB/A] | - | [A/AB - AB/A], [A/AB - AB/A - A/A], [B/A - A/A - A/B], [A/A - A/A] | - |
| 5 | O/O | - | - | [O/O - O/O] | - |
| 6 | B/O, O/B | [B/AB - AB/O], [B/A - A/O] | - | [B/AB - AB/O - O/B], [B/O - O/B], [B/O - O/O - O/B], [B/A - A/O - O/B] | - |
| 7 | A/O, O/A | [A/AB - AB/O], [A/B - B/O] | O/O, A/A | [A/AB - AB/O - O/A], [A/B - B/O - O/A], [A/O - O/O - O/A], [A/O - O/A], [A/O - O/A - A/A] | - |

Table 4
Matheuristic performance on PrefLib instances.

| Parameters | | Matheuristic | | | PIEF | | PICEF/Cycle | | B-P-C | |
|------------|------|--------------|--------|------|------|------|-------------|------|-------|------|
| L | n | OPT | LB | Time | OPT | Time | OPT | Time | OPT | Time |
| 3 | 512 | 10 | 322.7 | 0.2 | 10 | 8 | 10 | 82 | 10 | 1 |
| 3 | 1024 | 10 | 649.1 | 1.4 | 4 | 1304 | 5 | 1308 | 10 | 0 |
| 3 | 2048 | 10 | 1276.9 | 3.0 | 0 | 1800 | 0 | 1800 | 10 | 5 |
| 4 | 512 | 10 | 322.8 | 0.2 | 0 | 1800 | 0 | 1800 | 10 | 21 |
| 4 | 1024 | 10 | 649.3 | 0.7 | 0 | 1800 | 0 | 1800 | 9 | 800 |
| 4 | 2048 | 10 | 1276.9 | 3.4 | 0 | 1800 | 0 | 1800 | 0 | 1800 |

4. Improved generation of KE instances

This section introduces a new method for generating KE instances, based on the Saidman generator. This generator, which we describe in Section 4.2, is widely used to generate random KE instances for evaluating both computational techniques and policy changes. We develop a number of improvements to the Saidman generator, and then show that using our refined version creates instances that are substantially closer to the real-world instances from the UK in a number of important measures (such as number of edges or sizes of solutions) when compared to those constructed by the original Saidman generator. Each of our improvements is based on statistical evidence from the UKLSS, and accompanied by a justification based on real-world scenarios. This is followed by an experimental evaluation of these improvements, and a conclusion discussing the outcomes of the experiments.

4.1. Performance evaluation and data presentation

We evaluate the performance of our improvements by generating random instances and comparing them to real-world data from the UKLSS. We gathered all of our parameters, including distributions of blood groups, cPRA and the number of donors paired with a recipient from UKLSS instances from January 2012 to October 2019 (32 instances in total as UKLSS runs matchings quarterly), and these are summarised in Appendix B. As these real-world instances do not have a constant size, we cannot report the results of our experiments by comparing the averages of various parameters. Instead, for each real-world instance, we create 20 random instances with the same number of recipients and non-directed donors (the precise number of directed donors is determined by the generator as recipients may be paired with multiple donors). For each such randomly-generated instance and its associated real-world instance, and for each parameter we measure, we calculate the ratio of the value of parameter in the random instances divided by the value of the parameter in the real-world instance (e.g., if the random instance has 5000 edges, but the real-world instance has 4000 edges, we report the ratio $\frac{5000}{4000} = 1.25$). As each instance creates one ratio, this results in a set of 640 ratios for

each parameter measured. We then plot the distribution of these ratios as box-and-whisker plots. These plots show a coloured box denoting the interquartile range (IQR), a black vertical bar denoting the median, and a white cross denoting the mean. The horizontal lines extend a distance of 1.5 IQR from either side of the coloured box (if no data points extend this far, these lines stop at the last data point), with data points further away being marked as outliers with solid dots. We limit data and plots presented in this section to show only the ratios corresponding to the number of edges in the compatibility graphs and the size of the largest set of identified transplants. We include in Appendix C results on a much larger variety of parameters. The generator we develop is accessible from <https://wpetersson.github.io/kidney-webapp/#/generator>, and can be used in conjunction with the parameters described in this section and the parameters in Appendix B. The source for this generator is available from <https://github.com/wpetersson/kidney-webapp/>. The instances we generated are available from <http://dx.doi.org/10.5525/gla.researchdata.1213>.

4.2. The Saidman generator

The Saidman generator (Saidman et al., 2006) is widely used to generate instances of KEP problems (Constantino et al., 2013; Klimentova et al., 2016). We now describe how the Saidman generator functions, and then compare the output from the Saidman generator to real world instances.

4.2.1. Generation process

The Saidman generator takes as input the number of recipients, the number of non-directed donors, as well as the distributions of the donor and recipient blood groups (these may be identical, if distinct distributions are not known), the distribution of number of donors per recipient¹ and the distribution of cPRA values of the recipients. These distributions are commonly gathered from a long-running KEP; we will

¹ The UKLSS, and many other KEPs, allows a recipient to join a KEP with multiple donors, so as to increase their probability of finding a match.

be using values from the UKLKSS. The Saidman generator first produces donor–recipient pairs by independently drawing a blood group and cPRA for a new recipient, as well as drawing the number of donors paired with this recipient, and a blood group for each of these donors.

For any donor and recipient, their compatibility is determined as follows:

1. If the donor is blood-group incompatible with a given recipient, then they are not compatible.
2. Otherwise, draw uniformly at random a number r from $[0, 1]$. The donor and recipient are compatible if and only if $r > \frac{cPRA}{100}$.

If a donor is determined to be compatible with their paired recipient, the Saidman generator will discard both the recipient and all paired donors to avoid the generation of compatible pairs. Whilst never specifically addressed in the literature, it is commonly accepted that this is done so that the pool of donor–patient pairs more accurately reflects the real-world data. For instance, the proportion of type AB recipients in kidney exchange pools is often lower than the proportion of people with type AB in a general population, as it is easier for type AB recipients to find a willing and compatible donor because they are blood-group compatible with all potential donors. Discarding compatible pairs mimics this effect, and similar effects for recipients with a low cPRA.

Once sufficient recipients have been generated, the Saidman generator then draws a blood group for each non-directed donor required.

When this is complete, the compatibility of any donor with any recipient (other than their paired recipient, if applicable) is determined using the above algorithm, and a corresponding arc is added to the compatibility graph if the donor and the recipient are compatible.

There are two important factors to note when discussing the Saidman generator. Firstly, the blood groups of donors and recipients and the cPRA values are all assumed to be independent. This is but one of the aspects we will address later in this section. Secondly, the discarding of compatible pairs means that the distributions of the generated data will not match the input distributions (i.e., low cPRA recipients are more likely to be discarded, so low cPRA recipients will be under-represented in the generated data). To avoid this, the inputs to the Saidman generator can be tuned. This can be achieved by running the generator to produce an instance which is then analysed for some property or set of properties (e.g., the distribution of cPRA levels amongst recipients). If this property is deemed not sufficiently close to some desired target (e.g., the proportion of recipients with a low cPRA is deemed too low), then the inputs can be adjusted (e.g., the probability of a recipient drawing a low cPRA can be increased), and the process is repeated. Alternative tuning options include sweeping across a broad range of parameters to find the input distributions that give output distributions deemed best (Santos et al., 2017). In this paper, we tune the distributions of donor blood groups, recipient blood groups, recipient cPRA levels, and number of donors paired with a recipient.

4.2.2. Performance of the saidman generator

Recall from Section 4.1 that we generate 640 random instances, and for each random instance and each parameter we calculate the ratio comparing the parameter in the random instance to the parameter in the real-world instance with the same size. Fig. 3 gives boxplots of these ratios for two parameters, showing how the Saidman generator compares to the real-world data. Better generators therefore have means and medians closer to one, and with less variance, although some variance is to be expected. We see that the Saidman generator produced instances with substantially more edges than present within the UKLKSS data, and the largest set of transplants in an instance generated by the Saidman generator tends to be significantly larger than the largest set of transplants in an instance from the UKLKSS. A comparison of a larger set of parameters is given in Appendix C.

4.3. New generator configurations

In the rest of this section we introduce changes based on three different assumptions made by the Saidman generator. These are: the independence of donor and recipient blood groups and the recipient cPRA, the direct linear relationship between cPRA and donor compatibility, and the equivalence of all recipients with a given cPRA value.

4.3.1. Donor and recipient blood-groups and cPRA distributions

As noted earlier, the Saidman generator assumes that the blood groups of the donors and recipients, and the cPRA levels of the recipients, are independent. From a clinical point of view, however, it would not be surprising that recipients paired with blood-group compatible donors would be more likely to have tissue-type incompatibility, as donors and recipients who are both blood-group and tissue-type compatible are likely to be medically compatible and thus have proceeded with a transplant outside of a KEP. We ran a t -test of the corresponding Wald statistic (Buse, 1982) to compare our hypothesis that there is a correlation between a recipient having a blood-group compatible donor and the recipient's cPRA level against a null hypothesis. The resulting p -value of approximately 1.79×10^{-7} is a strong indicator that we can reject the null hypothesis, suggesting that different cPRA distributions should be used when generating recipients depending on whether a recipient has been generated with a blood-group compatible donor or not. We show in Fig. 4 the distributions of cPRA levels of recipients from the UKLKSS broken down according to whether or not the recipient is paired with a blood-group compatible donor.

An obvious requirement for modelling this phenomenon is the ability to accurately model the correct proportion of blood-group compatible pairs. To achieve this, we calculate the distribution of blood groups of the recipients. We then calculate five distributions of donor blood groups, one for each of the four possible recipient blood groups, and one for non-directed donors. This data is shown in Fig. 5, and we see that there are substantial differences between these distributions.

We implement a generator that includes these adaptations, and use the term SplitPRA to refer to this adaptation. Two box plots comparing the performance of the adaptation SplitPRA to the Saidman generator are shown in Fig. 6. We see from these that SplitPRA creates random instances that are closer to the real-world instances than those produced by the Saidman generator both in terms of number of edges and size of largest set of transplants.

4.3.2. Link between cPRA and transplant compatibility

For any donor and recipient who are blood-group compatible, the Saidman generator uses cPRA as a proxy for transplant compatibility. In real-world applications, transplant compatibility is closely tied to tissue-type compatibility (and cPRA is a good measure of this), but there are also external factors such as the size of the kidney, or the age or overall health of the donor. As such factors are much harder to randomly sample, the Saidman generator draws transplant compatibility uniformly at random with probability $(1 - cPRA/100)$. However, cPRA by definition only considers tissue-type compatibility, and ignores any other factors that may affect transplant compatibility. Additionally, while a cPRA value of 100 might seem to indicate that transplantation will be impossible, this is not necessarily the case. A nephrologist may allow the presence of certain “incompatible” HLA antibodies when determining whether to assess a potential donor for their recipient, as the risk associated with such transplants² may be worth taking considering the risk of waiting even longer for an ideal donor that may never arrive. Secondly, cPRA is a measure of how likely a random donor will be compatible with a given recipient, and is reported after being rounded

² Such transplants will often involve desensitisation techniques to avoid the rejection of the donated kidney.



(a) Distribution of the ratio of number of edges in randomly generated instances



(b) Distribution of the ratio of largest set of transplants in randomly generated instances.

Fig. 3. Boxplots outlining the performance of the Saidman generator.

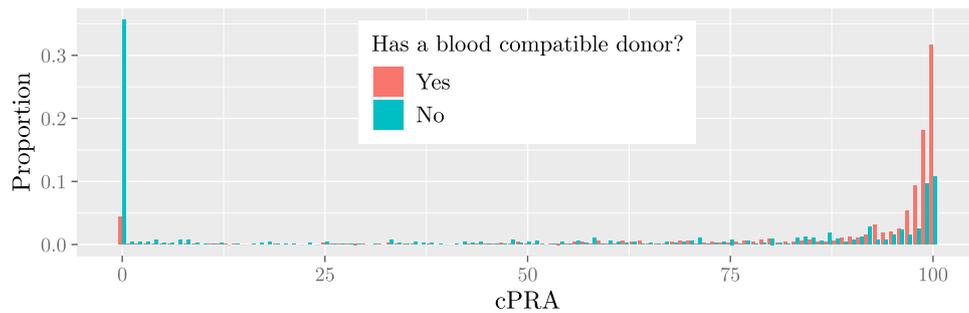


Fig. 4. Distributions of cPRA dependent on whether a recipient has a blood-group compatible donor.

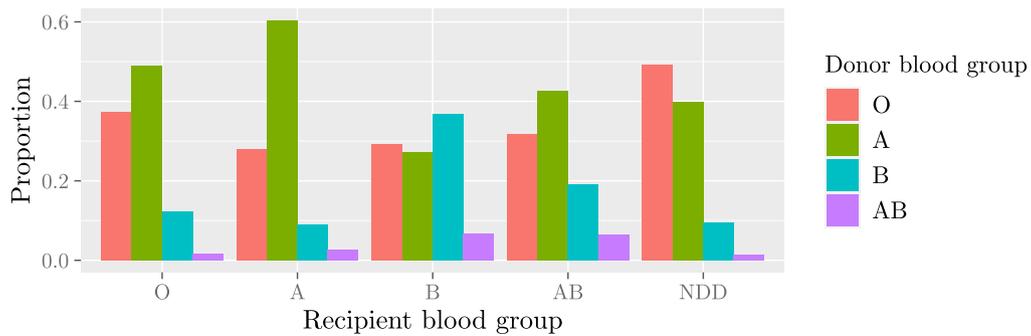
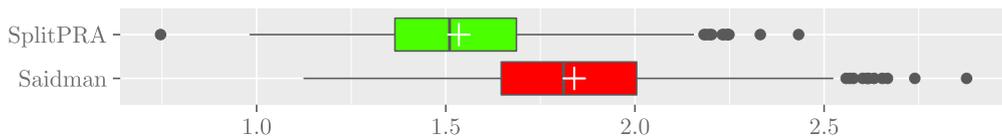
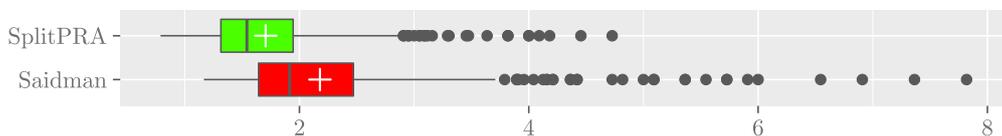


Fig. 5. Distributions of donor blood group by recipient blood group.



(a) Distribution of number of edges as proportion of edges in real-world instances



(b) Size of largest set of transplants

Fig. 6. Comparison of Saidman and SplitPRA generators.

to the nearest integer, which means that outliers are always possible. Indeed, we do see that in the real-world data recipients with a cPRA of 100 are compatible with approximately 1% of blood-group compatible donors. This highlights the fact that cPRA alone does not determine incompatibility, but merely is an indicator of incompatibility.

For each recipient in the UKLKSS data, we calculate the proportion of blood-group compatible donors appearing in a matching run simultaneously with the recipient that are also identified as potential donors for the recipient. We plot this data in Fig. 7. From this, it is clear to see that while there is a strong negative relationship between cPRA and compatibility (confirmed with a Spearman rank correlation coefficient of approximately -0.519 , which corresponds to a p -value of $<2.2 \times 10^{-16}$), it is not necessarily as simple as $(1 - \text{cPRA}/100)$. Indeed, a linear least-squares model weighted by the distribution of cPRA amongst the recipients gives the equation $P(\text{Compat.}) = 0.58 - 0.55 \times \text{cPRA}/100$. This linear model is shown as a solid blue line in Fig. 7, with the shaded region showing the 95% confidence interval. The triangles denote each data point, with one data point per cPRA value. Note that the number of samples represented by each data point is not visually represented in Fig. 7 to make the plot easier to read. As such, the reader should be aware that the level of correlation between cPRA and the variance of compatibility as seen in Fig. 7 may be misleading. By looking at the data we see that while there is a greater variance at lower values of cPRA (we discuss this in detail for recipients with a cPRA of exactly 0 in Section 4.3.3), much of the variance visually depicted in Fig. 7 is due to the relative difference in populations represented at each cPRA value (see Fig. 4 which shows the relative distribution of cPRA values among recipients).

We see that this linear model does approximate the relationship between cPRA and compatibility, but many data points are outside the 95% confidence interval. We also highlight that recipients with a cPRA of 0 (i.e., those recipients who theoretically should have no tissue-type compatibility problems) are only compatible with approximately 58% of the blood-compatible donors in their matching runs. We investigate this particular aspect further in Section 4.3.3. Other studies model the link between cPRA and tissue-type compatibility through the cumulative distribution function (CDF) of a standard normal distribution (Glorie, 2014). However, these studies only consider tissue-type compatibility, whereas we are aiming to model actual compatibility, giving one potential explanation for this difference. Based on the data we observe in Fig. 7, neither a linear relationship nor CDF nor any straight-forward model can accurately represent the link between cPRA and transplant compatibility, so we opt for the linear model which is able to more closely match the low but non-zero probability of compatibility that is observed with recipients with a cPRA of 100.

We test an adaptation to the generator that uses $P(\text{Compat.}) = 0.58 - 0.55 \times \text{cPRA}/100$ (i.e., the linear least-squares model) to determine compatibility between blood-compatible donors and recipients; we call this CalcXMatch. Full experimental results are shown in Appendix C, but Fig. 8 shows that CalcXMatch does reduce the number of edges (bringing it closer to the real-world instances) although it does so at the cost of significantly increasing the number of transplants identified. A closer examination of the compatibility equation shows that CalcXMatch will give recipients with a cPRA of 100 a 3% chance of compatibility, whereas the Saidman generator gives such recipients a 0% chance of compatibility. Thus, CalcXMatch has fewer incoming edges from the recipients with a low cPRA, and more incoming edges to recipients with cPRA values above 98, compared to the Saidman generator, resulting in more identified transplants.

We test a second adaptation that uses the formula $P(\text{Compat.}) = 0.55 - 0.55 \times \text{cPRA}/100$ to determine compatibility (forcing recipients with a cPRA of 100 to have a 0% chance of compatibility). We call this adaptation TweakXMatch, and its performance is included in Fig. 8. This, as expected, reduces the number of edges, but brings the mean number of edges very close to the mean number of edges in the real-world data. Whilst this looks impressive, we do highlight that

this particular generator, like the Saidman generator and just using SplitPRA, gives a recipient with a cPRA of 100 zero chance of being compatible, and so no recipients with a cPRA of 100 are ever identified for transplantation.

We also create a third adaptation that does not use one equation to determine compatibility for all recipients. Instead, we split the range of potential cPRA values into the bands $[0, 50)$, $[50, 95)$, $[95, 96)$, $[96, 97)$, $[97, 98)$, $[98, 99)$, $[99, 100)$ and 100. Then a linear regression is performed on the real-world recipients from the UKLKSS who fit into each band of cPRA values, giving one equation (or value, if the band contains just one cPRA value) per cPRA band. These equations are then used to determine compatibility for generated recipients, based on their cPRA values. We call this configuration BandXMatch, and its performance is also included in Fig. 8. We see that, as expected, it has a similar number of edges to CalcXMatch, but on average it has a smaller-sized largest set of transplants.

4.3.3. The compatibility of cPRA 0 recipients

One oddity that arises from Fig. 7 is that recipients with a cPRA of 0 are only compatible with approximately 58% of their blood-group compatible donors. This is somewhat less than may have been anticipated, so we investigated it further. Fig. 9 shows a histogram of the compatibility of real-world recipients with a cPRA of 0, where compatibility is calculated as the proportion of potential transplants in the instances divided by the number of blood-group compatible donors in said instances. We see that approximately 18% of these cPRA 0 recipients are compatible with all of their blood-group compatible donors, and that approximately 40% of these recipients are compatible with at least 95% of their blood-group compatible donors. However, almost 26% of recipients with a cPRA of 0 are compatible with less than 25% of their blood-compatible donors.

Some hypotheses have been presented for this phenomenon:

- Paediatric recipients may have stricter compatibility requirements, beyond the simple blood and tissue-type compatibility, or,
- if a recipient is known to have a low (or zero) cPRA, then this recipient (or their clinician) may be more stringent with regards to other parameters such as age or health of the donor.

We account for this phenomenon in our generator as follows. For each recipient with a cPRA of 0, we randomly sample from a smoothed variant of the distribution shown in Fig. 9 the compatibility proportion of this recipient. This compatibility proportion is then used directly as the probability that the recipient is compatible with a blood-group compatible donor. In this way, we generate instances wherein approximately 40% of recipients with a cPRA of 0 are compatible with over 95% of their blood-group compatible donors, and so on.

Fig. 10 shows the performance of this adaptation, which we denote by PRA0 in the plots. We see that when combined with either TweakXMatch or BandXMatch, this adaptation slightly reduces both the number of edges and size of the largest set of transplants.

4.4. Discussion

We have demonstrated a number of improvements to the process of generating random KE instances. We have shown that the recipients who have a blood-group compatible donor have a markedly different distribution of cPRA values to those who do not have a blood-group compatible donor, and that utilising this when generating instances will result in random instances that more closely resemble real-world instances. This approach can be utilised by simply calculating and using two distinct cPRA distributions for the two populations. We have also shown that the correlation between cPRA and transplant compatibility, whilst present, is not necessarily as simple as $P(\text{compatibility}) = 1 - \text{cPRA}/100$, and taking a more nuanced approach results in random instances that more closely resemble real-world instances. Lastly, we

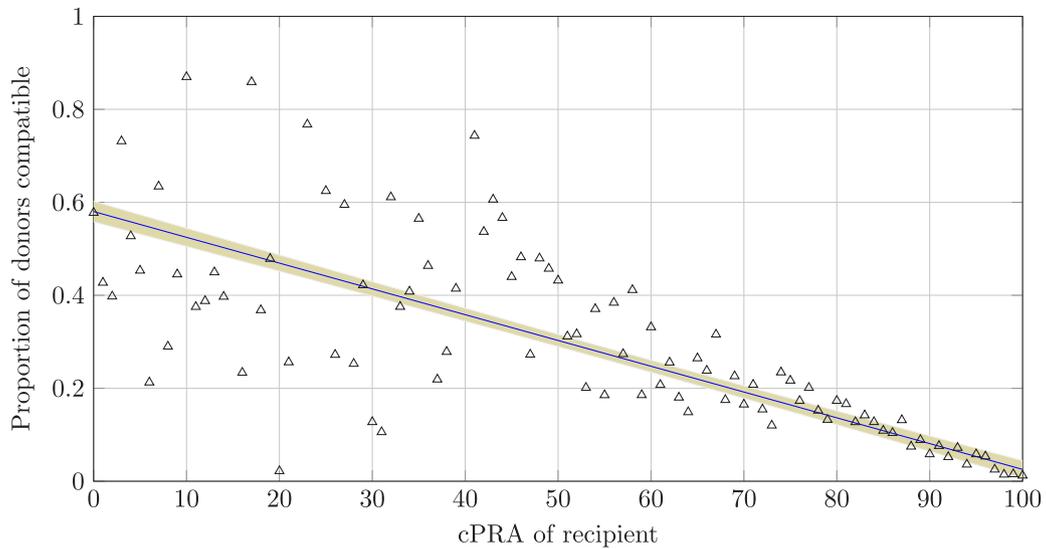
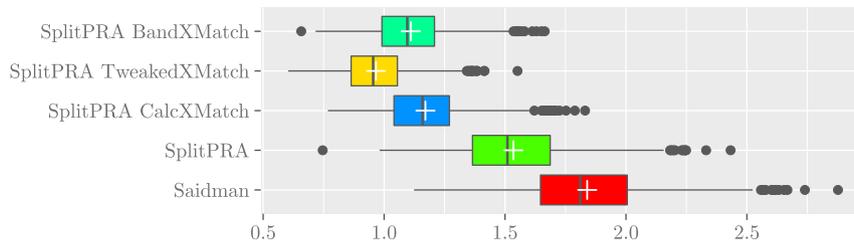
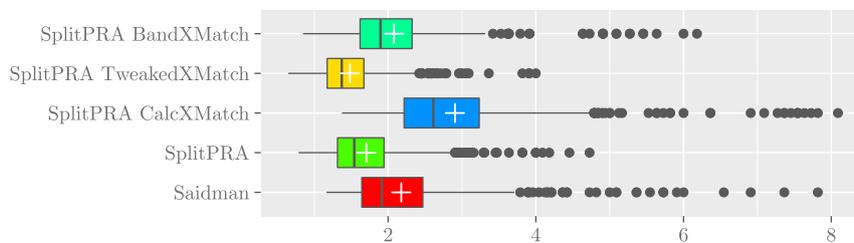


Fig. 7. Proportion of UKLKSS donors that are compatible with a blood-group compatible recipient.



(a) Distribution of number of edges as proportion of edges in real-world instances



(b) Size of largest set of transplants

Fig. 8. Comparison of four methods of calculating compatibility.

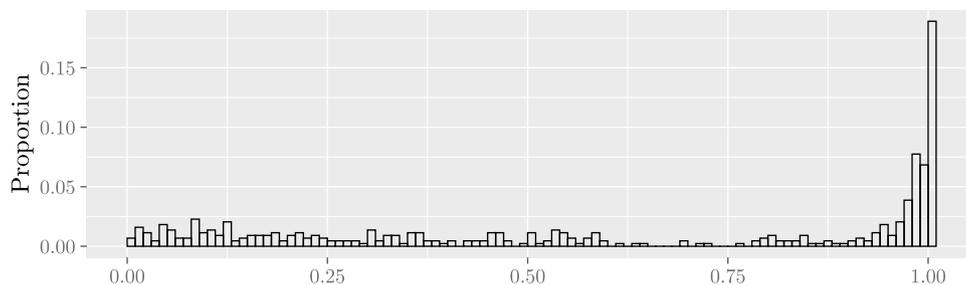
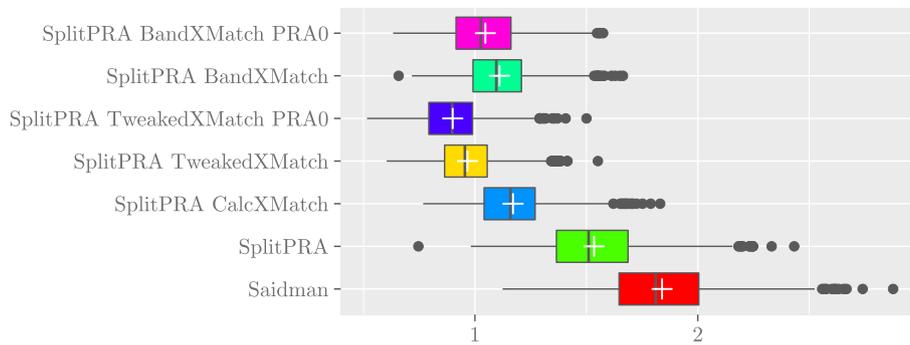


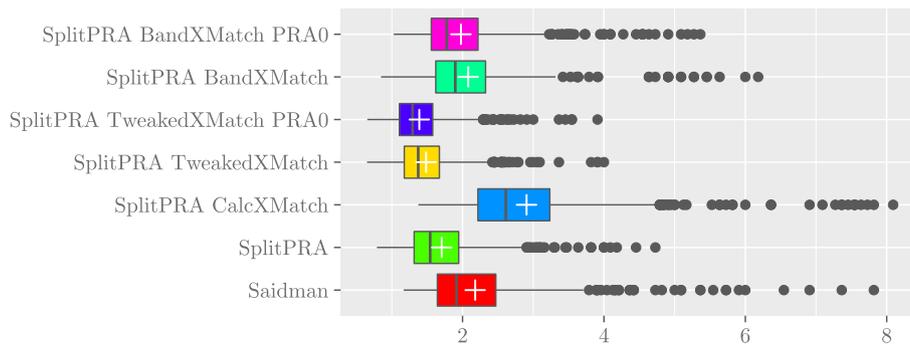
Fig. 9. Histogram of compatibility for cPRA 0 recipients.

show that, within the UKLKSS, recipients with a cPRA of 0 have an unusual distribution of transplant compatibility. Fig. 10 shows that two particular combinations of improvements, SplitPRA TweakedXMatch

PRA0 and SplitPRA BandXMatch PRA0, give the best performance. It is important to note that, while both Fig. 10 and further plots in Appendix C show that SplitPRA TweakedXMatch PRA0 generally produces



(a) Distribution of number of edges as proportion of edges in real-world instances



(b) Size of largest set of transplants

Fig. 10. Comparison of all six new methods of calculating compatibility.

instances with properties closer to instances from the UKLKSS, SplitPRA TweakXMatch PRA0 gives recipients with a cPRA of 100 zero chance of finding a compatible donor, while SplitPRA BandXMatch PRA0 gives each such recipient approximately 1% chance of being compatible with each blood-compatible donor in the pool. With that in mind, to simulate the UKLKSS, we recommend the use of SplitPRA BandXMatch PRA0.

We ran the first step of the matheuristic on instances generated with our new methods, and the bound obtained was generally found to be far above the optimal solution value (e.g., the maximum number of transplants, as determined by the cycle formulation, was 364.1 on average for 20 instances with 1024 recipient/donor pairs while the matheuristic was looking for a solution with value 668.2 on average). This indicates that, in such instances, the maximum number of transplants does not solely depend on the blood groups anymore, which explains why the value was overestimated in the first step of our approach. Note that one could update the matheuristic to find a feasible solution even if it does not match the bound found in the first step by modelling the “pairs to use” and the “cycles to complete” as parts of the objective function instead of within constraints. Even though such an adaptation could potentially provide good-quality solutions, the lack of a guarantee of optimality would be a major drawback, as exact approaches are strongly favoured for optimisation problems where the objective function may impact on human lives.

As our generator incorporates a wide variety of parameters, it would also be useful as one component in a dynamic kidney exchange programme simulation (e.g., Santos et al., 2017). However, we do not have access to, and therefore are unable to report on, either the rate at which identified transplants are performed, or the rate at

which recipients leave before being assigned a transplant. An accurate measure of both of these rates, and a study of any possible correlation with other parameters such as cPRA, would be required to complete such a dynamic kidney exchange programme simulator.

Many of the conclusions we have drawn are based on a study of UK data, but we expect that at least some of our adaptations (in particular, highlighting the correlation between blood compatibility and cPRA distributions, and the relationship between cPRA and transplant compatibility) will be useful when generating random instances for other KEs.

5. Conclusion

We have presented new upper bounds and new matheuristic for identifying large sets of transplants in KE instances. We show that, on a standard set of randomly-generated instances, combining our upper bound and our matheuristic achieves optimal solutions much faster than any algorithm from the literature. This particular set of instances, despite its common use in the literature, differs from real-world instances in a number of vital characteristics (particularly edge density), and our upper bound and matheuristic are not guaranteed to converge, and indeed do not converge, on all real-world instances. We then demonstrated a number of improvements to the standard method of generating random KE instances which are shown to produce instances that are much closer to real-world instances in many characteristics. We have made our generator, and relevant distributions and parameters from the UKLKSS available to allow people to test both

new KEP algorithms and new KEP policies on data that closely matches real-world instances from the UKLKS.

We hope to continue work with this generator by collaborating with other KEPs to determine which adaptations are applicable globally and obtain new adaptations for improved generation of KE instances. In particular, we believe that the insights explored in this paper can be used to improve dynamic modelling of kidney exchange programmes (as seen in Santos et al., 2017). Through careful study, correlations between the extra parameters required for such simulation (e.g., whether a selected transplant succeeds, the rate of departure from the programme, the chance of a positive crossmatch) and the parameters studied in this paper may be discovered, leading to more accurate dynamic modelling of KEPs.

CRedit authorship contribution statement

Maxence Delorme: Conceptualisation, Software, Writing – original draft, Writing – review & editing. **Sergio García:** Conceptualisation, Writing – review & editing, Supervision, Project administration, Funding acquisition. **Jacek Gondzio:** Conceptualisation, Writing – review & editing, Supervision. **Jörg Kalcsics:** Conceptualisation, Writing – review & editing, Supervision. **David Manlove:** Conceptualisation, Writing – review & editing, Supervision, Project administration, Funding acquisition. **William Pettersson:** Conceptualisation, Software, Writing – original draft, Writing – review & editing. **James Trimble:** Conceptualisation, Software.

Acknowledgements

The authors would like to thank the two anonymous reviewers for their valuable comments that have helped to improve the presentation of this paper, and NHS Blood and Transplant for sharing their data and experience with us. This research was supported by the Engineering and Physical Science Research Council, United Kingdom through grant numbers EP/P028306/1 (Manlove and Pettersson), EP/P029825/1 (Delorme, García, Gondzio, and Kalcsics), and EP/R513222/1 (Trimble).

Supplementary data

All appendices related to this article can be found online at <https://doi.org/10.1016/j.cor.2022.105707>.

References

- Abraham, D.J., Blum, A., Sandholm, T., 2007. Clearing algorithms for barter exchange markets: enabling nationwide kidney exchanges. In: Proceedings Of EC '07: The 8th ACM Conference On Electronic Commerce. ACM, pp. 295–304.
- Alvelos, F., Klimentova, X., Viana, A., 2019. Maximizing the expected number of transplants in kidney exchange programs with branch-and-price. *Ann. Oper. Res.* 272 (1–2), 429–444.
- Anderson, R., Ashlagi, I., Gamarnik, D., Roth, A.E., 2015. Finding long chains in kidney exchange using the traveling salesman problem. *Proc. Natl. Acad. Sci.* 112, 663–668.
- Ashlagi, I., Roth, A., 2011. Individual rationality and participation in large scale, multi-hospital kidney exchange. In: Proceedings Of EC '11: The 12th ACM Conference On Electronic Commerce. ACM, pp. 321–322.
- Ashlagi, I., Roth, A., 2012. New challenges in multihospital kidney exchange. *Amer. Econ. Rev.* 102 (3), 354–359.
- Axelrod, D.A., Schnitzler, M.A., Xiao, H., Irish, W., Tuttle-Newhall, E., Chang, S.-H., Kasiske, B.L., Alhamad, T., Lentine, K.L., 2018. An economic assessment of contemporary kidney transplant practice. *Am. J. Transplant.* 18, 1168–1176.
- Biró, P., Gyetvai, M., Klimentova, X., Pedroso, J., Pettersson, W., Viana, A., 2020. Compensation scheme with Shapley value for multi-country kidney exchange programmes. In: Steglich, M., Mueller, C., Neumann, G., Walther, M. (Eds.), Proceedings Of The 34th International ECMS Conference On Modelling And Simulation, ECMS 2020, Wildau, Germany, June 9–12, 2020. European Council for Modeling and Simulation, pp. 129–136.
- Biró, P., Haase-Kromwijk, B., Andersson, T., Åsgerirsson, E.I., Baltessová, T., Boletis, I., Bolotinha, C., Bond, G., Böhmig, G., Burnapp, L., Cechlárová, K., Ciaccio, P. Di, Fronek, J., Hadaya, K., Hemke, A., Jacquelinet, C., Johnson, R., Kieszek, R., Kuypers, D.R., Leishman, R., Macher, M.-A., Manlove, D., Menoudakou, G., Salonen, M., Smeulders, B., Sparacino, V., Spieksma, F.C.R., Valentín, M.O., Wilson, N., van der Klundert, J., 2019a. Building kidney exchange programmes in europe — An overview of exchange practice and activities. *Transplantation* 103, 1514–1522.
- Biró, P., van de Klundert, J., Manlove, D., Pettersson, W., Andersson, T., Burnapp, L., Chromy, P., Delgado, P., Dworczak, P., Haase, B., Hemke, A., Johnson, R., Klimentova, X., Kuypers, D., Costa, A., Smeulders, B., Spieksma, F., Valentín, M., Viana, A., 2019b. Modelling and optimisation in European kidney exchange programmes. *Eur. J. Oper. Res.* 291 (2), 447–456.
- Buse, A., 1982. The likelihood ratio, Wald, and Lagrange multiplier tests: An expository note. *Amer. Statist.* 36 (3a), 153–157.
- Carvalho, M., Klimentova, X., Glorie, K., Viana, A., Constantino, M., 2020. Robust models for the kidney exchange problem. *INFORMS J. Comput.*
- Constantino, M., Klimentova, X., Viana, A., Rais, A., 2013. New insights on integer-programming models for the kidney exchange problem. *Eur. J. Oper. Res.* 231, 57–68.
- Delorme, M., García, S., Gondzio, J., Kalcsics, J., Manlove, D., Pettersson, W., 2020. New algorithms for hierarchical optimisation in kidney exchange programmes. Technical report ERGO 20–005, University of Edinburgh.
- Dickerson, J., Manlove, D., Plaut, B., Sandholm, T., Trimble, J., 2016. Position-indexed formulations for kidney exchange. In: Proceedings Of The 2016 ACM Conference On Economics And Computation. In: EC '16, ACM, New York, NY, USA, pp. 25–42.
- Dickerson, J., Procaccia, A.D., Sandholm, T., 2012. Optimizing kidney exchange with transplant chains: Theory and reality. In: Proceedings Of AAMAS '12: The 11th International Conference On Autonomous Agents And Multiagent Systems. 2, International Foundation for Autonomous Agents and Multiagent Systems, pp. 711–718.
- Edmonds, J., 1965. Paths, trees, and flowers. *Canad. J. Math.* 17, 449–467.
- Glorie, K., 2012. Estimating the Probability of Positive Crossmatch After Negative Virtual Crossmatch. Technical Report EI 2012-25, (EI 2012-25), Erasmus University Rotterdam, pp. 1–8.
- Glorie, K., 2014. Clearing Barter Exchange Markets: Kidney Exchange and Beyond. (Ph.D. thesis). (EPS-2014-329-LIS), Erasmus University.
- Hart, A., Smith, J.M., Skeans, M.A., Gustafson, S.K., Stewart, D.E., Cheriak, W.S., Wainwright, J.L., Kucheryavaya, A., Woodbury, M., Snyder, J.J., Kasiske, B.L., Israni, A.K., 2017. OPTN/SRTR 2015 annual data report: Kidney. *Am. J. Transplant.* 17, 21–116.
- Jordan, S.C., Choi, J., Kahwaji, J., Vo, A., 2016. Progress in desensitization of the highly HLA sensitized patient. *Transplant. Proc.* 48 (3), 802–805.
- Klimentova, X., Pedroso, J., Viana, A., 2016. Maximising expectation of the number of transplants in kidney exchange programmes. *Comput. Oper. Res.* 73, 1–11.
- Kuhn, H.W., 1955. The hungarian method for the assignment problem. *Nav. Res. Logist. Q.* 2 (1–2), 83–97.
- Lam, E., Mak-Hau, V., 2020. Branch-and-cut-and-price for the cardinality-constrained multi-cycle problem in kidney exchange. *Comput. Oper. Res.* 115 (104852).
- Li, Z., Gupta, N., Das, S., Dickerson, J., 2018. Equilibrium behavior in competing dynamic matching markets. In: Proceedings Of The Twenty-Seventh International Joint Conference On Artificial Intelligence, IJCAI-18. International Joint Conferences on Artificial Intelligence Organization, pp. 389–395.
- Manlove, D., O'Malley, G., 2015. Paired and altruistic kidney donation in the UK: Algorithms and experimentation. *ACM J. Exp. Algorithm.* 19.
- Mattei, N., Walsch, T., 2020. Preflib.org: A library for preferences. <https://www.preflib.org/data/matching/kidney/> (accessed 23 September 2020).
- McElfresh, D., Bidkhorji, H., Dickerson, J., 2019. Scalable robust kidney exchange. *Proc. AAAI Conf. Artif. Intell.* 33 (01), 1077–1084, URL <https://ojs.aaai.org/index.php/AAAI/article/view/3899>.
- Mincu, R., Biró, P., Gyetvai, M., Popa, A., Verma, U., 2020. IP solutions for international kidney exchange programmes. *Cent. Eur. J. Oper. Res.* 1–21.
- Mulley, W., Kanellis, J., 2011. Understanding crossmatch testing in organ transplantation: A case-based guide for the general nephrologist. *Nephrology* 16 (2), 125–133.
- Roth, A.E., Sönmez, T., Ünver, M.U., 2004. Kidney exchange. *Q. J. Econ.* 119 (2), 457–488.
- Roth, A.E., Sönmez, T., Ünver, M.U., 2007. Efficient kidney exchange: Coincidence of wants in a market with compatibility-based preferences. *Amer. Econ. Rev.* 97 (3), 828–851.
- Saidman, S.L., Roth, A.E., Sönmez, T., Ünver, M.U., Delmonico, S.L., 2006. Increasing the opportunity of live kidney donation by matching for two and three way exchanges. *Transplantation* 81 (5), 773–782.
- Santos, N., Tubertini, P., Viana, A., Pedroso, J.P., 2017. Kidney exchange simulation and optimization. *J. Oper. Res. Soc.* 68 (12), 1521–1532.
- Wolfe, R.A., Roys, E.C., Merion, R.M., 2010. Trends in organ donation and transplantation in the United States, 1999–2008. *Am. J. Transpl.* 10, 961–972.