Finding long chains in kidney exchange using the traveling salesman problem

Ross Anderson[a], Itai Ashlagi[b], David Gamarnik[b], and Alvin E. Roth[c,1]

[a]Operation Research Center and [b]Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02142; and [c]Department of Economics, Stanford University, Stanford, CA 94305

As of May 2014 there were more than 100,000 patients on the waiting list for a kidney transplant from a deceased donor. Although the preferred treatment is a kidney transplant, every year there are fewer donors than new patients, so the wait for a transplant continues to grow. To address this shortage, kidney paired donation (KPD) programs allow patients with living but biologically incompatible donors to exchange donors through cycles or chains initiated by altruistic (nondirected) donors, thereby increasing the supply of kidneys in the system. In many KPD programs a centralized algorithm determines which exchanges will take place to maximize the total number of transplants performed. This optimization problem has proven challenging both in theory, because it is NP-hard, and in practice, because the algorithms previously used were unable to optimally search over all long chains. We give two new algorithms that use integer programming to optimally solve this problem, one of which is inspired by the techniques used to solve the traveling salesman problem. These algorithms provide the tools needed to find optimal solutions in practice.

kidney exchange | kidney paired donation | transplantation | algorithms | computation

As of May 2014 there were over 100,000 patients on the waiting list for a kidney transplant from a deceased donor. Many of these patients have a friend or family member willing to be a living kidney donor, but the donor is biologically incompatible. Kidney exchange, also called kidney paired donation (KPD), arose to allow these patients with willing donors (referred to as patient–donor pairs) to exchange kidneys, thus increasing the number of living donor transplants and reducing the size of the waiting list.

In KPD, incompatible patient–donor pairs can exchange kidneys in cycles with other such pairs so that every patient receives a kidney from a compatible donor (Fig. S1). Additionally, there are a small number of altruistic donors who are willing to donate their kidney to any patient without asking for anything in return. In KPD, these donors initiate a chain of transplants with incompatible pairs, ending with a patient on the waiting list that has no associated donor (Fig. S2).

Integrating both cycles and chains in KPD was proposed in ref. 1, allowing both chains and cycles to be of unlimited size. To ensure that every patient receives a kidney before her associated donor donates her kidney, cycles are conducted simultaneously (otherwise, if the intended donor is unable to donate to the patient whose associated donor already gave her kidney the pair not only did not receive a kidney, but also could not participate in future exchanges). Because organizing many surgeries simultaneously is logistically very complex, the first implementations of KPD by the New England Program for Kidney Exchange and other clearinghouses used only two-way cyclic exchanges. After a short period, clearinghouses have moved to allow three-way exchanges as well.

In ref. 2 it was proposed to relax the requirement of simultaneity to the weaker requirement that every patient–donor pair receive a kidney before they give a kidney. Although for cycles this restriction still required all surgeries be performed simultaneously, it did allow for nonsimultaneous chains. Note that these nonsimultaneous chains still protected patient–donor pairs from irreparable harm but allowed for the possibility of donors' backing out after their patient had received a transplant. Since the first nonsimultaneous chain was arranged (3) chain-type exchanges have accounted for a majority of the transplants in kidney exchange clearinghouses. [Approximately 75% of the transplants in the National Kidney Registry (NKR) and the Alliance for Paired Donation (APD) are done through chains.] Chains involving as many as 30 pairs have been performed in practice, capturing significant public interest (4). Very long chains are often planned in segments, in which the donor from the final pair in a segment is used to begin another segment after more patients arrive. Such donors are called "bridge donors." Once the segment is executed, the bridge donors and altruistic donors are essentially identical for the purpose of planning future transplants and are collectively referred to here as "nondirected donors" (NDDs). The problem of determining how to optimally select a set of cycles and chains to maximize the number of transplants performed is the focus of this work.

We refer to the problem of finding the maximum (possibly weighted) number of transplants for a pool of incompatible patient–donor pairs and NDDs as the kidney exchange problem (KEP). Optionally included in the problem are a maximum chain length and a maximum cycle length. Weights can be used to prioritize difficult-to-match patients. An example of a KEP instance is shown in Fig. S3. When there is no bound on the chain or cycle length, the problem can solved efficiently by reducing it to a maximum weighted perfect matching problem on a bipartite graph, as described in ref. 5. The special case in which only cycles

of length two are used can be exactly solved very efficiently, because it is trivially equivalent to the maximum matching problem. In ref. 5, the case of this problem where only cycles of length two and three are used was shown to be NP-hard, meaning that it is unlikely there will be an algorithm that always finds the optimal solution quickly in every instance of the problem; see ref. 6 for a stronger negative result in this special case. However, integer programming (IP) techniques have been used by a variety of authors to solve special cases of the KEP without chains or with chains of bounded length, as first proposed in ref. 7. In ref. 5, by improving the IP formulation of ref. 7 and devising good heuristics the authors were able to solve KEP instances with thousands of donor pairs, but without chains. Alternate IP formulations were further explored for this special case in ref. 8. In ref. 1, a heuristic to produce feasible solutions when using chains and cycles of bounded length was suggested. However, no optimization algorithm was given, opening a major algorithmic challenge to reliably solve large-scale instances of the general KEP. The technique of ref. 5 was extended in refs. 9–12 to solve large instances with bounded cycles and bounded chains. However, the algorithm became impractical when the maximum chain length was beyond four or five, because the formulation required a decision variable for each chain up to the maximum chain length. As seen in Table S1, the number of chains of length up to k grows very quickly with k, at least when creating instances using patient–donor pairs drawn at random from the historical dataset (a more thorough explanation is given in *Supporting Information*). When both chains and cycles are bounded by length k, the relaxed version of the IP formulation (as a linear program) provides a good approximation to the original formulation (13, 14), and a large body of literature suggests that such problems can be effectively solvable in practice. However, this approximation property does not seem to hold if chains are unbounded (15), and hence this problem is likely to be much more challenging than the bounded case, as confirmed by our computational experiments.

The primary challenge addressed in this paper is to find an algorithm to solve real instances of the KEP without bounding the maximum chain length. Solving this optimization problem is critical to the operations of KPD programs, which form long chains in practice but previously relied on heuristics that could lead to suboptimal solutions. We emphasize that we focus on real instances drawn from historical data, as opposed to synthetic instances, because the statistics of instances encountered in practice are quite complicated, making it difficult to generate representative random instances (a patient–donor pair waiting to be matched is on average more difficult to match than a patient–donor pair randomly selected from historical data, because the easy-to-match patients are on average matched more quickly). The second challenge we address in this paper is finding an algorithm that can solve instances using chains and cycles with an arbitrary bound on the maximum cycle length (with no maximum cycle length, the problem can be solved trivially). Although in practice cycles longer than three are rarely formed owing to logistical issues, there is some evidence that there may be benefits to considering longer cycles (16). Finally, in *Supporting Information* we demonstrate how our approach can be extended to efficiently deal with the case of bounded chain lengths as well as a stochastic version with "edge failures" (see refs. 17 and 18).

We propose two new algorithms for the KEP that address these challenges. The first algorithm is a recursion based on a direct IP formulation for the KEP. Despite the NP hardness of the underlying problem, the algorithm is surprisingly effective in practice in most instances, as we will show in a later section devoted to the computational results. Additionally, the algorithm does not break down as the maximum cycle length is increased. On the contrary, as the bound on the maximum cycle length increases, the performance improves. At the same time there are

several KEP instances of moderate size that have been encountered in practice on real data drawn from a particular KPD program that our algorithm was unable to solve. Thus, we have devised a second algorithm to reliably solve nearly all KEP instances of practical interest. The algorithm is motivated by an IP formulation of the so-called prize-collecting traveling salesman problem (PC-TSP), another NP-hard problem. The PC-TSP is a variant of the classical traveling salesman problem (TSP), one of the most widely studied NP-hard problems in combinatorial optimization. Although our PC-TSP–based algorithm was able to solve every instance we have encountered, somewhat surprisingly we have discovered that it is at times orders of magnitude slower than the direct recursive algorithm described above.

The two algorithms described above allow us to optimally allocate kidneys transplanted in KPD with arbitrary chain lengths. These algorithms were capable of solving to optimality all real data instances we have encountered, using only a single desktop computer. Furthermore, several of the most active KPD programs are using our implementations of these algorithms in conducting their kidney exchanges, including Methodist Specialty and Transplant Hospital in San Antonio, Texas and the Northwestern University Medical Center in Chicago.

The remainder of the paper is organized as follows. First, we formally define the KEP as an optimization problem on a graph. Then, we describe the recursive algorithm and PC-TSP–based algorithm for the KEP. Both algorithms are based on formulating associated IP problems. Next, we demonstrate the value of being able to find long chains by measuring the number of additional transplants they enable using simulations based on historical clinical data. We then compare the performance of our two algorithms in KEP instances found in actual data from the KPD programs, demonstrating the relative strength of the PC-TSP–based algorithm compared with the recursion-based algorithm. To provide some insight into the power of the PC-TSP–based algorithm we state a result showing that the formulation of the IP problem associated with the PC-TSP is stronger than the IP problem corresponding to the recursive algorithm. The proof of this result can be found in *Supporting Information*. Furthermore, we provide an example of a pathological instance of the KEP for which the recursive algorithm takes a very long time, whereas the PC-TSP–based algorithm solves this instance very quickly. Finally, we conclude with a summary of our results. Some technical results can be found in *Supporting Information*.

KEP

An instance of the KEP is described as follows:

- a list of NDDs,
- a list of patient–donor pairs (in which the donor wants to donate to the paired patient but is not compatible with this patient),
- the compatibility information between all donors and patients,
- the "weight," or priority, of each potential transplant, and
- a bound on the maximum cycle length.

The goal is then to find a set of transplants, organized into cycles and chains initiated by NDDs that uses each donor and patient at most once and maximizes the sum of the weights of all transplants performed. If all transplants have weight one, then we are simply trying to find the arrangement that maximizes the total number of transplants. In *Supporting Information* we will show how this definition can be supplemented to include an optional bound on the maximum chain length.

We now formalize this definition of the KEP in graph theoretic terms. We are given a directed graph $G = (V, E)$, a weight vector $\mathbf{w} = (w_1, \ldots, w_{|E|})$, and a nonnegative integer parameter k. The set V is partitioned into sets N (the NDDs) and P (the pairs of incompatible donors and patients). For $u, v \in V$; a directed

edge from u to v in E indicates that the donor in node u is compatible with the patient in node v. Because the nodes of N have no patient, they all must have in-degree zero (although there can be nodes in P with in-degree zero as well). The values w_e for each edge $e \in E$ are weights for the edges, indicating the importance of this transplant, and our goal is to find a maximum weight node disjoint cycle and chain packing, where the cycles can use at most k nodes and the chains must originate from nodes in N.

We introduce some notation. For each $v \in V$, let $\delta^-(v)$ be the edges pointing to v and $\delta^+(v)$ be the edges outgoing from v. Likewise, for a set of nodes $S \subset V$, let $\delta^-(S)$ be the set of edges into S and $\delta^+(S)$ be the set of edges out of S. For every $S \subset V$, let $E(S)$ be the set of all edges with both endpoints in S. For a set of edges $D \subset E$, let $V(D)$ be the set of vertices that are of the endpoints of edges in D. Let $\mathcal{C}$ be the set of all simple cycles in G, where each cycle $C \in \mathcal{C}$ is represented by a collection of edges, that is, $C \subset E$. Let $\mathcal{C}_k$ be the subset of $\mathcal{C}$ consisting of cycles that use k or fewer edges. For each $v \in V$, let $\mathcal{C}_k(v)$ be the cycles from $\mathcal{C}_k$ containing an edge incident to v. Given a cycle C, let $w_C = \sum_{e \in C} w_e$ be the total weight of the cycle according to our weight vector $\mathbf{w}$.

Recursive Algorithm for the KEP

We now describe the first of our two algorithms. The algorithm uses a straightforward IP formulation of the KEP, with a binary variable for every edge, and constraints so that each node is used at most once and no long cycles occur.

We first formulate the problem so that the objective is to maximize a weighted number of matches subject to the constraint that no cycle in the graph of length more than k is chosen. Because the number of such constraints is too large, the algorithm will only add these constraints when necessary. In particular, the algorithm first solves the relaxed problem where the bound on the cycle length is dropped, and then if it finds a solution with cycles longer than permitted it adds the corresponding constraints to eliminate these cycles and resolves the corresponding optimization problem.

Formally, we use decision variables y_e for $e \in E$ and f_v^i (flow in) and f_v^o (flow out) for $v \in V$, and solve

$$\max \quad \sum_{e \in E} w_e y_e$$

$$\text{s.t.} \quad \sum_{e \in \delta^-(v)} y_e = f_v^i \quad v \in V \qquad [1]$$

$$\sum_{e \in \delta^+(v)} y_e = f_v^o \quad v \in V$$

$$f_v^o \le f_v^i \le 1 \quad v \in P, \qquad [2]$$

$$f_v^o \le 1 \quad v \in N, \qquad [3]$$

$$\sum_{e \in C} y_e \le |C| - 1 \quad C \in \mathcal{C} \backslash \mathcal{C}_k, \qquad [4]$$

$$y_e \in \{0, 1\} \quad e \in E. \qquad [5]$$

Note that we introduce some auxiliary variables f_v^i and f_v^o for all $v \in V$ to simplify the formulation, although because they are defined by the equality constraints they can be eliminated. In words, **2** says that for the patient–donor pair nodes the flow out is at most the flow in, and the flow in is at most one; **3** says that for the NDD nodes the flow out is at most one; and in **4** the "cycle inequalities" say that for any cycle C of length greater than k the number of edges we can use is at most $|C| - 1$, thus prohibiting long cycles when $\mathbf{y}$ is integral.

The number of constraints in the IP above is exponential in $|E|$, owing to **4**. As a result, for large instances we cannot simply enumerate all of these constraints and give them directly to the IP solver. Instead, we use a simple recursive algorithm to solve the problem. First, we eliminate all of the constraints in **4** and solve the integer program to optimality. Then, we check whether the proposed solution contains any cycles of length greater than k. If so, we add the violated constraint from **4** and resolve. We repeat this procedure until our solution contains no cycles longer than k. This methodology is generally referred to using "lazy constraints." The effectiveness of this technique depends on our not having to generate too many of the constraints from **4** on a typical input. In the worst case, we might have to solve exponentially many IPs, and to solve each IP may take an exponential amount of time. However, as we will show, this technique is often quite effective in practice. Finally, note that the efficiency of this procedure relies on the fact that we can very quickly detect whether any of the constraints from **4** are violated for an integer solution, because we can (trivially) find the largest cycle in a degree-two graph in linear time.

PC-TSP–Based Algorithm for the KEP

We now describe our second algorithm for solving the KEP, inspired by a solution to the PC-TSP. Recall that in the TSP one is given a list of cities and the cost of going between pairs of cities, and the goal is to find a cycle visiting each city exactly once at the minimum cost (19). In the PC-TSP, again one must find a cycle visiting each city at most once, but now one has the option of skipping some cities entirely and paying a penalty (see ref. 20 or, more recently, ref. 21). Qualitatively, the PC-TSP problem is similar to the KEP in that one wants to find long paths in a graph (which the PC-TSP then closes off as a cycle) without the need to visit every node. As we will see below, our solution is similar to the solution for the PC-TSP.

For each cycle C of length at most k, we introduce a new variable z_C that indicates whether we are using the cycle C. We make the natural updates to **1** so the objective value does not change when the same edges are used and to **2** so that edges cannot be used both in a z_C variable and a y_e variable. Finally, we add **7** to block cycles longer than length k, similarly to **4**.

$$\max \quad \sum_{e \in E} w_e y_e + \sum_{C \in \mathcal{C}_k} w_C z_C$$

$$\text{s.t.} \quad \sum_{e \in \delta^-(v)} y_e = f_v^i \quad v \in V$$

$$\sum_{e \in \delta^+(v)} y_e = f_v^o \quad v \in V \qquad [6]$$

$$f_v^o + \sum_{C \in \mathcal{C}_k(v)} z_C \le f_v^i + \sum_{C \in \mathcal{C}_k(v)} z_C \le 1 \quad v \in P,$$

$$f_v^o \le 1 \quad v \in N,$$

$$\sum_{e \in \delta^-(S)} y_e \ge f_v^i \quad S \subseteq P, \ v \in S \qquad [7]$$

$$y_e \in \{0, 1\} \quad e \in E,$$

$$z_C \in \{0, 1\} \quad C \in \mathcal{C}_k.$$

The constraint **7** is very similar to the cut set inequalities for the TSP (19) as adapted to the PC-TSP by several authors (see ref. 21 and references therein). Fig. S4 provides a clarifying example explaining these constraints. Essentially, they work as follows. Suppose that a chain is reaching some node v, and as a result f_v^i equals one. Now suppose that we cut the graph into two pieces such that the half containing v does not contain any of the NDD nodes from N. Because every chain begins at some node in N (and thus does not begin in S), in order for our chain to reach

$v \in S$ it must use an edge that begins not in S and ends in S, that is, an edge $e \in \delta^-(S)$. Thus, our constraint requires that whenever there is flow into v, for every way that v can be cut off from the NDDs N there is at least this much flow over the cut.

Observe that the use of cycles is important for this formulation to be feasible, which makes it quite novel. For example, consider a graph that contains just a set of short disjoint cycles, which is therefore also the optimal solution. Using only constraint **7** with the set $S = V$ and any $v \in V$ implies that the flow going into each of the nodes has to be zero.

Again, the IP formulation has exponentially many constraints from **7**, so we cannot enumerate them and feed them directly to the IP solver. We could simply use the same recursive heuristic (lazy constraints) from the previous section to obtain a correct algorithm. Our solution still relaxes the constraints **7**, but instead we more aggressively attempt to find violated constraints and add them sooner, using a technique called "cutting planes" (or "user cuts"). To apply this method, we needed an efficient algorithm that given a potentially fractional solution can either find a violated constraint from **7** or determine that none exists. This problem is known in the field of optimization as the separation problem.

Theorem 1. *The separation problem for* **7** *can be solved by solving* $O(|P|)$ *network flow problems.*

A proof is given in *Supporting Information*. See ref. 19 for more on the separation problem and on the network flow problem. The solution is again very similar to the solution used to solve separation problems for the TSP and PC-TSP.

Measuring the Value of Long Chains with Clinical Data

To demonstrate the need for long chains in practice, we give an analysis of the number of additional transplants long chains enable, and the types of patients that receive these additional transplants. A patient's so-called percentage reactive antibody (PRA) measures the percentage of the donor pool that a patient has a tissue type incompatibility with (thus patients with a PRA near 100 are very difficult to match, whereas patients with a PRA near zero are easier to match). Patients with a PRA above 80 are commonly referred to as highly sensitized patients. In ref. 16 it was shown for a random model of kidney exchange that when there are many highly sensitized patients allowing for long chains will significantly increase the total number of transplants performed, and that the extra transplants will largely benefit the population of highly sensitized patients. In the data obtained from the APD and NKR exchange programs, we found that in the current exchange pools more than 35% of the patients are highly sensitized, so ref. 16 suggests that we should get some benefit out of allowing long chains.

To assess the added benefit of long chains, we used the following two-stage experiment. First, we repeatedly simulated the dynamic environment of a kidney exchange pool using clinical data (we use 650 patient–donor pairs and 33 altruistic donors from a period of 2.5 y and to create a large set of realistic snapshots of the exchange pool. We measured the effect of changing the maximum chain length on (*i*) the total number of transplants performed and (*ii*) the number of transplants performed on highly sensitized patients.

Note that this experiment is designed to demonstrate that long chains will often provide immediate opportunities to perform more transplants. This experiment, however, does not assess the long-run implications of using long chains or the effect of using long chains on the dynamics of the exchange pool (because in practice matched pairs leave the pool, whereas matched pairs may still be present in subsequent matchings during the simulation). Although demonstrating the latter would make a stronger case for the value of long chains, doing so requires dealing with several subtleties that are beyond the scope of this work. In particular, to understand long-run pool performance it is insufficient to simply look at the total number of transplants

performed over the entire time horizon, because given enough time in a model with no exit of patients and donors eventually everyone will be matched. To get the full picture, the time patients spend waiting to be matched should also be considered. However, to accurately model this requires more complicated simulations of the pool dynamics than are used here that account for real-world features such as exit, edge failures, and the time to schedule transplants. Thus, we choose not to demonstrate the value of long chains by computing the long-term gains from their use, because these calculations are sufficiently complicated to be a separate problem. We instead attempt the more modest goal of demonstrating that in the types of pools that would be generated by the natural kidney exchange dynamics using longer chains enables more immediate transplants.

Finally, note that to create snapshots that are representative of actual instances encountered by a KPD program it is insufficient to simply take altruistic donors and patient–donor pairs at random from a historical dataset, because the patients that are left over after each match run statistically tend to be harder to match than a randomly selected patient (the easy-to-match patients are more likely to be matched immediately).

We now describe our two-stage experiment design in full detail:

i) Generate a set of realistic snapshots of the exchange pool by performing 50 trials of the procedure below and saving all snapshots generated (every trial involves more than 100 snapshots):

- Assign each patient–donor pair and altruistic donor an arrival date by shuffling their historical arrival dates (i.e., select a permutation uniformly at random).
- Wait for the arrival of eight new participants to the exchange pool (either patient–donor pairs or altruistic donors) and record the current pool of those not yet matched as a snapshot.
- Solve the KEP for the current snapshot with the objective of maximizing the total number of transplants, a maximum cycle length of three, and a maximum chain length of five.
- Remove all matched patient–donor pairs from the current pool. The unmatched pairs stay in the pool. Now, if there are more arrivals remaining, we return to the second step (where we wait for new patient donor pairs), and we repeat in this manner until all pairs have arrived.

ii) Produce the following measures of improvement from allowing unbounded maximum chain length in comparison with a bounded chain length. In particular, for every bounded chain length j for $j \in \{3, 4, 5, 6\}$ we measure the following.

- Chance of more highly sensitized patients receiving a transplant: For each snapshot we count the increase percentage of highly sensitized patients (with PRA above 95) matched by allowing unbounded chains vs. bounded chain length of j.
- For each snapshot we measure the increase percentage of all matched patients by allowing unbounded chains vs. bounded chain length of j.
- We measure the number of instances in which there was at least one highly sensitized patient matched and measure the percentage of instances in which more highly sensitized patients were matched using unbounded chains vs. bounded chains of length j.

Note that in our method of generating snapshots bounding the maximum chain of length of five is somewhat arbitrary. Although this method is not ideal, because the method used to generate snapshots affects the statistics of subsequent snapshots, studying the dynamic optimization problem is beyond the scope of this paper.

The results of our experiment are summarized in Table S2. Allowing longer chains allows matching more of the most highly sensitized patients (with PRA at least 95) and also increases the

total number of matched patients. Note also that the percentage of instances in which more highly sensitized patients are matched using unbounded chains in comparison with bounded chains of length $j = 3, 4, 5, 6$ is nonnegligible (between 35% and 23%).

Algorithm Performance in Practice

We briefly summarize computational results for both the recursive formulation and the PC-TSP formulation described above. Table S3 provides the running time of both formulations on a series of "difficult" but realistic KEP instances encountered in practice.

To create these difficult but realistic instances we adapted the method from the previous section for generating realistic snapshots. However, when generating instances for this section we rotated through many combinations of the following three simulation parameters (all of these were fixed when generating snapshots previously):

i) We considered several objective functions giving varying degrees of priority for hard-to-match patients.
ii) We considered a maximum chain length of 3, 4, 5, 6, and unbounded.
iii) We varied the number of arrivals between reoptimizations from 1 to 64.

Finally, because we are interested in performance in difficult instances for Table S3 we only retained the instances where at least one algorithm had to generate at least 500 constraints of the type **4** or **7** in the recursive and PC-TSP algorithms, respectively. Thus, these instances represent more of a worst case than an average case for real data.

Table S3 contains the running time of both algorithms on these difficult instances, with a maximum attempted solve time of 20 min. In particular, in all instances with the reported running time less than 20 min the optimal solution was found. Looking at the table, we make the following observations in comparing the performance of the two algorithms:

- Both algorithms are able to solve most instances to optimality quickly.
- The PC-TSP formulation solves all instances to optimality.
- The PC-TSP algorithm is usually faster, although for the easier of these difficult instances the difference is sometimes negligible.
- On several inputs, the PC-TSP algorithm is orders of magnitude faster.

We stress again that these instances are the worst-case instances, in that we only showed results for problems where at least one algorithm had to generate a large number of constraints. These worst-case inputs are only a small fraction of all of the simulated inputs and, generally speaking, both algorithms can solve most of these instances to optimality very quickly.

To demonstrate that our algorithms can solve instances even larger than those occurring in implementing algorithms online we also ran our algorithms on the entire historical datasets for the KPD programs NKR and APD. Each dataset contains around 1,000 patients (although arriving over the span of several years), making these instances much larger than the instances described in Table S3. The running time for our algorithms on these instances is shown in Table S4. We see that *i*) both algorithms can solve both instances and *ii*) the recursive algorithm is much faster.

Although the second point may seem surprising, we do have some explanation as to why this is taking place. First, and most importantly, these instances are substantially different from the instances that KPDs encounter in practice, in that they do not contain a disproportionately high fraction of hard-to-match patients. As a result, there is a very large number of two and three cycles, making the number of variables in the PC-TSP formulation very large. Second, these instances are both "easy" instances, in that very few of the constraints **4** and **7** must be added by the algorithm to solve the integer programs, unlike the instances in Table S3. As we will see in the next section, the advantages of using the PC-TSP algorithm over the recursive algorithm depend on the constraints **4** and **7** being binding in the optimization problem.

For the purposes of comparing algorithms it would be preferable to have more realistic large-scale instances beyond the two described above, but the current historical data do not produce such large-scale instances. In an attempt to produce more realistic large-scale instances from the historical dataset we experimented with removing fractions of the altruistic donors at random. We found that these did not significantly change the performance of either algorithm.

Understanding the Power of the PC-TSP Formulation

We now describe a theoretical result that explains why the PC-TSP algorithm has much better performance than the recursive algorithm on the "difficult" instances from Table S3. To do so, we must first introduce some standard concepts from optimization theory.

Suppose two formulations of an integer program are given. Without the loss of generality, assume the underlying problem is of the maximization type. The linear programming relaxation of an integer program is the value of the optimal solution obtained when the integrality constraints are removed and the problem is solved as a linear programming problem. Let Z_1 and Z_2 be the optimal solutions to the linear programming relaxations for the given two formulations.

Definition 1: *If* $Z_1 \leq Z_2$ *for every problem instance, then formulation one is defined to be at least as strong as formulation two. We use the notation* $Z_1 \preceq Z_2$. *If in addition there exists a problem instance such that* $Z_1 < Z_2$, *then we say that formulation one is stronger than formulation two and use the notation* $Z_1 \prec Z_2$.

Very often in practice the stronger formulations greatly reduce the actual running time of the IP problems (19).

We compare the strength of our formulations presented using *Definition 1*. Let the optimal values of the linear programming relaxations of the IP problems corresponding to the recursive algorithm and the PC-TSP–based algorithms be given by Z_{rec} and Z_{tsp}, respectively.

Theorem 2. *The PC-TSP formulation is stronger than the recursive formulation, that is,* $Z_{\text{tsp}} \prec Z_{\text{rec}}$.

The proof can be found in *Supporting Information*.

The idea behind this theorem is demonstrated by a somewhat pathological instance of the KEP shown in Fig. S5. This instance has no altruistic donors and a maximum cycle length of three, and all edges have weight one. Compared with the real instances we considered, this is a small instance with only 30 nodes and 120 edges. The 30 nodes are arranged in a single large directed cycle, with a few additional edges. This graph has two important properties: There are no cycles of length three or less, but it has a very large number of cycles of length 30. Because there are no NDDs or short cycles, the optimal solution is zero. It is easy to see that the linear programming relaxation for the PC-TSP formulation is also zero. Thus, for this instance, solving the IP problem is no harder than solving a single linear programming problem with the PC-TSP formulation. Further, to solve the LP, we only need to generate a single cut from the exponential family of **7** (any cut where $S = P$ will do). As a result, we are able to solve this instance almost instantly with the PC-TSP algorithm. In contrast, for the IP problem corresponding to the recursive algorithm, the linear programming relaxation has the optimal value 29, because it can simply assign value 29/30 to every edge on the length-30 cycle. Worse yet, after 2 h of the running time

the best upper bound that the recursive formulation gives is still 30, because the high redundancy in the structure of the graph results in many possible cycles of length 30. The constraints for each of these cycles need to be added to obtain an upper bound of 29. For similar reasons, branch and bound is very ineffective. In the proof, we use two formulations, subtour elimination and a cycles formulation, which are shown to be weaker than the recursive formulation but stronger than the PC-TSP–based formulation (see also Figs. S6–S8).

Conclusion

We provide two new algorithms that solve realistic instances of the KEP encountered in practice. The first algorithm creates an integer program with a variable for each possible transplant and constraints ensuring that each donor and patient are in at most one transplant and then recursively resolves this IP and adds constraints to prevent the long cycles appearing in solutions. The second algorithm uses an IP based on formulation similar to the PC-TSP. To generate realistic instances of the KEP, we simulated dynamics of the patient pool in a kidney exchange using clinical data from the two major KPD Programs: NKR and the APD. Our algorithms have been able to solve all real data instances encountered in practice thus far. Our algorithms find optimal solutions for large-scale practical instances of the KEP with long chains after a very short running time, with the PC-TSP–based algorithm outperforming the recursion-based algorithm in many instances.

1. Roth AE, Sönmez T, Unver MU (2004) Kidney exchange. *Q J Econ* 119(2):457–488.
2. Roth AE, Sönmez T, Unver MU, Delmonico FL, Saidman SL (2006) Utilizing list exchange and nondirected donation through 'chain' paired kidney donations. *Am J Transplant* 6(11):2694–2705.
3. Rees MA, et al. (2009) A non-simultaneous extended altruistic donor chain. *New Engl J Med* 360(11):1096–1101.
4. Sack K (February 18, 2012) 60 lives, 30 kidneys, all linked. *NY Times*. Available at www.nytimes.com/2012/02/19/health/lives-forever-linked-through-kidney-transplant-chain-124.html. Accessed December 23, 2014.
5. Abraham DJ, Blum A, Sandholm T (2007) Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. *Proceedings of the 8th ACM Conference on Electronic Commerce* (Assoc Computing Machinery, New York), pp 295–304.
6. Biró P, Manlove DF, Rizzi R (2009) Maximum weight cycle packing in directed graphs, with application to kidney exchange programs. *Discrete Mathematics Algorithms Applications* 1(04):499–517.
7. Roth AE, Sönmez T, Ünver MU (2007) Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *Am Econ Rev* 97(3):828–851.
8. 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(1):57–68.
9. Dickerson JP, Procaccia AD, Sandholm T (2012) Optimizing kidney exchange with transplant chains: Theory and reality. *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems* (Assoc Computing Machinery, New York), Vol 2, pp 711–718.
10. Manlove DF, O'Malley G (2012) Paired and altruistic kidney donation in the UK: Algorithms and experimentation. *Experimental Algorithms*, Lecture Notes in Computer Science (Springer, Berlin), Vol 7276, pp 271–282.
11. Glorie KM, Wagelmans APM, Van de Klundert J (2012) Iterative branch-and-price for hierarchical multi-criteria kidney exchange. Econometric Institute Research Paper EI 2012-11 (Econometric Inst, Erasmus Univ, Rotterdam).
12. Klimentova X, Alvelos F, Viana A (2014) A new branch-and-price approach for the kidney exchange problem. *Computational Science and Its Applications – ICCSA 2014*, Lecture Notes in Computer Science (Springer, Berlin), Vol 8580, pp 237–252.
13. Füredi Z, Kahn J, Seymour PD (1993) On the fractional matching polytope of a hypergraph. *Combinatorica* 13(2):167–180.
14. Hazan E, Safra S, Schwartz O (2006) On the complexity of approximating k-set packing. *Computational Complexity* 15(1):20–39.
15. Björklund A, Husfeldt T, Khanna S (2004) Approximating longest directed paths and cycles. *Automata, Languages and Programming*, Lecture Notes in Computer Science (Springer, Berlin), Vol 3142, pp 222–233.
16. Ashlagi I, Gamarnik D, Rees MA, Roth AE (2012) The need for (long) chains in kidney exchange. (National Bureau of Economic Res, Cambridge, MA).
17. Dickerson JP, Procaccia AD, Sandholm T (2013) Failure-aware kidney exchange. *Proceedings of the Fourteenth ACM Conference on Electronic Commerce* (Assoc Computing Machinery, New York), pp 323–340.
18. Chen Y, et al. (2012) Graph-based optimization algorithm and software on kidney exchanges. *IEEE Trans Biomed Eng* 59(7):1985–1991.
19. Bertsimas D, Weismantel R (2005) *Optimization over Integers* (Dynamic Ideas, Belmont, MA), Vol 13.
20. Bienstock D, Goemans MX, Simchi-Levi D, Williamson D (1993) A note on the prize collecting traveling salesman problem. *Math Program* 59(1):413–420.
21. Goemans MX (2009) Combining approximation algorithms for the prize-collecting TSP. arXiv:0910.0553.