Proceedings of the 2020 Winter Simulation Conference
K.-H. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, eds.

A SIMULATION MODEL FOR THE MULTI-PERIOD KIDNEY EXCHANGE INCENTIVIZATION PROBLEM

Kurtis Konrad

Industrial and Systems Engineering
North Carolina State University
Raleigh, NC 27607, USA

ABSTRACT

Kidney exchanges provide an opportunity for individuals who need a new kidney to effectively trade a donor's incompatible kidney for a compatible one. We present a mechanism for fully dynamic kidney exchanges that incentivizes transplant centers to truthfully participate in global matchings through a credit-based weighting scheme. Our mechanism incorporates both cycles and altruistically initiated chains while allowing patients to remain in the system for multiple time periods. Using simulation, we demonstrate that this credit-based matching mechanism is strategy-proof, individually rational, and efficient for all transplant centers under the assumption that all offered matches are accepted.

1 INTRODUCTION

Every year, thousands of people find out that, for one reason or another, they need a new kidney. Humans are naturally born with two kidneys, but only require one kidney to survive, so it is common practice for friends and relatives to donate a kidney to their sick loved one. Unfortunately, many people who desire to donate discover that they are unsuitable matches. This typically occurs because of differences in blood type, although there can be several other reasons for incompatibility. For many years, this unfortunate circumstance meant that the sick individual had to wait for a kidney from a deceased donor.

In the past few years, however, doctors have started implementing a new approach, often referred to as the kidney exchange. At its heart, the kidney exchange is about individuals donating their kidneys to strangers, so their loved ones can receive a kidney. In the most basic case of a pairwise exchange, donor A wants to give a kidney to patient A, and donor B wants to give a kidney to patient B, but unfortunately both pairs are unsuitable for each other. Therefore, donor A gives a kidney to patient B and donor B gives a kidney to patient A. This idea can easily be extended to three-way exchanges, or even larger ones.

Every year, approximately 35,000 individuals learn that they are in need of a kidney transplant. According to the Organ Procurement and Transplantation Network (OPTN 2020), there are currently over 90,000 patients in need of a kidney, with fewer than 20,000 transplants performed each year. These transplants occur for one of three reasons: a deceased donor's kidneys are recovered, a patient finds a living donor, or an altruistic individual donates a kidney to a stranger. In a monetary market, effective market pricing mitigates such imbalances of supply and demand, but in 1984, the federal government passed the National Organ Transplant Act, which made it illegal to monetize the kidney allocation market.

1.1 Literature Review

The kidney exchange is an intricate market that contains many difficult matching situations and requirements. Every patient has preferences for kidneys based on acceptable blood and tissue types. Roth et al. (2005b) determined that preferences can effectively be evaluated as 0 or 1 because patients should be indifferent among all acceptable donor kidneys. They also explain that tissue type compatibility is related to whether a patient has antibodies that would reject a donor protein. Other factors that may be important to success include height, weight, sex, and overall health.

978-1-7281-9499-8/20/$31.00 ©2020 IEEE

Another intricacy that must be taken into consideration is the various sources of kidneys. In this market, there is only one group that can receive organs. However, kidneys are donated from three different groups. Some come from deceased individuals, who have no preferences about how their organs are utilized. Others come from living donors, who would prefer to give to their patient, but who are primarily focused on ensuring their patient receives a kidney. Finally, some kidneys come from altruistic donors, who have no true preferences regarding how their kidneys are used. Initially, Roth et al. (2004) suggested treating the kidney exchange market as a housing allocation problem. Roth et al. (2005b) then observed that this market cannot be modeled as a two-sided market of donors and patients because there are relatively few donors without patient preferences. They therefore use a matching mechanism based on graphing and variations on the top trading cycle algorithm. Integer programming models of this problem have also been created, with various efficient solution methods presented, including a solution based on the travelling salesman problem, suggested by Anderson et al. (2015).

Ausubel and Morrill (2014) observe an additional complication in this market; no individual can be contractually obligated to donate a kidney. Donation must therefore be properly incentivized, and the donor's surgery must begin at least as early as the corresponding patient's surgery to prevent the donor from reneging. Since this implies the simultaneity of all related surgeries, and most hospitals cannot accommodate more than four surgeries at once, only pair-wise swaps are possible. However, Ausubel and Morrill (2014) suggest that this constraint can be relaxed by only requiring the donation to occur no later than the receipt of the associated kidney. By routinely implementing three-way and higher exchanges, we can create more efficient matchings. Roth et al. (2007) find that the gains from three-way exchanges are substantial, and they more effectively utilize type O donors.

1.1.1 Market Participants

There are two different market participants to analyze when designing kidney exchange markets: the patients and the hospitals. Most patients are already incentivized to participate in the market because they cannot easily find a compatible donor. By participating in the market, they improve their chances of receiving an acceptable kidney. Top trading cycles and integer programming mechanisms can both be implemented in a strategy-proof way such that patients will always report their true preferences in the market. The largest patient incentivization difficulty is encouraging compatible pairs to join the exchange. These pairs are often satisfied that they are compatible, so they do not consider joining a kidney exchange registry. However, Gentry et al. (2007) found that by including compatible pairs, the match rate for incompatible pairs nearly doubled. Since no proposed matching can be binding, these compatible pairs cannot lose anything by joining the market. In the worst-case scenario, they will simply refuse the proposed matching and donate as they had originally planned.

Santos et al. (2017) develop a simulation model of the kidney exchange to help policy holders fairly and effectively perform matchings. They utilize a fully dynamic methodology, allowing for many different types of patients and donors to participate. They also incorporate the possibility that offered matches fail, and they provide methods for reincorporating these failed matches back into the pool. Offered matches are those suggested by the exchange process, and they can fail for a variety of reasons, such as donor reneging, significant health changes, and failed tissue cross-matches.

The other group to consider is the hospitals; from a management perspective, hospitals must generate revenue. This basic business requirement often leads hospitals to join organ exchange programs, but to only reveal difficult-to-match patients to the exchange. Hajaj et al. (2015) suggest a strategy-proof, individually rational, and efficient mechanism that incentivizes hospitals to truthfully participate in the market through a credit program (ie. reveal all patients). Theirs is a quasi-dynamic model, in that hospitals participate in the market for multiple periods, but patients do not. For every pair that they reveal to the market, they are awarded arbitrary bookkeeping units called credits based on their expected number of pairs. The system of credits is useful because as a hospital gains additional credits, it improves the numbers of matches that it receives. These credits carry through time, and thus hospitals have no incentives to cheat, and a strategy-proof and efficient mechanism has been created. This is only possible because hospitals are allowed to exist between matchings.

1.2 Our Contribution

In this paper, we will examine ways to expand the market design proposed by Hajaj et al. (2015). Their setup is only quasi-dynamic because the authors choose to limit their scope to static patients. If a patient does not receive a match, he or she is discarded from the model. We expand their model into a truly dynamic one, where unmatched patients remain in the model until they receive a kidney or renege from the market. By creating a fully dynamic model that allows both patients and hospitals to exist for multiple periods, we present a more accurate picture of reality as dialysis enables kidney transplant patients to wait several periods for a successful match.

We propose a mechanism modification and use simulation to show its effectiveness, loosely following the example of Santos et al. (2017). We will attempt to show that this mechanism is strategy-proof, individually rational for the hospitals, and pareto efficient for the kidney exchange market and that this is a robust mechanism with respect to potential input parameters.

The remainder of the paper is organized as follows. We describe a model of the kidney exchange market with both multi-period hospitals and patients. We explain the steps of our credit mechanism and the matching process. Finally, we present a simulation that implements our mechanism for the kidney exchange and fulfills our desired mechanism characteristics.

2 DYNAMIC MODEL OF THE EXCHANGE

As with most kidney exchange problems, and the formulation used by Hajaj et al. (2015) in particular, we consider a set of n transplant centers $T = \{1, 2, \ldots, n\}$. In each period, we assume each center receives a non-empty set of patient-donor pairs $V_i = \{v_i^1, v_i^2, \ldots, v_i^N\}$. To motivate participation in the centralized matching mechanism, we assume that there is a positive probability for each center in each time period that it receives a pair that it cannot match internally.

There are two types of matches that can be created by the mechanism: cycles and chains. Chains are initiated by altruistic donors, who are treated as a standard patient-donor pair. In this case, the patient is a representation of either the deceased-donor waiting list, or a bridge donation to the next period. Therefore, the last donor in the chain donates to this 'altruistic patient'. By assuming this surgery also occurs at the center where the altruist began the chain, we can show that the number of surgeries that a center receives is identical to the number it would have received if the altruist was a regular patient-donor pair. We can therefore treat altruists as pairs, and we can say that every pair that arrives to the model v_i^j has a constant independent probability of being altruistic.

Each center has a known capacity K_i, which is the number of active patients that a center can have. Hajaj et al. (2015) suggest that the capacity is a function of the average number of arrivals k_i, $K_i = f(k_i) = \alpha \cdot k_i$, where they let $\alpha = 2$ for exposition purposes. They treat the arrival rate distribution of pairs to a center t_i as of a general form with this mean k_i. We innocuously rearrange the capacity function to model the mean number of arrivals as a function of capacity, $k_i = f(K_i) = K_i/\alpha$.

The most general process for the model proposed by Hajaj et al. (2015) is as follows and is shown below in Figure 1. Patients and donors arrive at the transplant centers, which then report these pairs to the model. Based on the number of pairs reported, the center receives an updated credit balance. The mechanism then determines the maximum total number of matches that are possible in the reported pool. In doing so, it uses a myopic approach that is common practice in most kidney exchanges. Using this global maximum as an additional model constraint, the mechanism determines the best and worst possible matchings for each center. If that center has a positive credit balance, it receives the higher number of matches. Otherwise, it gets fewer matches. Its credit balance is then updated based on the spread of the best and worst matchings. The process then repeats for the next time period.

2.1 Desired Mechanism Properties

We have stated that our goals are to guarantee individual rationality, efficiency, and strategy-proofness for the transplant centers. Individual rationality means that the economic utility from participating in the exchange mechanism is at least as high as the utility from not participating. As we are in a dynamic setting, we must

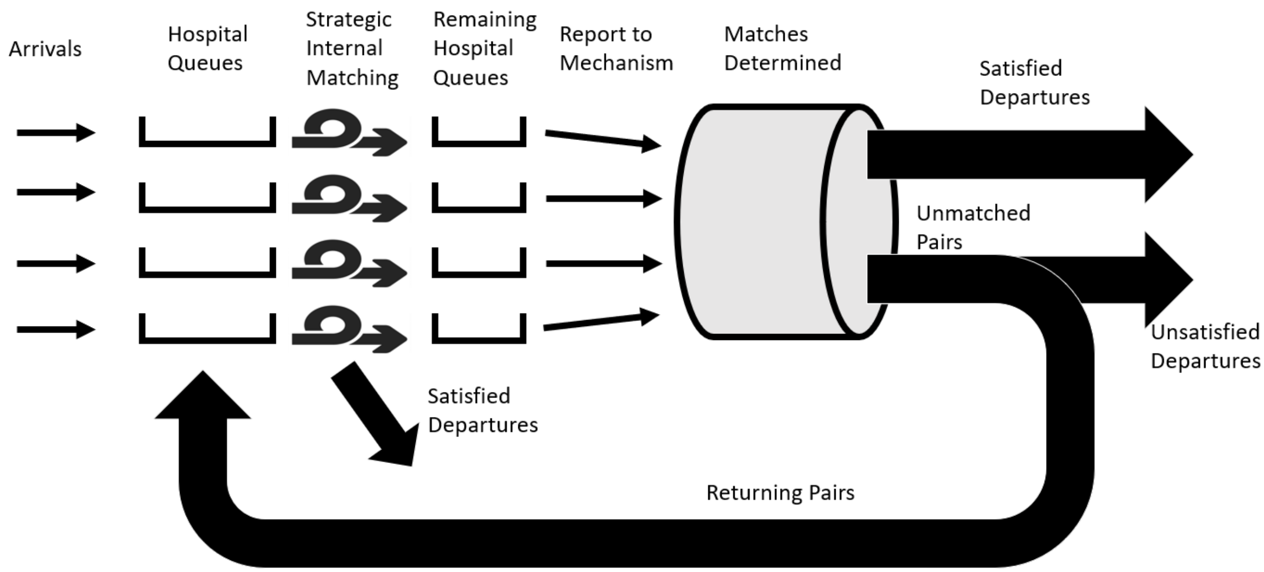

Figure 1: Queueing model of kidney exchange.

also define long-term versus immediate individual rationality. Long-term individual rationality means that participation has an equivalent or larger expected utility. Immediate individual rationality, however, is stricter in that it requires a transplant center to be guaranteed at least an equivalent utility in every individual period.

Efficiency is effectively equivalent to the total social welfare achieved by the solution. Our mechanism is efficient if it assigns kidneys so that it maximizes the global utility of the allocation; in other words, it maximizes the total number of matches that are produced in each period. A variation of efficiency that we may also consider is that of individually rational efficiency. This is a relevant measure of efficiency when we want to ensure immediate individual rationality, and it is therefore defined as the maximum possible matching in a period that still guarantees immediate individual rationality.

Our final criterion of interest is that of strategy-proofness. We define a mechanism to be strategy-proof if every center is no worse off by reporting all its pairs to the mechanism than reporting a subset of its pairs in every state of the world. When transplant centers do not reveal all their pairs to the model, total efficiency is compromised. Therefore, we seek to ensure no center has an incentive to misreport the pairs it receives.

2.2 Dynamic Credit Mechanism with Multi-Period Patients

Since unmatched patients have a non-zero probability of departing the model in each period, the credit mechanism must begin by modeling patient arrivals and departures at a given transplant center. Therefore, our mechanism performs the following steps: 1) a prediction of patient arrivals and departures, 2) an initial credit balance update, 3) a global matching, and 4) an end-of-match credit balance update as shown below in Figure 2.

2.2.1 Step 1: Predicting Patient Departures and Arrivals

At the very end of a matching, the mechanism knows each hospital has a set of unmatched pairs b_i. Then, patients arrive and depart until the next matching occurs. Until an unmatched pair departs, it is a reasonable assumption to treat them as taking up space in their assigned hospital. Therefore, although each hospital has a capacity K_i, an arriving pair may face a restricted capacity $\widetilde{K}_i(\tau) = K_i - b_i(\tau) - A_i(\tau)$, where τ is the time the pair arrives, $b_i(\tau)$ is the total number of unmatched pairs remaining in the hospital at time τ, and $A_i(\tau)$ is the total number of arrivals to the hospital that occur by time τ . It is reasonable to assume that all patients arrive and depart independently, and it is not unreasonable to believe that this process is independent of the amount of time since the last arrival or departure. We also assume that pairs are not allowed to depart before

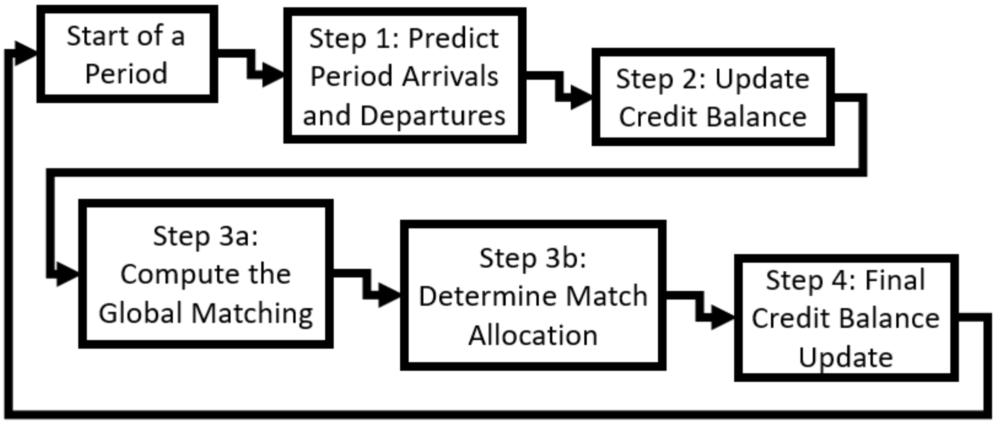

Figure 2: Mechanism process for one period.

their first matching attempt. This is reasonable as most hospitals will not add a patient to the list if the patient cannot wait until a matching is performed.

Therefore, our mechanism models both arrivals and departures as Poisson processes. Since arriving pairs may be blocked if $\widetilde{K}_i(\tau) = 0$, we suggest that the overall number of patients in a hospital $b_i(\tau) + A_i(\tau)$ can be modeled as a Markovian arrival process. Thus, we can develop an analytical representation of this process without excessive difficulty. We create a two-state Phase-type distribution to model this information, and present a sample transition rate matrix $\mathbf{Q}$ for a hospital where $K_i = 3$:

$$\mathbf{Q} = \left(\begin{array}{cccc|ccc|cc|c}
1 & 0 & 0 & 0 & & & & & & \\
\lambda & -\lambda & 0 & 0 & & \mathbf{0} & & & \mathbf{0} & \mathbf{0} \\
0 & \lambda & -\lambda & 0 & & & & & & \\
0 & 0 & \lambda & -\lambda & & & & & & \\
\hline
0 & \mu & 0 & 0 & -\mu & 0 & 0 & & & \\
0 & 0 & \mu & 0 & \lambda & -\lambda-\mu & 0 & & \mathbf{0} & \mathbf{0} \\
0 & 0 & 0 & \mu & 0 & \lambda & -\lambda-\mu & & & \\
\hline
 & \mathbf{0} & & & 0 & 2\mu & 0 & -2\mu & 0 & \mathbf{0} \\
 & & & & 0 & 0 & 2\mu & \lambda & -\lambda-2\mu & \\
\hline
 & \mathbf{0} & & & & \mathbf{0} & & 0 & 3\mu & -3\mu
\end{array}\right)
\begin{array}{c}
(0,3) \\ (0,2) \\ (0,1) \\ (0,0) \\ \hline (1,2) \\ (1,1) \\ (1,0) \\ \hline (2,1) \\ (2,0) \\ \hline (3,0)
\end{array}$$

From this matrix, we can see that we have defined μ to be an individual departure rate whereas λ is the overall arrival rate. We can also see our state space is defined as the pair (j,k), where j represents the number of patients who remain in the system from the previous matching period, and k is the number of new patients. The pattern of this matrix can be easily extended for hospitals that have a larger capacity. Using properties of Continuous Time Markov Chains (CTMC), we can derive the transition probability matrix $\mathbf{P}(\tau) = e^{\mathbf{Q}\tau}$, where $e^{\mathbf{X}}$ represents the matrix exponential. If we allow our rates to be defined in terms of the matching frequency (i.e., a matching is performed every $\tau = 1$ units), then the mechanism can compute the expected number of arrivals to a hospital given how many unmatched patients it had remaining at the end of the previous matching:

$$\mathbb{E}(A_i(1)|b_i(0) = \beta) = \sum_{j,k} k \cdot \mathbf{P}_{(\beta,0),(j,k)}(1).$$

2.2.2 Step 2: Initial Credit Balance Update

In the second step, the mechanism cycles through every transplant center and updates the center's credit balance based on the number of pairs reported to the clearinghouse. For a model with single period patients, a center's credits are increased at the beginning of a period by $\Delta c_i^{begin} = 4k_i(V_i^{disclosed} - k_i)$, where k_i is the average number of pairs that arrive at center i and $V_i^{disclosed}$ is the number of patients that the center reports to the model. The $4k_i$ term is a weighting factor that Hajaj et al. (2015) use to ensure strategy-proof-ness.

When multi-period patients are included, our discussion above leads us to consider a slightly more general definition of the credit increase. Namely, we let $\Delta c_i^{begin} = 4\mathbb{E}[A_i(1)|0](V_i^{disclosed} - \mathbb{E}[A_i(1)|b_i(0) = \beta])$. We have broken k_i into two separate definitions. $\mathbb{E}[A_i(1)|0]$ is the average number of pairs that arrive if a hospital starts out empty. $V_i^{disclosed}$ is now the number of new patients reported to the model, and $\mathbb{E}[A_i(1)|b_i(0) = \beta]$ is the average number of pairs that would be expected to arrive to a center given that the center had β patients remaining after the last matching. To aid in maintaining strategy-proof-ness, we note that $\mathbb{E}[A_i(1)|0] \geq \mathbb{E}[A_i(1)|b_i(0) = \beta]$. Therefore, we will not underestimate the required credit increase.

Under this formulation, we have extended the initial credit balance update that Hajaj et al. (2015) used to allow patients to remain for multiple periods. All information that the model uses is available and cannot be strategically misreported. We make the innocuous assumption that patients are reported with a unique identifier (such as Social Security Number) such that a hospital cannot report the same patient twice.

2.2.3 Step 3: Global and Subsequent Matchings

We constructed our matching mechanism based on the binary integer linear programming (BILP) techniques discussed by Anderson et al. (2015). Since Ausubel and Morrill (2014) showed that cycles need not be simultaneous to be incentive compatible, we have chosen not to limit the size of a chain or cycle. It is a relatively straightforward exercise to add these limits, if desired. In practice, these chains and cycles are usually limited in scope because larger groups have a higher chance of a failure occurring in one of the matches. This either causes the entire cycle or the remaining chain to collapse. We assume that offered matches do not fail. However, we have added tests for these collapses, to examine their effects on the system. Incorporating size restraints adds complication to the model but can only improve the results.

In each matching, we seek to maximize the weighted total number of matches subject to various market constraints. We have modeled altruistic donors as donors who have incompatible patients with type AB blood. This interpretation means that the altruistic "patient" is essentially a place holder for a bridge donation or a donation to the kidney waiting list. Initially, we find the maximum total number of exchanges possible. The constraints are based on compatibility matrix and flow conservation. Decision variables are X_{ij}, where $X_{ij} = 1$ if donor i gives to patient j and zero otherwise. Each decision variable must be less than or equal to the compatibility of donor i and patient j, C_{ij}. In a given period, once the set of pairs is disclosed by the transplant center, the mechanism constructs this compatibility matrix $\mathbf{C}$, where C_{ij} is a binary parameter equal to 1 if donor i is compatible with patient j, and 0 otherwise. The total flow out of donor i must be less than or equal to 1 since each donor can only give one kidney. The total flow into patient j must be less than or equal to 1 since each patient is only allowed to receive one kidney. Finally, the total flow into a patient must be equal to the total flow out of its respective donor. This is required as every donor has a corresponding patient and vice versa; this is true even for altruistic pairs. This is also required as a basic rule of conservation since every kidney must originate somewhere and end somewhere else and every pair must give a kidney to receive a kidney. w_{ij} is the weighting factor associated with a donation from donor i to patient j. In our model, $w_{ij} = 1, \quad \forall i, j$ because we seek to maximize the total number of matches. This is the simplest weighting scheme and most ethically straightforward, although it is possible to imagine a weighting scheme that gives priority to different patients. This could be due to many factors including age, health, amount of time waiting for a kidney, and difficulty of matching. We leave this weighting system for future research. The linear program formulation is a straight forward assignment problem as shown below.

$$
\begin{aligned}
\text{maximize:} \quad & \sum_i \sum_j w_{ij} X_{ij} & \\
\text{subject to:} \quad & \sum_j X_{ij} \leq 1 & \forall i \\
& \sum_i X_{ij} \leq 1 & \forall j \\
& \sum_i X_{ik} = \sum_j X_{kj} & \forall k \\
& X_{ij} \leq C_{ij} & \forall i, j
\end{aligned}
$$

Once the maximum number of matches is determined, this requirement is added to the constraints for the future matchings that are performed during the period. The algorithm then computes a random selection order for the centers, a process shown below in Figure 3. It runs two BILPs for each center. The first seeks to maximize the number of matches that the center receives, whereas the second seeks to minimize that number. In both cases, it is constrained by the global number of required matches and the number of matches assigned to other transplant centers. The mechanism then examines the center's credit balance. A non-negative balance is rewarded by giving the center its best matching possible. Negative credit balances are punished, and the center receives its worst possible match. As the algorithm determines how many matches a center will receive, this information is also added as a constraint.

2.2.4 Step 4: End-of-Match Credit Balance Update

After a center has been allocated a certain number of matches for the period, its credit balance is updated to reflect the spread in best and worst matches. Let u_i^+ be the number of matches center i receives in the best case, and u_i^- be the number it receives in the worst case ($u_i^+ \geq u_i^-$). If the initial credit balance is non-negative, the new credit balance is $c_i = c_i - (u_i^+ - u_i^-)$, whereas a negative initial credit balance leads to a new credit balance of $c_i = c_i + (u_i^+ - u_i^-)$.

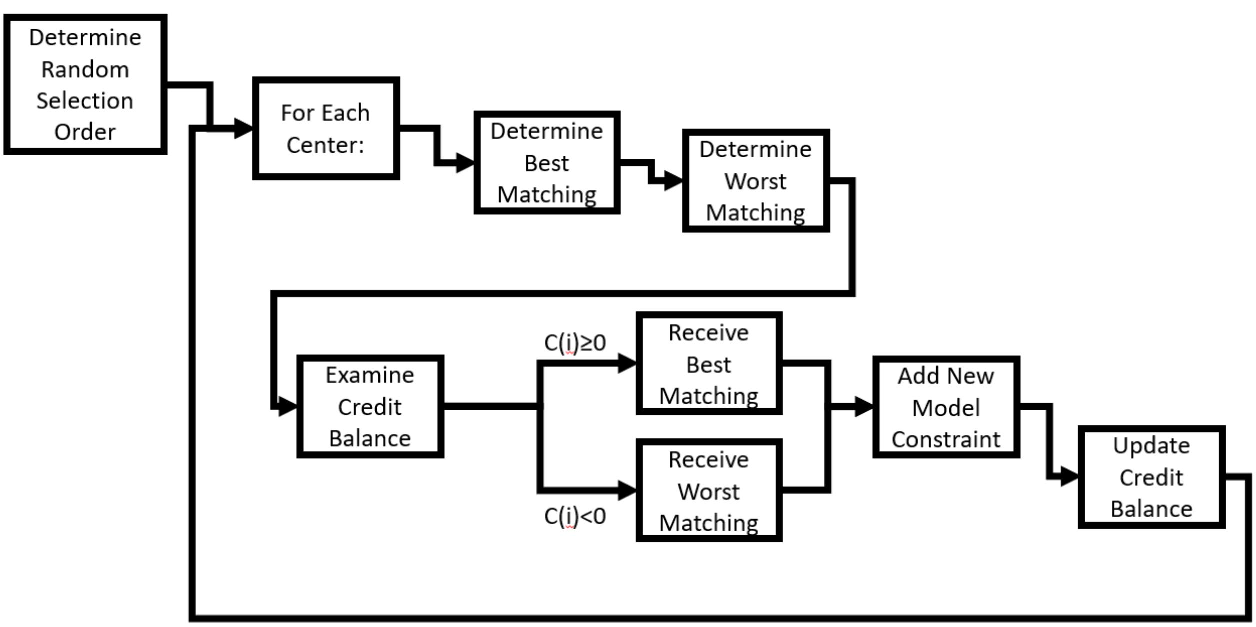

Figure 3: The match allocation process.

2.3 Mechanism Intuition and Reasonableness

Although we do not analytically prove that our mechanism is strategy-proof, we provide some brief intuition that we will build from in future work to prove its optimality. First, we note that the mechanism is indeed socially efficient because it produces the largest number of matches possible in every period. The mechanism is also long-run individually rational, because each center receives more matches than it would if it did not participate. Suppose a center performed an internal matching and then reported its unmatched patients to the exchange. In the long run for that strategic center, the probability of getting at least one match from the exchange is one. Therefore, it is individually rational for a strategic center to participate. If the mechanism is strategy-proof, then being truthful is more beneficial than being strategic, and therefore the mechanism is individually rational.

Hajaj et al. (2015) proved their mechanism is strategy-proof in the case with single-period patients. Our mechanism reduces to theirs if our patients only stay one period. Our mechanism also only rewards credits for new patients, but it does not punish hospitals for holding multi-period patients. We therefore believe that although our mechanism may not be perfectly strategy-proof, it is reasonably close, and we will demonstrate its practical appeal in the next section.

3 SIMULATION FRAMEWORK AND DESIGN

We construct a simulation model using the 2019b MATLAB release and Gurobi 8.1.0. We test this mechanism on several different environments, where our objective is to demonstrate this mechanism is strategy-proof. To do this, we construct a homogeneous environment, where all centers in a simulation run have the same capacity and are all either strategic or truthful. At the beginning of a period, strategic centers create a match internally and only report their unmatched pairs to the clearinghouse, whereas truthful centers report all their pairs to the clearinghouse. We assume that strategic centers have access to any of their patients who were unmatched in the previous period and can use these pairs when performing their internal matches. This is a realistic assumption as patients who are reported to a clearinghouse are free to leave the market at any time, and they will do so if their transplant center offers them an internal match.

We assume that 5% of the arriving pairs are altruistic donors, as Hajaj et al. (2015) do in their Saidman calculations. We also assume that these donors have a deterministic and finite time horizon of two periods, after which they will depart. This is based on the idea that altruistic donors have limited time in which to donate a kidney and a person's egotistical belief that if they are not used quickly, they must not be needed.

Based on an average of distributions from Roth et al. (2005a), Ashlagi and Roth (2014), Sönmez and Ünver (2014), and OPTN (2020), we assume a blood type distribution where 46% of individuals have type O blood, 37% have type A blood, 13% have type B blood, and 4% have type AB blood. In order to show worst-case performance, we assume that no compatible pairs are included in the model. Therefore, if a patient and corresponding donor have the same blood type, we assume they must have failed a tissue-type compatibility test. We treat the attempted pair arrival rate to a center as a Poisson distribution with an arrival rate $\lambda_i = \frac{K_i}{\alpha}$, where we let $\alpha = 2$. Thus, the mean number of attempted arrivals in a given period is half of the center's unrestricted capacity. This arrival rate could be adjusted to any distribution, but we have chosen the Poisson for simplicity, and tractability from both a simulation and analytical perspectives.

We test three experimental frameworks, as shown in Table 1; each formulation is then parameterized for the number of centers in the exchange, from $T = \{2, \ldots, 10\}$ and for the capacity of each center $K_i = \{5, 10, 20\}$. Each setup is then simulated with 50 replications of 60 time periods in each simulation run. We can thus view each simulation run as a five year time span, where the exchange assigns matches every month. Our three experimental frameworks are as follows. The first (Baseline Validation) is a validation of both our model and the results reported by Hajaj et al. (2015). In this case, all offered matches are successful, and the probability of departure parameter for unmatched pairs is 1, meaning that all pairs exist for only one period. The second framework (Primary Simulation) maintains the assumption that all matches are successful, but it changes the probability of departure parameter to 0.1. This parameter was arbitrarily chosen to show the mechanism is strategy-proof, and it suggests that unmatched patients will remain in the model for an average of about 10 periods. The final framework (Assumption Relaxation) also sets the probability of departure parameter to 0.1, but it relaxes the assumption that all offered matches are successful, instead assuming they are successful only 85% of the time. We include this framework to analyze the validity of a perfect matching assumption.

Table 1: Simulation frameworks and relevant parameters.

Simulation Framework	Baseline Validation	Primary Simulation	Assumption Relaxation
Probability of Unmatched Pair Departure	100%	10%	10%
Probability of Successful Matches	100%	100%	85%

4 NUMERICAL RESULTS

To determine the benefit of being truthful, we analyze the percentage differences between truthful and strategic centers for four criteria of interest. We compare the total number of transplants that occur in each period, and the percentage of all arrivals that are matched in each period. We expect that both measures should be larger for truthful centers than strategic ones. We also compare the number of periods that it takes for a patient to receive a kidney and the number of patients who leave the system unsatisfied, which should both be smaller for truthful centers. We also split all measurements by patient blood type. All graphical measurements are calculated as $\frac{Meas_{Truthful}}{Meas_{Strategic}} - 1$, and so represent percentage improvements; the 95% confidence intervals are included at each point. In addition, we have color-coded the figures so that a green background represents improvement for truthful centers, and a red background represents worse performance for truthful centers.

Figure 4's results confirm that truthful centers do better on average than strategic centers in the setup defined by Hajaj et al. (2015), thus showing that our simulation is a valid representation of the model we have described. We note that the average time to receive a kidney in this setting must be one period, regardless of whether a center is truthful or strategic. Therefore, there is no difference in the two for this measurement.

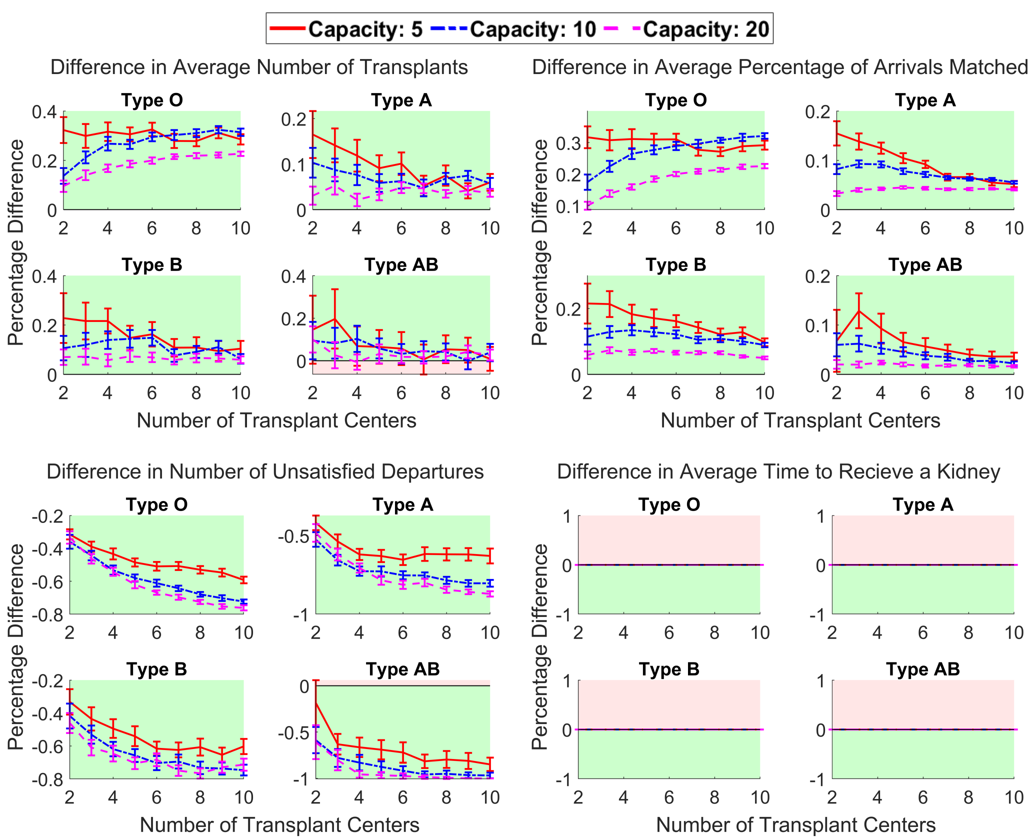

Figure 4: Baseline validation: departure probability= 100%, matched success probability= 100%.

Figure 5 shows that our mechanism provides significant incentives for transplant centers to be truthful. On average, truthful centers matched 10%-20% more pairs than strategic centers. They also matched a larger percentage of their arrivals, which was particularly beneficial for type O patients. Truthful centers also experienced about a 50% reduction in the number of unsatisfied departures from the system than strategic centers. Finally, patients spend about 10% less time waiting to receive a kidney when centers are truthful. Almost all confidence intervals are different from zero in the direction of improvement, suggesting this mechanism is indeed strategy-proof, individually rational, and efficient.

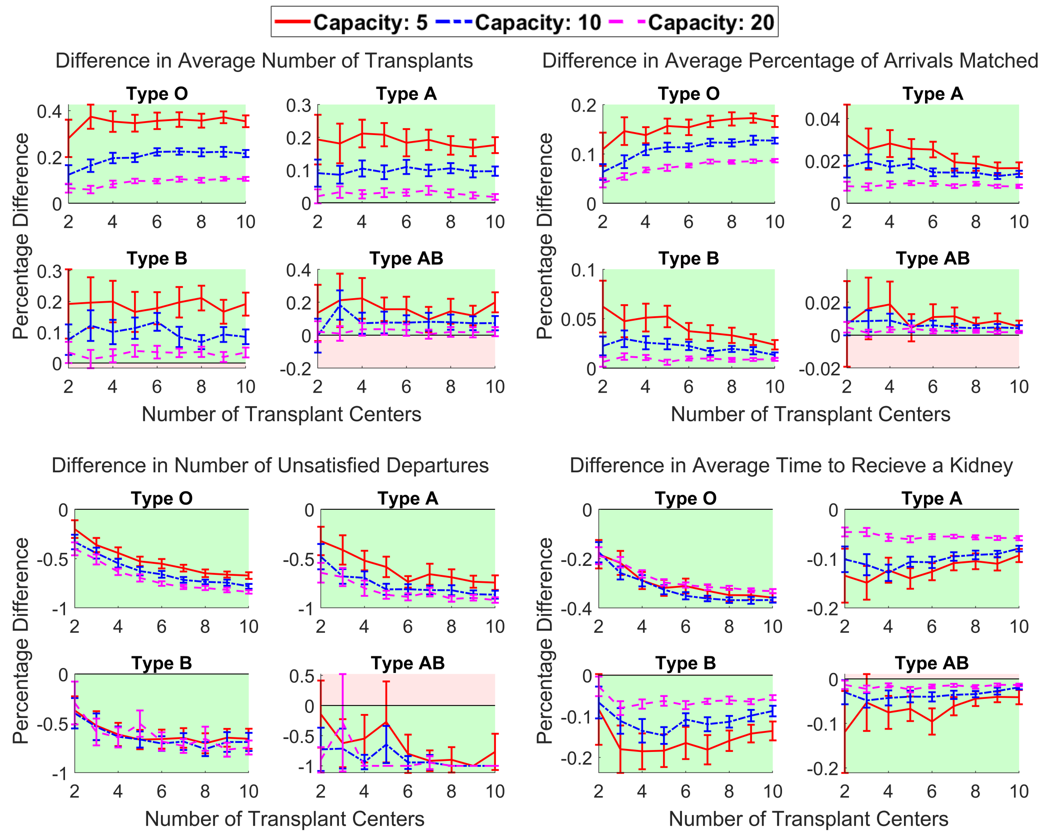

Figure 5: Primary simulation: departure probability=10%, matched success probability=100%.

When there is a chance that offered matches fail, as in Figure 6, very different simulation results appear. In this case, our results indicate that truthful centers actually receive fewer matches and match a smaller percentage of their arriving pairs. Truthful centers also have many more unsatisfied departures than strategic centers. These results indicate that assuming all matches are successful is critical to the success of our mechanism. One potential reason strategic centers could be performing better is because on average, smaller groups have a higher chance that all matches in the cycle or chain will be successful. Under our assumptions, the larger clearinghouse will tend to create larger chains and cycles which have a higher chance of failing. These results show that incorporating tissue type compatibilities and other cause of failure into the mechanism is an important direction for future work.

5 CONCLUSIONS AND FUTURE DIRECTIONS

We have created a credit based matching mechanism that properly incentivizes transplant centers to truthfully participate in a kidney exchange when patients have the possibility of existing for multiple periods. Our mechanism expands on the work done by Hajaj et al. (2015) to ensure that when patients remain in the system, truthfully reporting all arriving patients is strategy-proof, individually rational, and globally efficient. Using simulation, we have shown that our modifications to the original credit system do indeed fulfill these desired qualities, given the assumption that all offered matches are successful.

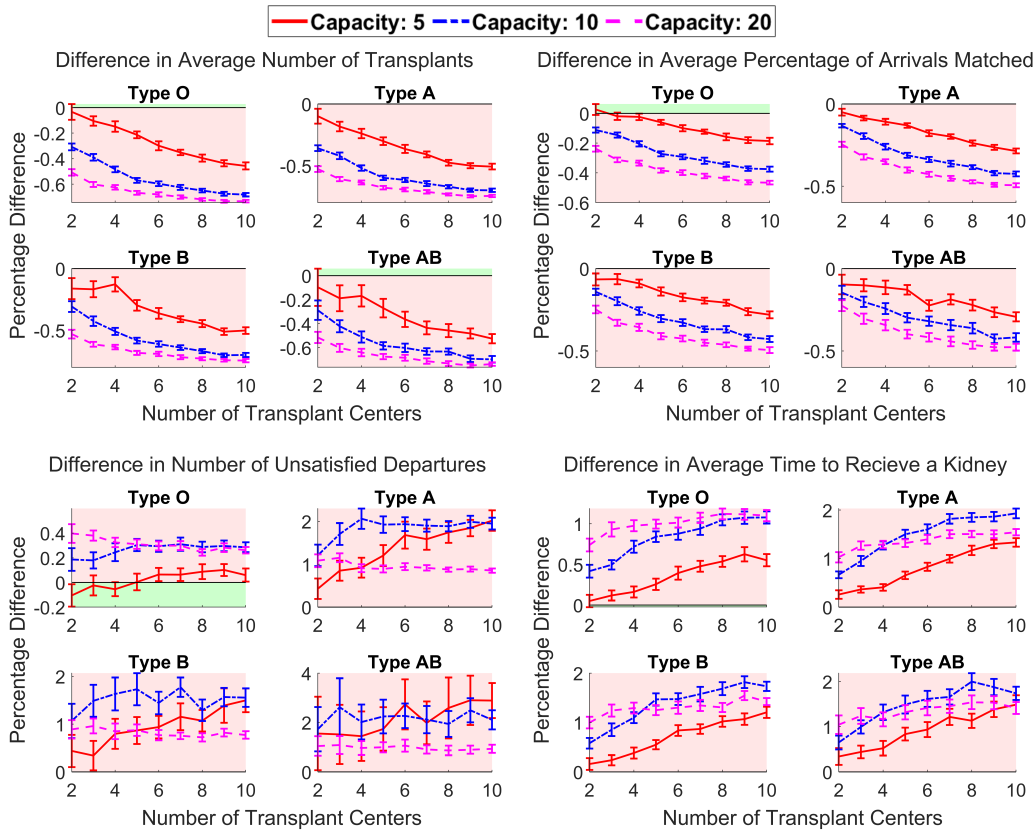

Figure 6: Assumption relaxation: departure probability=10%, matched success probability=85%.

A credit-based mechanism has great potential not only for the kidney exchange, but also for other matching markets. Our research can be expanded in many additional directions. Most importantly, we must derive rigorous proofs that this mechanism is indeed strategy-proof, individually rational, and efficient. Currently, this mechanism is globally efficient and individually rational in the long run. However, transplant centers may desire short run individual rationality, in which they always receive at least as many matches in a period as they would by conducting internal matches themselves. Future research should investigate how this mechanism performs under this new criterion.

As our simulations showed, we should also seek to find ways to relax the assumption of perfectly successful matches. Without any further research on the mechanism itself, there appears to be value in knowing the true match compatibility before matches are offered. We suggest further support should be given to the investigation of performing virtual cross matches and other compatibility factors. For the mechanism, there may be value in implementing a mechanism that borrows some ideas from the deferred acceptance algorithm. In our mechanism, the time between matchings is intentionally ambiguous for the purpose of generalization. If this time period is sufficiently large enough (a month or more), a match could initially be proposed, and tissue type compatibility could then be tested in a relatively short segment of time. Then, this information could be used to update a compatibility matrix and the matching could be performed. Intuitively, this would reduce the total number of failed matches in the final step since initial matches would have already been tested, which could lead to better outcomes for truthful centers and worse ones for strategic centers.

Additionally, more research should be done to test the robustness of the mechanism in nonhomogenous environments; that is, ones where centers have different capacities and not all follow the same strategy (truthful or strategic). As Hajaj et al. (2015) suggest, the mechanism should also be tested when individual centers have a probability of being strategic in a given period. We also suggest analyzing the model in the face of different arrival rate distributions, different blood type distributions, different weighting factors on the decision variables in the linear program, and real exchange data. Finally, future work should consider additional compatibility factors that may complicate the model.

Kidney exchanges have proven to be beneficial for many individuals suffering from kidney failure. As exchanges continue to grow, more individuals will benefit from the increased thickness of the market. By creating, and ultimately implementing better matching mechanisms, we will continue to increase the effectiveness of kidney exchanges.

ACKNOWLEDGMENTS

The author is extremely grateful to Dr. Yunan Liu for his support, encouragement, and guidance throughout the project.

REFERENCES

Anderson, R., I. Ashlagi, D. Gamarnik, and A. E. Roth. 2015. "Finding Long Chains in Kidney Exchange Using the Traveling Salesman Problem". *Proceedings of the National Academy of Sciences* 112(3):663–668.

Ashlagi, I., and A. E. Roth. 2014. "Free Riding and Participation in Large Scale, Multi-Hospital Kidney Exchange: Free Riding". *Theoretical Economics* 9(3):817–863.

Ausubel, L. M., and T. Morrill. 2014. "Sequential Kidney Exchange". *American Economic Journal: Microeconomics* 6(3):265–285.

Gentry, S. E., D. L. Segev, M. Simmerling, and R. A. Montgomery. 2007. "Expanding Kidney Paired Donation Through Participation by Compatible Pairs". *American Journal of Transplantation* 7(10):2361–2370.

Hajaj, C., J. P. Dickerson, A. Hassidim, T. Sandholm, and D. Sarne. 2015. "Strategy-Proof and Efficient Kidney Exchange Using a Credit Mechanism". In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, edited by B. Bonet and S. Koenig, 921–928. Palo Alto, California: Association for the Advancement of Artificial Intelligence Press.

OPTN 2020. "Organ Procurement and Transplantation Network". https://optn.transplant.hrsa.gov/, accessed 29th July.

Roth, A. E., T. Sönmez, and M. U. Ünver. 2004. "Kidney Exchange". *The Quarterly Journal of Economics* 119(2):457–488.

Roth, A. E., T. Sönmez, and M. U. Ünver. 2005a. "A Kidney Exchange Clearinghouse in New England". *American Economic Review* 95(2):376–380.

Roth, A. E., T. Sönmez, and M. U. Ünver. 2005b. "Pairwise Kidney Exchange". *Journal of Economic Theory* 125(2):151–188.

Roth, A. E., T. Sönmez, and M. U. Ünver. 2007. "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences". *American Economic Review* 97(3):828–851.

Santos, N., P. Tubertini, A. Viana, and J. P. Pedroso. 2017. "Kidney Exchange Simulation and Optimization". *Journal of the Operational Research Society* 68(12):1521–1532.

Sönmez, T., and M. U. Ünver. 2014. "Altruistically Unbalanced Kidney Exchange". *Journal of Economic Theory* 152:105–129.

AUTHOR BIOGRAPHIES

KURTIS KONRAD is a PhD student in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. He holds a masters degree in Industrial and Systems Engineering from North Carolina State University. His research interest is in stochastic modeling, simulation, and online learning, with interests in applying these methodologies to health care, transportation networks, and other service systems. His email address is kekonrad@ncsu.edu.