

Robust Dynamic Optimization with Application to Kidney Exchange

(Doctoral Consortium)

John P. Dickerson
Computer Science Department, Carnegie Mellon University
Pittsburgh, PA, USA
dickerson@cs.cmu.edu

ABSTRACT

Kidney exchange provides a life-saving alternative to long waiting lists for patients in need of a new kidney. We derive a variety of mathematical models for the kidney exchange optimization problem, where the general goal is to maximize some form of social welfare vis-à-vis transplanting kidneys. We explore the implications of making the optimization problem dynamic (considering the future evolution of the exchange pool when optimizing now), failure-aware (where possible post-algorithmic match failures are accounted for), and fairness-aware (losing overall efficiency at the cost of a more balanced matching). Our goal is to provide an empirically grounded framework that combines each of these dimensions in a theoretically sound way. We support our models with real results from one of the largest fielded kidney exchanges in the world.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems;
J.4 [Social and Behavioral Sciences]: Economics

Keywords

Kidney exchange, stochastic optimization, equity vs. efficiency

1. INTRODUCTION

The preferred treatment for kidney failure is transplantation; however, the demand for donor kidneys is far greater than supply. For example, 34,837 people were added to the US national waiting list in 2012, while only 15,938 left it due to receiving a kidney [11]. Demand is increasing worldwide.

Complementing standard deceased and living donation is kidney exchange, which allows patients with a willing but medically incompatible living donor to swap their donor with other patients. Roth, Sönmez, and Ünver [10] were the first to address the economics of large-scale kidney exchange, motivating the need to consider properties like efficiency, incentive compatibility, and individual rationality when designing and fielding exchange mechanisms. Since then, a number of multi-hospital kidney exchanges have been fielded, prompting other economists, computer scientists, and medical professionals to study the kidney exchange problem.

We look at kidney exchange primarily through the lens of computational economics, but inform our research via collaboration with the United Network for Organ Sharing (UNOS) National Kid-

ney Paired Donation (KPD) Pilot Program, which includes 133 hospitals, or roughly 58% of all transplant centers, in the United States.

2. TOWARD ROBUST DYNAMIC ORGAN EXCHANGE

From a computational point of view, one can view an n -patient kidney exchange as a directed compatibility graph $G = (V, E)$. Here, vertices represent patient-donor pairs, and a directed edge e exists from vertex v_i to v_j if the patient belonging to v_j is compatible with the donor at v_i . A donor is willing to give her kidney if and only if her associated patient also receives a kidney; thus, we are interested in forming cycles c in the graph G , where each vertex in c obtains the kidney of the previous vertex.

In practice, cycles are capped in length by some small constant L (for example, the UNOS exchange uses $L = 3$). This cap is due to all transplants in a cycle necessarily being performed simultaneously so that no donor backs out after his patient has received a kidney but before he has donated his kidney. A recent innovation relaxes this constraint. Chains are initiated by a non-directed (altruistic) donor with no paired patient donating his kidney to a patient, whose paired donor donates her kidney to a patient, and so on. Transplants in chains need not be executed simultaneously; if a donor backs out after his paired patient receives a kidney, the chain is unfortunately broken but no remaining pair is harmed.

A matching M is a collection of vertex-disjoint cycles and chains in the graph G . Note that the elements of the matching must be disjoint because no donor can give more than one of his kidneys. Then, given the set of all legal matchings \mathcal{M} , the general *clearing problem* is to find some matching $M^* \in \mathcal{M}$ that maximizes a utility function $u : \mathcal{M} \rightarrow \mathbb{R}$. Formally:

$$M^* = \operatorname{argmax}_{M \in \mathcal{M}} u(M)$$

Abraham, Blum, and Sandholm [1] created the first scalable algorithm for solving a basic (utilitarian, static, deterministic, no long chains) version of this problem. Their algorithm leverages integer programming, specifically branch-and-price, a technique that proves optimality by incrementally generating only a small part of the model during tree search. Our work significantly extends both the formulation they considered and the algorithmic techniques used to find a matching that is provably or approximately optimal.

2.1 New dimensions in kidney exchange

Translating a real-world problem into an abstract mathematical model inevitably results in some loss of correspondence between the final computed solution and whatever is implemented in reality. We explore a variety of generalizations of the standard kidney exchange model that reduce this loss in expressiveness; we quantify

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

the theoretical behavior of each generalization and validate them empirically on generated and real data from the UNOS exchange.

2.1.1 Considering the future

Fielded exchanges currently match myopically, maximizing the number of patients who get kidneys in an offline fashion at each time period (e.g., weekly in the UNOS exchange). This is sub-optimal; the clearing problem is dynamic since patients and donors appear and expire over time. Thus, the matching algorithm should take distributional information about possible futures into account in deciding what action to take now. This is typically done by drawing sample trajectories of possible futures at each time period, but may require a prohibitively large number of trajectories or prohibitive memory and computation to decide what action to take.

To counter the computational complexity of full dynamic optimization, we proposed to learn *potentials* of elements (e.g., vertices) of the current problem [4]. After learning potential, at run time we must only run an *offline* matching algorithm at each time period, but subtracting out in the objective the potentials of the elements used up in the matching. We theoretically compared the power of using potentials on increasingly large elements, and empirically showed significant gains over myopic matching (while scaling to much larger graphs than can be handled by other dynamic optimization techniques).

2.1.2 Failure-aware kidney exchange

Successful transplantation of a kidney relies on tissue-type compatibility between the donor organ and patient, among other medical and logistical factors. This is determined through a tissue-type *crossmatch* between a donor and patient's blood; if the two types differ substantially, the patient's body will reject the donor organ.

For a variety of practical reasons, including crossmatch failures, most planned kidney exchange transplants do not go to transplant in reality. In [6], we analyzed the inclusion of probabilistic match failures in standard kidney exchange models. We address two presently unsolved problems in kidney exchange: first, how the efficacy of altruist-initiated donor chains changes as chain length increases (which we also considered theoretically and empirically in a deterministic model in [5]), and second, how to match robustly in the face of post-algorithmic match transplant failure.

We showed that failure-aware kidney exchange can significantly increase the expected number of lives saved in theory, on random graph models; on real data from UNOS kidney exchange match runs; and on synthetic data generated via a model of dynamic kidney exchange [6]. We also designed a branch-and-price-based optimal clearing algorithm for the probabilistic exchange clearing problem and showed that this new solver scales well on large simulated data, unlike prior clearing algorithms.

2.1.3 Balancing efficiency and equity

Some patients are *highly-sensitized*; there is a very low probability that their blood will pass a crossmatch test with a random organ. For these patients, finding a kidney is quite difficult (and median time on the waiting list jumps by a factor of three over less sensitized patients [11]). Roughly 17% of the adult patients on the waiting list for *deceased* donor kidneys are highly-sensitized [8]. Conversely, the percentage of highly-sensitized patients in fielded kidney exchanges is quite high; roughly 60% of the UNOS nationwide kidney exchange is highly-sensitized.

Motivated by these highly-sensitized pools, we proposed two natural criteria for balancing fairness and efficiency in kidney exchange [7]. We performed a preliminary theoretical analysis of the *price of fairness*—the relative loss in efficiency due to considering fairness—in dense kidney exchange graphs, and showed empirically that these results do not align with real data.

2.2 General robust dynamic optimization

Our completed work has largely addressed each of these outstanding issues in fielded exchanges *independently*. Moving forward, the kidney exchange community would benefit immensely from *combined* approaches to handling not just dynamic matching, match failures, and fairness in the optimization problem, but also game-theoretic and legal considerations in the design of the matching mechanism itself. We plan to draw on previous work from the operations research and economics literature to move in this direction: for example, Hooker and Williams present a general methodology for balancing a particular form of fairness (that we feel would not be the criterion of choice in kidney exchange) and efficiency [9]; Bertsimas, Farias, and Trichakis formalized a proposal for balancing fairness and efficiency in *deceased* organ allocation [3]; and Ashlagi, Jaillet, and Manshadi theoretically address dynamic exchange in a reduced model [2]. A general parameterized model of kidney exchange will increase the efficacy of fielded exchange and aid in the widespread adoption of new exchanges in differing legal and political environments.

3. ACKNOWLEDGMENTS

This research was supported by the National Science Foundation under grants IIS-0905390, IIS-0964579, and CCF-1101668 and by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575 and the Blacklight system at the Pittsburgh Supercomputing Center (PSC). The author acknowledges Intel Corporation and IBM for gifts. The author also thanks the UNOS KPD Pilot Program staff.

4. REFERENCES

- [1] D. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *ACM Conference on Electronic Commerce (EC)*, pages 295–304, 2007.
- [2] I. Ashlagi, P. Jaillet, and V. H. Manshadi. Kidney exchange in dynamic sparse heterogeneous pools. In *ACM Conference on Electronic Commerce (EC)*, pages 25–26, 2013.
- [3] D. Bertsimas, V. F. Farias, and N. Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.
- [4] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Dynamic matching via weighted myopia with application to kidney exchange. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1340–1346, 2012.
- [5] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Optimizing kidney exchange with transplant chains: Theory and reality. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 711–718, 2012.
- [6] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Failure-aware kidney exchange. In *ACM Conference on Electronic Commerce (EC)*, pages 323–340, 2013.
- [7] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Price of fairness in kidney exchange. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2014.
- [8] HHS/HRSA/HSB/DOT. OPTN/SRTR annual data report, 2011.
- [9] J. N. Hooker and H. P. Williams. Combining equity and utilitarianism in a mathematical programming model. *Management Science*, 58(9):1682–1693, 2012.
- [10] A. Roth, T. Sönmez, and U. Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119(2):457–488, 2004.
- [11] United Network for Organ Sharing (UNOS). <http://www.unos.org/>.