Interfaces

Kidney Exchange and the Alliance for Paired Donation: Operations Research Changes the Way Kidneys Are Transplanted

Ross Anderson, Itai Ashlagi, David Gamarnik, Michael Rees, Alvin E. Roth, Tayfun Sönmez, M. Utku Ünver

THE FRANZ EDELMAN AWARD
Achievement in Operations Research

Kidney Exchange and the Alliance for Paired Donation: Operations Research Changes the Way Kidneys Are Transplanted

Ross Anderson, Itai Ashlagi, David Gamarnik
Massachusetts Institute of Technology, Cambridge, Massachusetts 02142
{rma350@gmail.com, iashlagi@mit.edu, gamarnik@mit.edu}

Michael Rees
The Alliance for Paired Donation, Maumee, Ohio 43537, michael.rees2@utoledo.edu

Alvin E. Roth
Stanford University, Stanford, California 94305, alroth@stanford.edu

Tayfun Sönmez, M. Utku Ünver
Boston College, Chestnut Hill, Massachusetts 02467 {sonmezt@bc.edu, unver@bc.edu}

Many end-stage renal disease sufferers who require a kidney transplant to prolong their lives have a relative or friend who has volunteered to donate a kidney to them, but whose kidney is incompatible with the intended recipient. This incompatibility can sometimes be overcome by exchanging a kidney with another incompatible patient-donor pair. Such kidney exchanges have emerged as a standard mode of kidney transplantation in the United States. The Alliance for Paired Donation (APD) developed and implemented nonsimultaneous extended altruistic donor (NEAD) chains, an innovative technique that allows a previously binding constraint (of simultaneity) to be relaxed; thus, it permits longer chains and better-optimized matching of potential donors to patients, greatly increasing the number of possible transplants. Since 2006, the APD has saved more than 220 lives through its kidney exchange program, with more than 75 percent of these achieved through nonsimultaneous chains. Other kidney exchange programs have adopted the technology and methods pioneered by APD, resulting in more than 1,000 lives already saved, with the promise of increasing impact in coming years. In 2013, the percentage of transplants from nonsimultaneous chains reached more than six percent of the number of transplants from living donors. In this paper, we describe the long-term optimization and market design research that supports this innovation. We also describe how a team of physicians and operations researchers worked to overcome the skepticism and resistance of the medical community to the NEAD innovation.

Keywords: matching: transplantation; market design; optimization; design; integer programming.

This paper describes how the Alliance for Paired Donation (APD) used operations research (OR) and market design methodologies to make improved matches between kidney disease sufferers and potential kidney donors. To date, the APD, a nonprofit organization, has saved more than 220 lives through its development and implementation of nonsimultaneous extended altruistic donor (NEAD) chains that we describe next. In addition, other kidney transplant institutions and kidney exchange programs have adopted the APD methodology, saving more than 1,000 lives to date. In the future, we expect that thousands of lives will be saved annually. The percentage of transplants from nonsimultaneous chains reached more than six percent of the number of transplants from living donors in 2013. We estimate that as a result of using these chains, the number of transplants conducted through kidney exchange programs has increased by more than 10 percent and the percentage of very highly sensitized patients transplanted

(i.e., those for whom it is most difficult to find a compatible kidney) has increased by more than 50 percent.

This project involved a blend of OR techniques, including optimization algorithms for finding a large number of matches between donors and recipients and market design mechanisms to facilitate collaboration across organizational boundaries by hospitals, institutions, and surgeons who otherwise compete in the $50 billion end-stage renal disease (ESRD) industry.

Background

In the United States, about 100,000 sufferers of ESRD are currently on the waiting list for a kidney transplant from a deceased donor. Transplantation is the preferred treatment for this severe disease. Sadly, about 4,000 of the patients on the waiting list will die this year before receiving a transplant, and another 2,500 will be removed from the queue as "too sick to transplant." In addition, about 5,000 patients will receive transplants from living donors and 10,000 will receive transplants from deceased donors. The problem is growing worse, as the waiting list increases by about 7,000 patients per year, making ESRD one of the most expensive problems in the U.S. healthcare system.

These long waiting times for transplants have two dominant causes. The first is the shortage of kidneys available for transplant—kidneys that are available from either deceased donors or living persons who are willing to donate one of their two healthy kidneys. The second cause is potential incompatibility between living donors and their intended recipients. The thrust of kidney exchange programs is to significantly increase the number of living-donor kidney transplants by bringing together incompatible donors and recipients and conducting exchanges so that each recipient receives a compatible kidney. These programs essentially develop a marketplace for kidney exchanges, taking into account that a monetary market is forbidden by law. Specifically, kidney exchange allows a potential living donor whose kidney is incompatible with the intended recipient to donate a kidney to another patient so that the donor's recipient receives a compatible kidney from another donor. This is done by exchange between two or more incompatible patient-donor pairs; each donor gives a compatible kidney to another donor's intended recipient (see Figure 1(a)). Rapaport (1986) first suggested the idea of
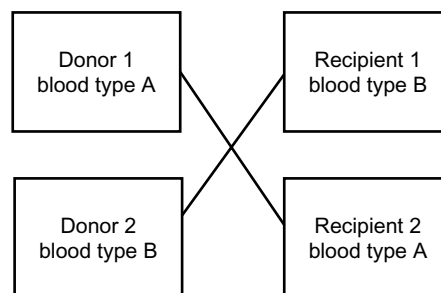

Figure 1(a): This diagram illustrates a two-way cyclic exchange between two blood-type-incompatible recipient-donor pairs, Recipient 1-Donor 1 and Recipient 2-Donor 2.

swapping donors between incompatible pairs, and the first such exchanges took place in Korea in the 1990s (Park et al. 1999). The first exchange in the United States occurred in 2000 at Rhode Island Hospital.

Originally, most kidney exchanges were conducted through simple cyclic exchanges, as Figure 1(a) depicts. Because of fear of possible reneging, such cyclic exchanges have been conducted simultaneously, making the exchange process a significant logistical challenge: four operating rooms and four surgical teams are required for the two nephrectomies and two transplants involved in the simplest exchange between two patient-donor pairs. For these reasons, cyclic exchanges with more than three patient-donor pairs are rarely conducted. Another type of exchange took the form of a chain. Chains using nondirected donors (NDDs) (i.e., kidney donors who decide to donate without having an intended recipient) were also conducted simultaneously at first, thus involving only two or three transplants at a time and ending with a transplant to a patient on the waiting list (see Figure 1(b)).

Contribution

The APD initially adopted design and optimization techniques for identifying short cycles and chains. It was, however, APD's introduction of an innovative and highly controversial approach based on NEAD chains that significantly increased the number of transplants that can be achieved through kidney exchange.

The added value derived from using NEAD chains is threefold. First, current clinical practice shows ample empirical evidence, substantiated by conclusions
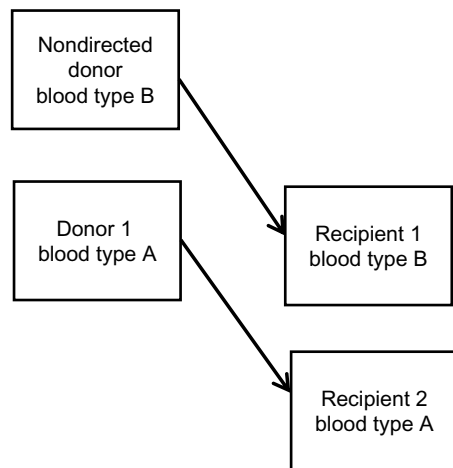
Figure 1(b): This diagram illustrates a two-transplant chain involving one nondirected donor, one blood-type-incompatible recipient-donor pair, Recipient 1-Donor 1, and one recipient, Recipient 2, who has no living donor to continue the chain.

drawn from mathematical models reflecting the composition of contemporary patient-donor pools (see Figure 4), that when exchanges are conducted in an optimal way, approximately 70 percent of transplants are conducted via NEAD chains. Thus, NEAD chains are responsible for the majority of successful kidney transplants conducted via kidney exchange at both APD and within other major exchange programs. By allowing only short (at most three) chains, the number of transplants would have been reduced by at least 12 percent, and more than 30 percent of the recipients would have been transplanted after a longer waiting period. The number of very highly sensitized transplanted patients (patients who have many antibodies and are thus not likely to accept a donor's kidney) would have been reduced by approximately 50 percent.

Second, unlike cyclic exchanges, chains can be conducted nonsimultaneously, while assuring that each patient in each patient-donor pair receives a kidney before the donor of the same pair donates the kidney. The relaxation of the simultaneity requirement significantly reduces the logistical overhead, making long chains possible and accomplishing many transplants. The longest NEAD chain to date—by the National Kidney Registry (NKR)—achieved 30 transplants, and involved 60 people, 30 donors, and 30 patients (Sack 2012).

Finally, NEAD chains can provide transplants for many hard-to-match patients. A patient receiving a donor's kidney needs to be both blood- and tissue-type compatible (i.e., the patient does not have an antibody that would lead to the rejection of the donor's kidney). Chains offer hope to highly sensitized patients who are likely to be incompatible with the people who wish to donate a kidney to them. Highly sensitized patients typically wait unusually long periods before a compatible kidney from a deceased donor becomes available, if one ever does. For many such patients, although no cyclic exchanges could be identified, chains involving these patients and originating with NDDs were found and implemented, providing their only kidney transplant option.

Conducting kidney exchanges through cycles or chains introduces the challenge of finding a maximal set of compatible matches, a classical OR combinatorial optimization problem, which we solve using integer programming techniques and insights from the theory of random graphs. Allowing chains to be long significantly increases the size of the optimization models and the resulting computational challenges.

Moreover, because kidney exchange is decentralized (the United States has more than 200 transplant centers and innumerable dialysis clinics), organizing kidney exchange is both an optimization problem and a serious market design and coordination problem. A successful broadscale approach to designing kidney exchange clearinghouses and protocols needs to take into account the goals and interests of the hundreds of existing institutions in the $50 billion-per-year business of caring for ESRD patients, because participating agents often have conflicting objectives. Thus, market design issues are central to this work.

We describe both the optimization approach and elements of the market design that the APD developed for conducting kidney exchange. In addition to technical issues, we emphasize the cultural and political obstacles to implementing long chains and our work to overcome them. We survey the steps, beginning in 2004, that led our team toward the APD's current successful practices. We detail its direct impacts through the APD, describe how other organizations have adopted nonsimultaneous chains, and provide estimates of the overall impact of these innovations.

Operations Research for Kidney Exchange: Market Design and Optimization

Early Experience: Exchanges with Short Simultaneous Cycles and Chains

In 2000, Michael Rees, the founder-to-be of the APD, and his father, Alan, wrote the first computer program to heuristically identify all possible two-way cyclic kidney exchanges within a pool of incompatible pairs. This program applied a utility score to each exchange and produced a rank-ordered list of all possible exchanges. In 2003, Michael Rees and Jonathan Kopke (at the University of Cincinnati) modified the program to be Web-based and track participating patients and donors from registration through transplantation. The Ohio Solid Organ Transplantation Consortium (OSOTC) living-donor kidney exchange program utilized this software from 2004 until 2006, when the OSOTC was disbanded.

OR methods were first applied to kidney exchange when Roth et al. (2004) proposed organizing kidney exchange on a large scale, including integrating exchange cycles and chains. They proposed an allocation mechanism that made it safe for participating patients and their surgeons to reveal relevant information (e.g., their medical data and that of all their willing donors, and which donors in the pool are acceptable to donate to them). That proposal involved algorithms that could generate large cycles and long chains. The mathematical treatment of those algorithms dates back to the top-trading-cycle algorithm introduced by Shapley and Scarf (1974), who attributed it to David Gale. Roth (1982) established its dominant strategy incentive properties (i.e., properties that allow participants to safely reveal their true preferences) that would be relevant in the kidney exchange context. Abdulkadiroğlu and Sönmez (1999) adapted the algorithm to more general settings by modeling the allocation of dormitory rooms in a way that later provided a natural bridge to models for allocating kidneys.

In 2004, however, the surgical culture and infrastructure permitted only pairwise simultaneous exchanges. To avoid the possibility that one pair would give a kidney and then not receive one (e.g., because of reneging), it was (and remains) the practice that all surgeries in a cyclic exchange be conducted simultaneously. So, even the simplest exchange between two patient-donor pairs required four operating rooms and four surgical teams for simultaneously conducting the two nephrectomies and transplants.

Responding to this limitation, Roth et al. (2005a) considered how to organize kidney exchange around only pairwise exchanges. This constraint allows the application of classic graph theory methods, including Edmonds' matching algorithm (Edmonds 1965) and the Gallai-Edmonds decomposition (Gallai 1964). It also requires that the sets of matchable nodes in a graph constitute a ~~matroid~~. Roth et al. (2005a) were able to show how to achieve a maximal set of kidney transplants in a way that continues to give patients and surgeons incentives to reveal all relevant information (e.g., the number of willing donors associated with a patient in cases with more than one such donor). In contrast to the simple OSOTC rank-ordered program, the New England Program for Kidney Exchange (NEPKE), established in 2005, was the first multihospital kidney exchange program to optimize pairwise kidney exchanges (Roth et al. 2005b). In collaborating with NEPKE and gaining experience with pairwise exchanges, Saidman et al. (2006) and Roth et al. (2007) showed that efficiency gains could be achieved by incorporating chains (still short ones) and larger cyclic exchanges that would require only relatively modest additional surgical infrastructures. Three-way cycles and short simultaneous chains quickly began to be part of regular kidney exchange practice in almost all programs, including those of the APD. Roth et al. (2007) estimate that the potential gains from three-way exchanges result in at least 20 percent more transplants; under different assumptions in the evolution of the exchange pool, Ashlagi et al. (2013) estimate that gains are at least 50 percent more when three-way chains are included.

These proposals were instrumental in stimulating the founding of the APD in 2006 under the direction of Michael Rees. The APD's predecessor, the OSOTC, identified all two-way cycles, greedily ranked them, and thus did not find the largest number of two-way cycles; in contrast, the APD adopted optimization software, designed by Roth, Sönmez, and Ünver, to maximize the weighted number of transplants in existing pools. It uses weights to define and account

for priorities between different matches. This algorithm finds optimal solutions consisting of cycles and chains whose lengths are bounded by three or four, depending on the circumstances. The algorithm is based on solving an integer program in which each feasible cycle and chain is assigned a binary decision variable. Until 2007, all surgeries in a cycle or chain were implemented simultaneously, thus significantly limiting the length of these cycles and chains. For example, a three-way exchange required six operating rooms and six surgical teams (for three nephrectomies and three transplants).

Relaxing the Simultaneity Constraints and the Benefit of Long Chains

In 2006, Roth, Sönmez, and Ünver, in conjunction with Drs. Frank Delmonico and Susan Saidman, observed that the growing number of altruistic NDDs would allow the simultaneity constraint to be relaxed for transplant chains initiated by NDDs (Roth et al. 2006). When a chain is initiated by an NDD, it can be organized so that the NDD makes the initial donation and no patient-donor pair has to donate a kidney before receiving one. This reduces the cost of a broken link, because no pair would be left without a kidney after donating one. By reducing the cost of a broken link, we begin to realize the potential benefits of chains in which operations are conducted nonsimultaneously. The benefit is that nonsimultaneous chains can be

longer, because more operating rooms and surgical teams can be assembled over a longer time frame.

Nevertheless, practitioners in the transplant community raised objections about the use of nonsimultaneous chains. They were concerned that a broken chain or chains (e.g., because of reneging) would cause patients and the public to lose trust in the kidney exchange system. Dr. Rees disagreed and felt that the utilitarian gains would outweigh the risk of equity losses. Therefore, after careful planning to minimize this danger, the APD implemented a NEAD chain and began to carry out the first nonsimultaneous, long chain of kidney transplants in July 2007. In this first chain, short segments were identified and the indicated transplants were carried out simultaneously. The APD designated the last donor of each identified segment of the chain to be the bridge donor for future segments of the chain, after the intended recipient had already received a transplant. It trusted bridge donors to "pay it forward" in the future to the next segment of the chain when it was identified. Segments of the first NEAD chain were identified dynamically as the process unfolded, and transplants were sometimes separated by months and thousands of miles. By March 2008, this first NEAD chain included 10 transplants and 11 donors, 10 of whom had donated kidneys to 10 patients they did not know. Rees et al. (2009) reported it in the New England Journal of Medicine, making it easier for others to adopt this innovation (see Figure 2).

	MI	July 2007 AZ 1	July 2007 OH 2	Sept 2007 OH 3	Sept 2007 OH 4	Feb 2008 MD 5	Feb 2008 MD 6	Feb 2008 MD 7	Feb 2008 NC 8	March 2008 MD 9	March 2008 OH 10
Recipient PRA		62	0	23	0	100	78	64	3	100	46
Recipient ethnicity		Cauc	Cauc	Cauc	Cauc	Cauc	Hisp	Cauc	Cauc	Cauc	AA
Relationship		Husband Wife	Mother Daughter	Daughter Mother	Sister Brother	Wife Husband	Father Daughter	Husband Wife	Friend Friend	Brother Brother	Daughter Mother

* This recipient required desensitization to blood group (AHG Titer of 1/8).

\# This recipient required desensitization to HLA DSA by T and B cell flow cytometry.
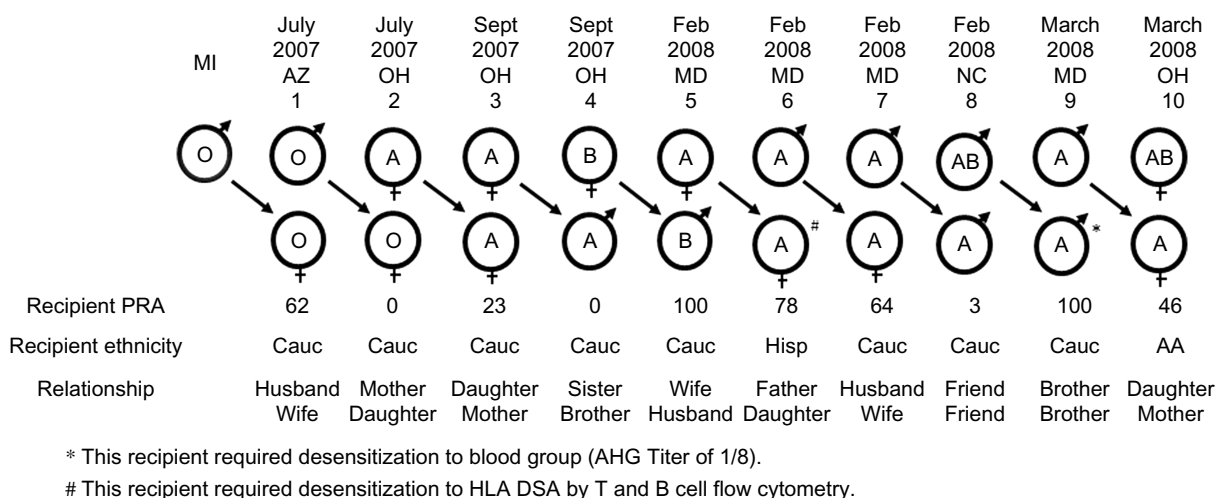
Figure 2: The graph shows the first NEAD chain, as reported by Rees et al. (2009).

The end of the chain was the 11th donor, whose blood type was AB and the last bridge donor at the time Rees et al. (2009) was published. She eventually waited three years before she donated to continue the chain. AB is the least frequent blood type in the population and most AB patients who join kidney exchange are highly sensitized, because their blood types are never incompatible with those of any donor; thus, AB patients can only be incompatible with their donors if their tissue types are incompatible. Exchange programs today seldom designate an AB donor as a bridge donor to continue a chain within the kidney exchange pool, but rather close the chain with a donation of the AB donor's kidney to a patient on the (long) waiting list for deceased donors. That initial chain was eventually extended to 16 transplants, and only ended when the APD participated with the United Network for Organ Sharing (UNOS) in an exchange, and UNOS chose to end the chain with a donation to someone on the deceased-donor waiting list. Today, after the completion of many more nonsimultaneous chains, the incidence of reneging by bridge donors continues to be very low (about two percent). Note that chains can break for reasons other than a choice on the part of the bridge donor; for example, the medical situation of the donor or the intended recipient could change. In the latter case, the chain can often be repaired by finding a new recipient for the bridge donor.

Since 2007, several other kidney exchange matching services have adopted NEAD chains for arranging kidney exchanges in the United States. Most notably, the NKR built its matching service around this approach. Of the more than 1,000 transplants that the NKR has conducted, more than 67 percent of them have used NEAD chains. Although long chains may seem intuitively attractive, their advantages were not clear initially. Long chains might not increase efficiency if they simply identify transplants that could otherwise have been done in multiple short cycles. In addition, as we previously mention, many in the transplant community continued to worry that nonsimultaneous chains were fraught with the danger of eroding trust in the entire kidney exchange enterprise if broken chains became common.

Allowing chains to be nonsimultaneous and thus longer than three, meant optimizing over a huge solution space. For example, in a pool with 150 pairs, but only one NDD, the number of chains bounded by length three can reach up to a few thousand, and could include more than one million chains if we permit them to be of length six. We discuss further computational issues and algorithms in the *Algorithms, Optimization, and Implementation* section.

Simultaneous or Nonsimultaneous Chains? Criticism and Objections

Whether to proceed with NEAD chains or continue with short simultaneous chains became a major debate in the medical kidney transplant community in 2009. An influential team from Johns Hopkins Hospital argued that based on computer simulations, nonsimultaneous chains should not be conducted, because even a modest risk of a broken chain implied that more transplants would be accomplished via short simultaneous chains (Gentry et al. 2009), which they called domino-paired donation (DPD) chains. DPD chains include up to two incompatible pairs and end with a patient on the waiting list for deceased donors. Because the waiting list always contains many patients (today about 100,000), finding a compatible patient on the waiting list for the last kidney in a chain is always possible; in addition, if this transplant is done simultaneously with the others, a bridge donor who might break the chain by failing to donate at a later date becomes impossible.

We engaged in the debate and noted that the Hopkins team failed to capture some of the potential for NEAD chains to accomplish many transplants, because its model constrained NEAD chains to have three or fewer transplants in each round of optimization (Ashlagi et al. 2011a). We showed that the conclusions of the Hopkins teams were reversed when long chain segments were allowed. Moreover, most of the additional patients who received transplants from long chains rather than shorter ones would be very highly sensitized. This debate within the medical academic community continued, as Gentry and Segev (2011) and Ashlagi et al. (2011b) discuss, but did not prevent the continued adoption of long chains by the APD and a few other transplant networks. As additional favorable clinical experience and data accumulated, the debate was settled in favor of NEAD chains.

A team from Carnegie Mellon University also conducted simulations that led it to conclude that

although NEAD chains result in more transplants than DPDs, NEAD chain segments should be capped at four transplants (Dickerson et al. 2012). The national UNOS pilot program has adopted this constraint on chains.

As of this writing, the accumulated clinical success of NEAD chains, including some very long ones, has become so apparent that early opponents (e.g., Johns Hopkins Hospital) have now adopted them (Rector and Cohn 2013), making exchanges based on nonsimultaneous, potentially long chains a standard practice.

The Power of Long Chains: Mystery Explained

We now explain what makes kidney exchanges so potent based on nonsimultaneous chains , how this realization came to be, and how it has changed the practice of kidney exchange. The key reason that long chains increase the number of transplants is the high percentage of very highly sensitized patients in kidney exchange pools. Patient sensitivity is measured by the panel-reactive antibody (PRA), which is the likelihood that the patient will be tissue-type incompatible with a random donor. For example, if a patient has a PRA of 99 percent, then 99 percent of the blood-type-compatible potential donors for this patient will not be tissue compatible. By simulating a pool using a statistical process based on empirical medical characteristics, less than 10 percent of the general population patients are expected to have a PRA level of 80 or more. This is the case for the pool of ESRD patients on the waiting list for a deceased donor. However, in kidney exchange pools, and in particular in the APD pool, more than 40 percent of the patients have a PRA of at least 80. In addition, more than 65 percent of the APD patients with PRA levels of 80 to 100 have a PRA greater than 95. Figure 3 plots the percentage of patients who are registered in the APD pool and have a PRA of at least 95. The percentages of patients newly registered to the pool and with PRAs greater than 95 in 2010, 2011, and 2012 were 33, 33, and 30, respectively; therefore, both the stock of patients in the pool and the flow of new patients are much more highly sensitized than the general patient population.

This means that the directed graph depicting the compatibilities between members of a kidney exchange pool is very sparse. Patients in kidney
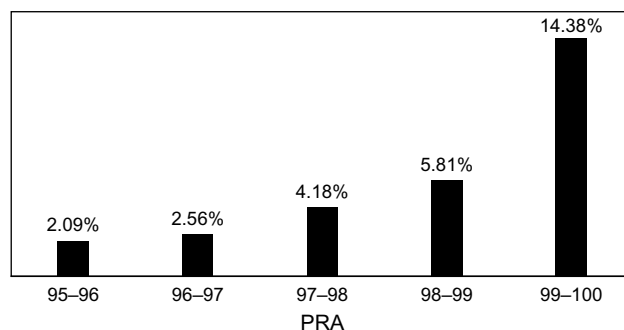

Figure 3: The percentage of patients in the APD pool who have a panel-reactive antibody (PRA) of at least 95 is heavily skewed toward the most highly sensitized patients.

exchange pools have such high PRAs, in contrast to the general patient population, for two primary reasons. First, as kidney exchange grew, hospitals became strategic players. Hospitals (or directors of transplant centers) can choose from larger sets of possible actions than can individual surgeons or patients, because hospitals see multiple patient-donor pairs. Because many major transplant centers have acquired experience with kidney exchange, hospitals often match their easy-to-match pairs internally (Roth 2008, Ashlagi et al. 2013), and only refer their difficult-to-match pairs to the kidney exchange networks. As a result, kidney exchange pools contain very large percentages of highly sensitized patients. Second, most highly sensitized patients and hard-to-match pairs naturally accumulate over time in the pool (up to a steady state, because patients cannot indefinitely remain in the pool).

Pools consisting primarily of highly sensitized patients, coupled with the shortage of NDDs, lead to the need for exchanges based on long chains. This is because when we represent which donors can give kidneys to which patients via a compatibility graph (in which the nodes are incompatible patient-donor pairs, and a directed edge goes from one pair to another if the kidney from the donor in the first pair is compatible with the patient in the second pair), the resulting compatibility graphs become sparse; as a result, exchanges that are based only on short cycles and chains are substantially suboptimal. Mathematically, the benefit of long chains in such pools can be illustrated by some classical facts from the theory of

random graphs. In the sparse Erdos-Renyi random graphs, the number of short cycles involving highly sensitized patients is very small. At the same time, there exists with significant probability a long chain covering a substantial number of nodes in the graph. Ashlagi et al. (2013) analyze a random graph with a mixture of dense (low-PRA patients) and sparse (high-PRA patients) parts. They show that in such graphs, allowing for long chains amounts to an order-of-magnitude increase in the number of patients participating in the exchanges.

To further illustrate the typical properties of pools of patients encountered in practice, consider a snapshot of APD data (see Figure 4). This figure describes the compatibility graph induced by the subset of patient-donor pairs such that all patients and all donors have blood type A. In particular, no blood type incompatibilities exist among these pairs. At the moment, represented by this compatibility graph, 38 such pairs have some compatibilities among themselves (18 such pairs were not compatible as patients or donors with any other of these pairs); each of the 38 pairs is represented by a node of the graph. An arrow points from one node to another if the kidney from the donor in the first pair is compatible with the patient in the second pair. Of the 38 pairs, 30 contain patients with high PRAs, which we depict as white nodes. We see that these nodes have few incoming edges. The nodes containing the eight low-PRA patients have substantially more incoming edges and are shaded in the figure. The dashed edges are parts of cycles. Note that no cycle contains only high-PRA patients, and only one cycle includes one high-PRA patient.

NEAD chains have become the regular practice at the APD, which now uses an optimization algorithm in which the length of chains is unconstrained. This is now also the practice at the NKR, the highest-volume kidney exchange program today. In the *NKR and UNOS KPD*: *Other Kidney Exchange Programs* section, we review data from two other kidney exchange programs.

Benefits

Results at the APD
Before describing the impact, we discuss how the typical matching process works at the APD. More than 80 transplant centers are currently registered with the APD (centers do not commit to enroll all of their pairs or accept matches). Each weekday after new incompatible pairs are entered into the APD pool or a potential match is abandoned (e.g., when a pair does not accept a match offer), the APD conducts a match run to find a maximum weighted number of matches, allowing both cyclic and long nonsimultaneous chain exchanges. The solution is based on the apparent compatibilities between patients and donors in the pool. The APD pool contains approximately 200 patient-donor pairs at any given time. Finding a solution typically takes three minutes or less; the *Algorithms, Optimization, and Implementation* section includes details. The solution is determined according to apparent compatibilities based on the medical information of patients and donors (blood tests must still be done before the transplant takes place). After the software finds an optimal solution, a central lab
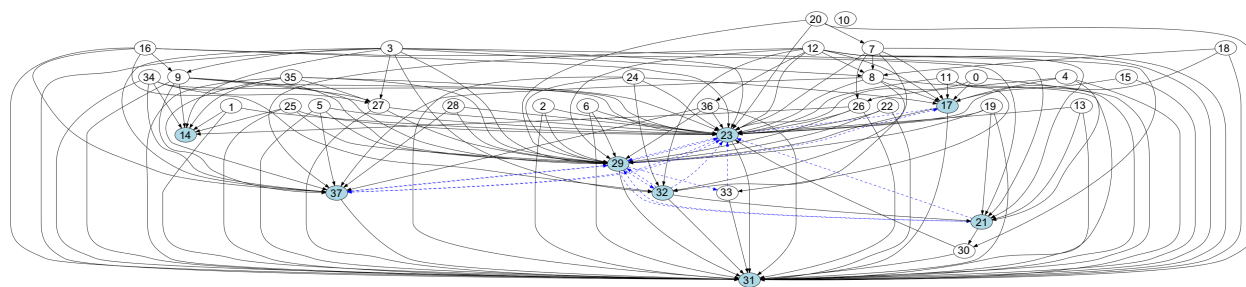

Figure 4: (Color online) This compatibility graph is of the patient-donor pairs, each with blood type A, from APD data. Most of these patients are highly sensitized; therefore, they cannot participate in cyclic exchanges, but can be reached via chains.

conducts crossmatch tests on the blood samples of the patients and donors to determine whether each recipient in the solution is tissue-type compatible with the intended donor. The transplant center of each patient who is part of the solution is then informed of the details about the patient's matched donor so that the center can accept the donor. Centers and (or) patients can refuse to accept offers. An exchange (chain or cycle) proceeds to the transplantation stage if a blood-test failure or donor rejection does not occur during these stages.

Significant benefits accrue directly to the transplanted recipients. They are freed from being tied to dialysis for many hours each week (e.g., three weekly sessions of four hours each), regain their health and ability to work, and have their lives extended. Financial benefits also accrue to the medical system and to third-party payers. The dialysis treatments that most ESRD patients use cost Medicare $87,272 per year per patient (United States Renal Data System 2013) and cost private insurers much more. Transplanting a single kidney-failure patient, at an average yearly cost of $32,922 (averaged over five years), rather than having that patient remain on dialysis, saves Medicare $271,750 over five years. Note the average survival rate for a dialysis patient is about five years.

Since 2007, the APD has conducted more than 220 transplants through kidney exchange, and more than 168 of these have used chains. As a result, we can estimate the financial impact of these transplants. If each transplant saves $271,750 over five years, the transplants of the APD, which exceed 200, have produced savings in excess of $54 million for the U.S. healthcare system. In a more conservative analysis of the savings from a single living-donor kidney transplant, which takes into account that transplant patients live longer than dialysis patients and therefore incur medical costs over a longer period, Matas and Schnitzler (2004) estimate the savings at $94,579, and the quality-adjusted life years (QALYs) gained at 3.5. Adding the value of QALYs, a living-unrelated-donor transplant saves $269,319, assuming society values additional QALYs from transplantation at the rate paid per QALY achieved by dialysis. Using this analysis, 168-plus living-donor transplants (through chains facilitated by APD) have saved $16 million and produced a value to society of $45 million.

The analysis in Wolfe et al. (1999) shows that the average patient who received a deceased-donor kidney transplant lived 10 years longer than that patient would have lived had he (she) stayed on dialysis. Therefore, APD transplants have given these 168-plus patients at least 1,600 extra years of life. These numbers are conservative, because the average living-donor kidney lasts twice as long as the average deceased-donor kidney.

Impact at Other Transplant Networks in the United States

In the United States, more than 2,600 kidney paired-donation (KPD) transplants have been performed since 2000 when the first two KPD transplants were performed at Rhode Island Hospital, and more than 2,250 KPD transplants have been done since 2008. To demonstrate the impact of NEAD chains, consider that in 2006, when exchanges were still constrained to be simultaneous, 74 KPD transplants and 68 nondirected living-kidney donations were performed in the United States. In 2012, more than 200 of the 528 paired-donation transplants performed in the United States resulted from NEAD chains.

According to UNOS, the number of transplants from kidney exchange in the United States has reached 2,658. We believe this number is a lower bound, because many centers conduct internal exchanges, which they might not report; in addition, the 2013 numbers are subject to change, because centers are still reporting their fourth-quarter numbers. Figure 5 shows the progress by year.

NKR and UNOS KPD: Other Kidney Exchange Programs

The NKR and UNOS KPD pilot programs both run multihospital exchange programs in the United States. The NKR, a private organization, was founded by Garet Hil in 2007 to facilitate actively recruiting NDDs and running NEAD chains. It is currently the highest-volume clearinghouse in the United States and has facilitated more than 1,000 transplants. UNOS is the federal contractor responsible for administering the waiting list for deceased-donor organs. Because it has working relationships with every U.S. transplant center, it is a natural candidate to organize a federally
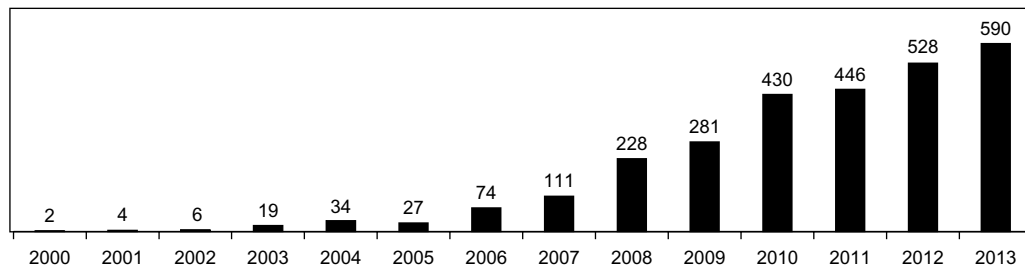
Figure 5: Since 2000, transplants through kidney exchange have increased.

sponsored kidney exchange that would unite all U.S. transplant centers. In October 2010, it launched a pilot KPD program, which has yet to become a major producer of kidney exchange transplants; however, since UNOS began to allow chains, it has been making progress toward this objective. Through 2013, it had conducted 79 transplants (52 of them in 2013). Since the launch of the NKR and UNOS programs, more than 1,600 pairs have enrolled in the NKR and more than 1,000 pairs have enrolled in UNOS (the enrollment rate was very low in the first years). The pools of both programs seem to have stabilized to approximately 220–250 pairs. Authors Ashlagi, Rees, and Roth are presently advisors to UNOS and members of the UNOS KPD workgroup. Ashlagi, Roth, and Rees have advised the NKR, and Ashlagi and Roth collaborated with the NKR to develop and validate their current optimization algorithm.

Chains play a significant role in both programs. Approximately 88 percent of the transplants facilitated by the NKR have been achieved through chains (176 chains overall). More than 28 percent of the transplanted patients had PRAs greater than 80 and more than 15 percent had PRAs greater than 95; in the past two years, this latter percentage increased to 20

percent. As of January 30, 2014, the NKR had conducted 179 chains, which accomplished 823 transplants, and 44 cyclic exchanges, which accomplished 103 transplants. More than one-fourth of these transplants were accounted for by the 16 chains involving 10 or more transplants, while the 96 chains involving three or fewer transplants accounted for only about 10 percent of these transplants. Therefore, although the long chains are rare (fewer than 10 percent of the chains), they account for a disproportionate share of the transplants.

Figure 6 shows the lengths of all chains conducted by the NKR. Note the long tail; most chains are short, but the longest 2.3 percent of chains account for more than 11 percent of all transplants. More than 13 percent of all transplanted patients with PRAs greater than or equal to 97 have been transplanted through those chains. In addition, 73 percent of the transplanted patients with PRAs greater than or equal 97 were transplanted through chains or cycles that have a length of at least four, while only 67 percent of all transplanted patients have been transplanted through chains and cycles with lengths of at least four.

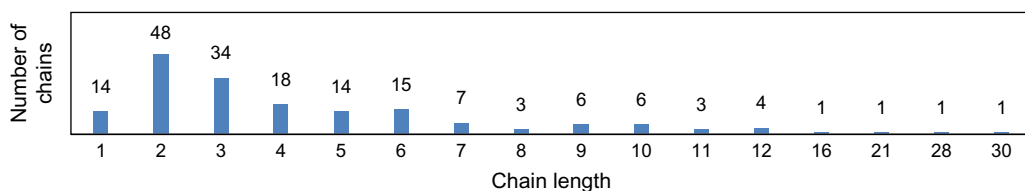The experience of the UNOS KPD pilot program differs somewhat from that of the NKR. At first,

Figure 6: (Color online) This graph shows the number and lengths of chains that the NKR has achieved.

UNOS searched for only two-way exchanges; however, in 2011, the UNOS algorithm was adapted to find chains unrestricted by length. As a result of match-offer refusals and crossmatch failures, UNOS had a high failure rate after identifying matches, but prior to performing transplants; therefore, since October 2012, it has limited chain segments to four and often nonsimultaneously conducts such chains. About 15 percent of the transplants UNOS has facilitated have been achieved through two-way cycles, 40 percent through three-way cycles, and 45 percent through chains; even with the restriction on chain length, most UNOS transplants have come from chains and three-way cycles.

UNOS and the NKR had different experiences in executing their kidney exchange programs, leading them to implement slightly different policies; however, NEAD chains play a major role in successful transplants in both programs. The NKR has been successful in attracting NDDs. This success is largely because it provides incentives to transplant centers to share the NDDs who approach it; it guarantees a center that it will end a chain with a patient within that center. Not all chains need to be long when so many altruistic donors exist. Most of the very short chains conducted by the NKR are done because they include a patient who is very highly sensitized, the quality of the match is very high, or the NDD is an AB donor, in which case is very unlikely to match an AB patient within a pair.

We cannot overemphasize the importance of matching the most highly sensitized patients through long chains, because many of these patients have little or no prospect of receiving a deceased-donor kidney. Data from the Organ Procurement and Transplantation Network for the United States show that the median waiting time for patients with PRAs greater than or equal to 80 was 13.5 years for candidates listed in 1999–2000; the median time for such patients who began waiting in 2003–2004 cannot be computed yet because fewer than half of them have been transplanted. As we have seen, NEAD chains serve a disproportionate share of very highly sensitized patients (i.e., patients with PRAs greater than or equal to 95).

Finally, we note that the number of NDDs has significantly increased since the advent of NEAD chains—in part because of the added incentive of being able to help more than one patient as a result of one's altruism, and perhaps also because of the publicity that sometimes accompanies long chains.

Algorithms, Optimization, and Implementation

The APD and other kidney exchange programs organize transplants by regularly searching the compatibility graph generated by the current pool of patients and donors for the maximum weighted number of transplants that can be achieved through cycles and chains. It is convenient to think of the pool as a compatibility network described by a directed graph $G(V, E)$. In this graph, each patient-donor pair and each NDD is a node $v \in V$. There is an edge between two nodes v_1 and v_2 if the donor of node v_1 is compatible with the patient of node v_2 (according to the medical data). Edges are associated with weights to capture some priority (i.e., importance, urgency) for conducting the corresponding matches. At the present time, weights are often assigned based on the subjective judgments of a medical board, but efforts are underway to make the choice of weights depend more on well-defined objectives for the overall set of matches. In addition to medical factors, a board may consider other factors, including the geographic residences of the parties and the length of time a patient has been on the waiting list. Typically, higher-sensitized patients are assigned higher weights.

The optimization problem considered prior to the introduction of nonsimultaneous chains had the following form (the OSOTC program did not optimize and the APD used optimization and unrestricted length of chains and cycles from the outset):

Maximize weighted number of matched pairs,
s.t. each pair is matched at most once via a cycle or chain;
cycle length at most 3;
chain length at most 3.

Existing pools typically contain about 200 pairs and up to 10 NDDs; some NDDs are bridge donors and some are altruistic donors. An altruistic donor initiates a new chain.

At first, kidney exchange programs considered algorithms that searched for an optimal solution, allowing up to three-way cycles and chains, as we

describe in the previous formulation. The problem is solved using an integer programming technique by introducing for each cycle a binary decision variable with a length of at most three and each chain with a length of at most three. This formulation works well when optimizing over the pools observed in practice.

Abraham et al. (2007), and earlier Kevin Cheung and Michel Goemans (personal comm.), have shown that solving the optimization problem we describe previously with bounded-chain or cycle lengths is an NP-hard problem. In practice, however, computational issues have not been a serious issue when both chains and cycles are bounded by a short length. Abraham et al. (2007) developed an algorithm based on column generation to solve such problems for very large pools. The APD briefly used their algorithm and UNOS now uses an extended version of it. It allows chains with segments with lengths of no more than three; UNOS has changed its policy to incorporate these chains.

Allowing for longer chains makes the previous formulation hard to solve in both theory and practice, because the potential number of chains grows very fast, as the allowable chain length grows; for example, in a historical pool of 150 pairs, the number of chains of length 6 beginning with an O donor easily reaches over five million.

In *The Recursive Algorithm* section next, we describe two algorithms that we developed to solve the optimization problem when long chains are permitted. These algorithms scale very well on the existing pools and find optimal solutions in a short time in all encountered instances. We use both algorithms to find matches.

The Recursive Algorithm

We use this algorithm to solve an optimization problem using constraint generation without assigning a variable for every chain, but instead introducing flow-conservation constraints. In the graph $G(V, E)$, we use a binary decision variable y_e that determines whether edge e (from a donor to a compatible recipient) will be chosen. Let $in(v)$ be the set of edges that point to node v and let $out(v)$ be the set of outgoing edges from node v. For each edge e, let w_e be the weight associated with edge e. We further add an edge, with zero weight, from every node corresponding to an incompatible pair to every node corresponding to a

NDD. This is done so that every chain originating from a NDD can be extended to a cycle beginning and ending with a NDD; as a result, our model and algorithm need to explicitly find only cycles (not chains). Our algorithm uses the following formulation (see the appendix for a formal formulation):

Recursive:

Maximize weighted flow,

s.t. Total flow out of a pair is at most the total flow that goes into a pair;

Total flow out of an altruistic donor is at most 1;

Total flow that goes into a pair is at most 1;

Flow on each edge is binary.

This is an integer programming formulation, which finds the largest disjoint cycle cover of a graph. As such, this formulation ignores for now bounds on the length of the cycles and chains. To incorporate a bound on the length of the cycles (e.g., to be at most 3) and also allow chains originating from NDDs to be arbitrarily long, we run the algorithm recursively by eliminating long cycles one at a time. Specifically, our recursive algorithm works as follows: It first solves the previous formulation . It then inspects the optimal solution. If the optimal solution outputs a cycle C of size more than 3, which does not include a NDD (and thus is not a feasible cycle), we add a constraint to resolve the recursive formulation with an additional constraint that eliminates this cycle. This constraint has the form in which the total flow in the edges of the cycle is at most the size of the cycle minus one.

We then solve the integer programming problem again with this additional constraint, and repeat the procedure until we identify a solution without cycles with length 4 or more. This recursive algorithm works well in many instances (see *The Prize-Collecting Traveling-Salesman-Problem-Based Algorithm* section next); however, in many other instances, the running time is too long (e.g., exceeds 20 minutes). To solve the remaining instances, we use an alternative integer programming formulation based on the prize-collecting traveling-salesman problem (PC-TSP). Next, we describe this formulation and the running times corresponding to both formulations.

The Prize-Collecting Traveling-Salesman-Problem-Based Algorithm

In the TSP, one is given a graph and must find a route that visits each node exactly once at minimum cost. In

the PC-TSP, the goal is to find a route in the directed graph that visits each node at most once, while paying some penalty for each node that is not visited. Qualitatively, the PC-TSP is similar to the optimization problem we are facing; we want to find long paths in a graph without visiting each node. In the appendix, we present our solution, which is very similar to the solution to the PC-TSP. As in the PC-TSP, the formulation has an exponential number of constraints that have a "cut" interpretation. The solution relaxes those constraints, but more aggressively attempts to find violating constraints by using a technique called cutting planes; the appendix provides details.

Table 1 describes the running times of both algorithms for what we call difficult instances—instances involving at least 500 constraints in the recursive

Instance information			Running time (secs)	
NDDs	Pairs	Edges	Recursive	PC-TSP
3	202	4,706	0.18	0.255
10	156	1,109	4.425	1.069
6	263	8,939	16.186	11.055
5	284	10,126	28.063	16.03
6	324	13,175	143.432	
6	328	13,711	150.877	27.67
6	312	13,045	1,200*	1,200*
10	152	1,125	10.388	0.245
3	269	2,642	13.896	0.056
10	257	2,461	16.206	0.113
7	255	2,390	16.7	0.108
6	215	6,145	44.101	2.237
10	255	2,550	103.112	0.136
1	310	4,463	177.582	0.151
11	257	2,502	201.154	0.154
6	261	8,915	340.312	3.829
10	256	2,411	347.791	0.119
6	330	13,399	522.619	6.507
10	256	2,347	683.949	0.121
7	291	3,771	1,200*	0.163
8	275	3,158	1,200*	0.306
4	289	3,499	1,200*	0.376
3	199	2,581	1,200*	1.943
7	198	4,882	1,200*	8.255
2	389	8,346	1,200*	16.076

Table 1: The table describes the performance of the recursive and PC-TSP algorithms for difficult instances (i.e., instances that generated at least 500 violating constraints). *Note.* **Instances for which the optimal solution was not found within 20 minutes are indicated by** *. **In instances above the midline, both algorithms have running times within the same order of magnitude; for instances below the midline, the PC-TSP-based algorithm was faster by at least an order of magnitude.**

algorithm and analogous cut constraints in the PC-TSP algorithm. We generated these instances while running each algorithm on real data from the APD and NKR programs.

Anderson et al. (2014) developed software, which the APD is currently testing, based on currently running both algorithms described previously and generating the optimal solution produced by the algorithm that finds this solution. Other exchange programs, including those at Northwestern Medical Center, Methodist Hospital in San Antonio, and Georgetown Medical Center, presently use this software. This software is also used to suggest kidney exchanges as a result of merging the NKR, APD, and the Methodist Specialty and Transplant Hospital pools; therefore, it identifies matches in the largest combined pool in the United States—more than 600 pairs.

Conclusions

Kidney exchange has become a standard part of transplantation in the United States, and the innovative ideas of many researchers and practitioners have played an important role in this success. One distinct innovation that has had a profound effect in increasing the number of transplants has been the APD's introduction of long nonsimultaneous chains.

Since nonsimultaneous chains were introduced, more than 75 percent of the 2,600 transplants have been conducted through NEAD chains. The APD was both the first to implement such chains and the first to optimize matches that incorporated them; the APD team also collaborated in assisting other centers and kidney exchange programs to adopt this methodology. Over time, other clearinghouses have also adopted NEAD chains; notably, the NKR, which is presently the highest-volume clearinghouse for kidney exchange, has facilitated more than 85 percent of its transplants through nonsimultaneous chains. Early opponents of long nonsimultaneous chains have adopted them, and no longer question their effectiveness. Furthermore, early ethical concerns for choosing to give a NDD kidney to a patient other than the first patient on the deceased-donor waiting list have been overridden by the profound increase in the number of transplants resulting from NEAD chains.

The very highly sensitized patients would seem to benefit most from the use of long chains, because

such chains are often the only way these patients can receive a transplant. The nature of exchange is that it benefits all parties, and the benefits flow not only to all participants in kidney exchange, but also to patients who do not have a donor and are waiting for a deceased-donor kidney. Some of those patients benefit directly when they receive the last-donor kidney from a long chain; however, even patients who will eventually receive a deceased-donor kidney benefit from kidney exchange, because each transplant accomplished through kidney exchange potentially relieves the recipient waiting for a deceased-donor kidney from having to wait as long as would be necessary if no kidney exchange existed. Kidney exchange delivers living-donor kidneys to patients who are on the waiting list for deceased-donor kidneys; therefore, it shortens the wait for those who have no other option.

Important design issues remain to be solved if kidney exchange is to continue to grow. One important issue is to provide more incentives for transplant centers to fully participate (i.e., offer kidney exchange to their easy-to-match pairs and to those who are difficult to match). Ashlagi and Roth (2014) explore how a bonus mechanism, similar to the spirit of a frequent-flyer program, could be developed to provide appropriate incentives for centers to enroll their easy-to-match pairs. Incentives for centers to enter their NDDs are being implemented in the NKR and at the APD (a chain terminates with a patient who is associated with a center that entered a NDD). Most chains today end with patients on the deceased-donor waiting list.

Another important design issue concerns the financial arrangements involving hospitals that participate in kidney exchange; because each hospital has a different cost, these differences can present obstacles when hospitals are exchanging kidneys. One direction we are pursuing is to establish a standard acquisition charge for obtaining a kidney for transplant (Rees et al. 2012).

Market design draws on both optimization and game theory to address the operational aspects that constrain exchange and the incentive constraints that may also have to be satisfied. In the case of kidney exchange, the operational constraints involve assembling the resources to coordinate and conduct complicated surgical procedures across many hospitals.

The incentive constraints entail assembling the medical information from patients, surgeons, and hospitals, and ensuring that the information they reveal (e.g., the number and characteristics of their willing donors) is protected. These constraints also entail the incentives provided to hospitals for enrolling both their hard-to-match and easy-to-match patients and donors. The task of solving these problems is evolving as logistical capabilities improve and as the medical and economic environments change. For now, however, kidney exchange has made possible several thousand life-saving and life-enhancing surgeries, yielding millions of dollars in savings and thousands of additional years of life.

In closing, we note that when operations research began to emerge as a distinct discipline in the years following World War II, it was closely linked to economics; however, those links grew more tenuous over time. The development of kidney exchange and the field of market design more generally illustrate how close those links can still be, as Roth (2002) discusses.

Appendix

The Recursive Formulation
The following is a formulation used in the recursive algorithm:

Recursive:

$$\text{Maximize} \quad \sum w_e y_e,$$

$$\text{s.t.} \quad \sum_{e \in \text{out}(v)} y_e \leq \sum_{e \in \text{in}(v)} y_e \quad \text{for every pair } v;$$

$$\sum_{e \in \text{out}(v)} y_e \leq 1 \quad \text{for every altruistic donor } v;$$

$$\sum_{e \in \text{in}(v)} y_e \leq 1;$$

$$y_e \in \{0, 1\}.$$

Note that the recursive formulation does not restrict cycles to be short. After solving the algorithm, if the solution finds a cycle of length C that is too long, the following

constraint to rule out that long cycle is added before resolving the problem:

$$\sum_{e \in C} y_e \leq |C| - 1.$$

The Prize-Collecting Traveling-Salesman-Problem-Based Algorithm

We first present the algorithm for unbounded chain length, and then describe modifications to solve for any bounded chain length. To write the formulation, we need some notations. For each cycle C of length at most k, we introduce a binary variable z_C, which indicates whether we are using the cycle C. Each cycle C is also associated with a weight w_C, which is the sum of weights in all edges in the cycle. We use flow binary variables y_e for each edge e, as in the recursive algorithm, to indicate if we use the edge e. We define decision variables f_v^i and f_v^o to be the flow in and out of node v respectively. We define N and P to be the set of NDDs and pairs in the graph, respectively. Define $C_k(v)$ to be all the cycles of length at most k that include node v. We can now write the formulation as follows:

$$\max \left\{ \sum_e w_e y_e + \sum_C w_C z_C \right\},$$

$$\text{s.t.} \sum_{e \in \text{in}(v)} y_e = f_v^i, \quad v \in V, \tag{1}$$

$$\sum_{e \in \text{out}(v)} y_e = f_v^o, \quad v \in V, \tag{2}$$

$$f_v^o + \sum_{C \in C_k(v)} z_C \leq f_v^i + \sum_{C \in C_k(v)} z_C \leq 1, \quad v \in P, \tag{3}$$

$$f_v^o \leq 1, \quad v \in N \tag{4}$$

$$\sum_{e \in \text{out}(S)} y_e \geq f_v^i, \quad S \subseteq P, v \in S, \tag{5}$$

$$y_e \in \{0, 1\}, \, e \in E, \, z_C \in \{0, 1\}, \, C \in C_k(v).$$

Constraints (3) make two statements: first, the amount of flow out of pair v is at most the amount of flow that goes into pair v; second, the sum of the amount of flow that goes into pair v and the number of cycles to which that pair v is assigned is at most 1. Constraints (4) say that the amount of flow leaving each NDD is at most 1.

Constraints (5) are very similar to the cutset inequalities for the TSP, as adapted to the PC-TSP (Goemans 2009). Essentially, they work as follows. Suppose a chain is reaching some node v, and as a result, f_v^i equals 1. Now suppose that we divide the graph into two pieces, such that the half containing v, which we denote as S, does not contain any of the NDD nodes from N. Because each chain begins at some NDD (and thus not in S), in order for the chain to reach $v \in S$, it must use an edge that does not begin in S and ends in S. Thus, the constraint requires that whenever there is flow into v, for every way that v can be cut off from the NDDs, there is at least this much flow over the cut.

As in the recursive algorithm, there are exponentially many constraints of type (5); therefore, we can use the same recursive heuristic as in the recursive algorithm. Instead, although we relax the constraints (5), we use an efficient algorithm to identify several violating constraints and add them to the formulation sooner than in the recursive algorithm. Identifying these constraints is known as the separation problem in optimization. The separation problem for constraints (5) can be solved by solving $O(P)$ network-flow problems (Anderson et al. 2014).

When chains are required to be bounded by length, we slightly modify the formulation as follows. For each NDD n and each edge e, we introduce auxiliary variables y_e^n and $f_v^{i,n}, f_v^{o,n}$ indicating flow that must begin at node n. We then add a few categories of constraints. First, we enforce that for each edge e, the flow on y_e must come from exactly one of the y_e^n. Second, we define the same relationship between $f_v^{i,n}, f_v^{o,n}$, and y_e^n as in constraints (1) and (2) from the PC-TSP formulation. Third, we require that for each NDD n and each vertex v, we have $f_v^{o,n} \leq f_v^{i,n} \leq 1$. Finally, we add the maximum chain-length constraint, saying that for each NDD n, we use a total number of edges y_e^n that is at most the maximum chain length.

References

Abdulkadiroğlu A, Sönmez T (1999) House allocation with existing tenants. *J. Econom. Theory* 88(2):233–260.

Abraham DJ, Blum A, Sandholm T (2007) Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. Accessed July 4, 2013, http://www.cs.cmu.edu/~sandholm/kidneyExchange.EC07.withGrantInfo.pdf.

Anderson RI, Ashlagi I, Gamarnik D, Roth AE (2014) Using the traveling salesman problem to optimize kidney paired donation. Working paper, Massachusetts Institute of Tehnology, Cambridge, MA.

Ashlagi I, Gamarnik D, Rees MA, Roth AE (2013) The need for (long) chains in kidney exchange. Working paper, Massachusetts Institute of Technology, Cambridge, MA.

Ashlagi I, Gilchrist DA, Roth AE, Rees MA (2011a) Nonsimultaneous chains and dominos in kidney-paired donation-revisited. *Amer. J. Transplantation* 11(5):984–994.

Ashlagi I, Gilchrist DS, Roth AE, Rees MA (2011b) Nead chains in transplantation. *Amer. J. Transplantation* 11(12):2780–2781.

Ashlagi I, Roth AE (2014) Free riding and participation in large scale, multi-hospital kidney exchange. *Theoret. Econom.* 9(3):817–863.

Cheung K, Goemans M (2014) Personal communication with Alvin Roth.

Dickerson JP, Procaccia AD, Sandholm T (2012) Optimizing kidney exchange with transplant chains: Theory and reality. Accessed July 4, 2013, http://www.cs.cmu.edu/~sandholm/chains.aamas12.pdf.

Edmonds J (1965) Paths, trees, and flowers. *Canad. J. Math.* 17(3):449–467.

Gallai T (1964) Maximale Systeme unabhängiger Kanten. *Magyar Tud. Akad. Mat. Kutató Int. Közl.* 9(9):401–413.

Gentry SE, Montgomery RA, Swihart BJ, Segev DL (2009) The roles of dominos and nonsimultaneous chains in kidney paired donation. *Amer. J. Transplantation* 9(6):1330–1336.

Gentry SE, Segev DL (2011) The honeymoon phase and studies of nonsimultaneous chains in kidney-paired donation. *Amer. J. Transplantation* 11(12):2778–2781.

Goemans MX (2009) Combining approximation algorithms for the prize-collecting TSP. Accessed July 4, 2013, http://arxiv.org/abs/0910.0553.pdf.

Matas AJ, Schnitzler M (2004) Payment for living donor (vendor) kidneys: A cost-effectiveness analysis. *Amer. J. Transplantation* 4(2):216–221.

Park K, Moon JI, Kim SI, Kim YS (1999) Exchange-donor program in kidney transplantation. *Transplantation* 67(2):336–338.

Rapaport FT (1986) The case for a living emotionally related international kidney donor exchange registry. *Transplantation Proc.* 18(3):5–9.

Rector K, Cohn M (2013) At Hopkins, kidney transplants occur in chain reactions. Accessed March 1, 2014, http://www.baltimoresun.com/health/bs-hs-kidney-donations-20131230-story.html.

Rees MA, Kopke JE, Pelletier RP, Segev DL, Rutter ME, Fabrega AJ, Rogers J, et al. (2009) A nonsimultaneous, extended, altruistic-donor chain. *New England J. Medicine* 360(11):1096–1101.

Rees MA, Schnitzler MA, Zavala E, Cutler JA, Roth AE, Irwin FD, Crawford SW, Leichtman AB (2012) Call to develop a standard acquisition charge model for kidney paired donation. *Amer. J. Transplantation* 12(6):1392–1397.

Roth AE (1982) Incentive compatibility in a market with indivisibilities. *Econom. Ltr.* 9(2):127–132.

Roth AE (2002) The economist as engineer: Game theory, experimental economics and computation as tools of design economics. *Econometrica* 70(4):1341–1378.

Roth AE (2008) What have we learned from market design? *Econom. J.* 118(527):285–310.

Roth AE, Sönmez T, Ünver UM (2004) Kidney exchange. *Quart. J. Econom.* 119(2):457–488.

Roth AE, Sönmez T, Ünver MU (2005a) Pairwise kidney exchange. *J. Econom. Theory* 125(2):151–188.

Roth AE, Sönmez T, Ünver UM (2005b) A kidney exchange clearinghouse in New England. *Amer. Econom. Rev. Papers Proc.* 95(2):376–380.

Roth AE, Sönmez T, Ünver UM (2007) Effcient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *Amer. Econom. Rev.* 97(3):828–851.

Roth AE, Sönmez T, Ünver UM, Delmonico FL, Saidman SL (2006) Utilizing list exchange and nondirected donation through "chain" paired kidney donations. *Amer. J. Transplantation* 6(11):2694–2705.

Sack K (2012) 60 lives, 30 kidneys, all linked. Accessed February 1, 2014, http://www.nytimes.com/2012/02/19/health/lives-forever-linked-through-kidney-transplant-chain-124.html.

Saidman SL, Roth AE, Sönmez T, Ünver UM, Delmonico FL (2006) Increasing the opportunity of live kidney donation by matching for two- and three-way exchanges. *Transplantation* 81(5):773–782.

Shapley LS, Scarf H (1974) On cores and indivisibility. *J. Math. Econom.* 1(1):23–28.

United States Renal Data System (2013) Annual data report: Atlas of chronic kidney disease and end-stage renal disease in the United States, Part 2. National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Diseases, Bethesda, MD.

Wolfe RA, Ashby VB, Milford ML, Ojo AO, Ettenger RE, Agogoa LYC, Held PJ, Port FK (1999) Comparison of mortality in all patients on dialysis, patients on dialysis awaiting transplantation, and recipients of a first cadaveric transplant. *New England J. Medicine* 341(23):1725–1730.

Ross Anderson is a PhD student in the Operations Research Center at MIT. His research interests are in applied probability, optimization, and healthcare applications. He graduated from Cornell University with a BS in operations research and computer science.

Itai Ashlagi is an assistant professor of operations management at the Sloan School of Management of Massachusetts Institute of Technology. He received M.Sc. and PhD in operations research from the Israel Institute of Technology. He was a post doctorate at Harvard University. He is interested in topics on the border of microeconomic theory and operations research and in particular his research focuses on topics in market design and game theory. He received the outstanding paper award from ACM Conference on Economics and Computation.

David Gamarnik is a professor of operations research at the Sloan School of Management of Massachusetts Institute of Technology. He received a BA in mathematics from New York University in 1993 and a PhD in operations research from MIT in 1998. Since then he was a research staff member of IBM T.J. Watson Research Center, before joining MIT in 2005. He is a recipient of the Erlang Prize and the Best Publication Award from the INFORMS Applied Probability Society, IBM Faculty Partnership Award, and several NSF-sponsored grants. He is an area editor of *Operations Research*, associate editor of *Mathematics of Operations Research, Stochastic Systems*, and *Queueing Systems*.

Michael Rees obtained his MD from the University of Michigan and his PhD in immunology and transplant surgery fellowship at the University of Cambridge, England. He directs the University of Toledo transplant program and the Alliance for Paired Donation. He initiated the first nonsimultaneous kidney paired donation chain of transplants and the first international paired exchange chain. He also studies innate immunity in xenotransplantation.

Alvin E. Roth is the Craig and Susan McCaw Professor of Economics at Stanford University, and the George Gund Professor *Emeritus* of Economics and Business Administration at Harvard. He received a BS in operations research from Columbia University in 1971 and a PhD in operations research from Stanford University in 1974. He directed the redesign of the National Resident Matching Program, through which approximately 25,000 doctors a year find

their first employment as residents at American hospitals. He has also helped in the reorganization of the market for more senior physicians, as they pursue subspecialty training, and in other labor markets. He helped design the high school matching system used in New York City to match approximately 90,000 students to high schools each year. He also helped redesign the matching system used in Boston Public Schools, for students of all ages. More recently he has helped design school choice systems in several other large American cities. He is one of the founders and designers of kidney exchange in the United States that helps incompatible patient-donor pairs find life-saving compatible kidneys for transplantation. He shared the 2012 Nobel Memorial Prize in Economic Sciences.

Tayfun Sönmez is a professor of economics at Boston College and a distinguished research fellow at Koç University. His research agenda focuses on the theory and practice of market design, with an emphasis on its applications in matching markets. He is best known for his contributions to the formulation, analysis, and practical design of kidney exchange and school choice. He is one of the pioneers of the market design approach to school choice that led to adoption of improved matching mechanisms in several U.S. school districts including New York City and Boston. He is one of the pioneers of the market design approach to kidney exchange and a founder of the New England Program for Kidney Exchange. He is a fellow of the Econometric Society.

M. Utku Ünver is a professor of economics at Boston College and a Distinguished Research Fellow at Koç University. His research focuses on the theory and practice of matching market design. Being one of the pioneers of market design approach for kidney exchange clearinghouses, his research found wide application in the United States and abroad. He is one of the founders of the New England Program for Kidney Exchange and one of the founding collaborators for the Alliance for Paired Donation. He also served as an at-large representative for the United Network for Organ Sharing task committee on kidney exchange. His other market design related research focuses on adoption of kids, tuition exchange, house allocation and exchange, and course bidding at business schools. He is the president of the Society for Economic Design.