18.204: STRATEGIES FOR PAIRWISE KIDNEY EXCHANGE

GREG LUBIN

Abstract. Because kidney donors are often not compatible with their intended patients, exchanges can be arranged between one incompatible patient-donor pair and another. Given a set of constraints, there are certain strategies by which we can maximize the efficiency of this process among a pool of patient-donor pairs while incentivizing hospitals to submit truthful data. I will present a randomized matching mechanism which has an approximation ratio of 3/2 to the maximum cardinality matching, and is also expected to motivate truthfulness. This is an improvement over previous mechanisms, and is more efficient than any possible deterministic truthful mechanism.

1. Introduction

There are currently 100,791 kidney patients in U.S. on the waiting list to receive a cadaver kidney for a life-saving transplant. Median waiting time is 3.6 years, and in 2014, 4,761 patients died while waiting for a kidney transplant. Since there aren't nearly enough cadaver kidneys available to fulfill the need for transplants, programs have been launched in the US to facilitate exchanges with live donors. Since healthy people can survive with only one kidney, these programs connect kidney patients with donors who are willing to donate one of their kidneys. However, a healthy donor who would like to give his kidney to a specific patient is often not able to do so because of immunological or blood type incompatibilities. National kidney exchange programs aim to implement exchanges between two incompatible patientdonor pairs u and v so that the donor of pair u donates her kidney to the patient of pair v and vice versa. This requires four simultaneous operations between the two pairs, which is why arranging exchanges involving more pairs would be far more logistically difficult.

In a pairwise kidney exchange program, each incompatible patient-donor pair would enroll through their hospital in the hopes of being matched with another pair that is compatible for an exchange. Ideally, each hospital would truthfully report all of its patient-donor pairs to a national system, and then an algorithm would process the compatibility data to produce a matching between as many pairs as possible. However, there are certain reasons why hospitals might not be truthful in reporting all of their patients to the system, since each hospital's incentives are to maximize the number of their patient-donor pairs that are matched. If a hospital expects that they can increase their number of matched patients by matching some of them internally, then they will not submit them to the central matching algorithm. This might then cause fewer patients overall to be matched, since the algorithm is not able to find the maximum matching over all the patients.

I will go over some of the previously attempted truthful algorithms (referred to as *mechanisms*), and introduce a recently designed randomized matching algorithm that is truthful in expectation and achieves a greater number of matched patients than previous algorithms. In this problem, the objective is to maximize the number of total kidney exchanges while remaining $truthful$, i.e. incentivizing hospitals to report all their patient-donor pairs to the system. Thus, any truthful mechanism must return a number of matches for each hospital that is expected to be greater than if they matched any of their patients internally.

2. Problem and constraints

I will model the matching problem through representing it as a graph, in which each node represents a patient-donor pair. In the initial compatibility graph, each edge between two nodes means that those two pairs are compatible for an exchange. After a matching mechanism is applied to the graph, an edge between two nodes in the matching graph will mean that these two pairs were matched by the algorithm. Thus, the matching graph will be a set of disjoint edges, in which each node shares an edge with at most one other node, and some nodes will have a degree 0 if they were not successfully matched with another pair. I will refer to the *gain* of a hospital as the number of nodes from that hospital that are matched, i.e. the number of nodes that are incident to edges in the matching graph.

It is apparent that the optimal solution just involves performing a maximal matching computation on a compatibility graph containing all the nodes. This mechanism is unfortunately not truthful however, because hospitals might choose to match some of their nodes internally in order to maximize their gain. In these cases, the gain of the hospital would be the number of internally matched patients plus the patients that it submitted to the system and were matched by the mechanisms, which on average would be greater than their gain from submitting all their nodes to the maximum cardinality mechanisms. A mechanism is only truthful if the gain for each hospital in all cases is greater when they report all of their nodes to the graph. In the case of a randomized mechanism, it is truthful in expectation when the expected gain for each hospital will be maximized when they submit all of their nodes to the graph. Thus, we want to design another mechanism that will maximize cardinality of the overall graph while assuring that no hospital would have the incentive to hide any of their nodes. The efficiency of a truthful mechanism is evaluated through its *approximation ratio*, which is the maximum ratio over all possible situations of: size of the maximum cardinality matching / expected size of the matching returned by the mechanism.

3. Previous Matching Algorithms

3.1. **Maximum Cardinality Matching.** As mentioned above, it is clear that the maximum cardinality matching is not truthful, despite returning the most possible edges in the matching graph.

FIGURE 1. Example of maximum cardinality matching

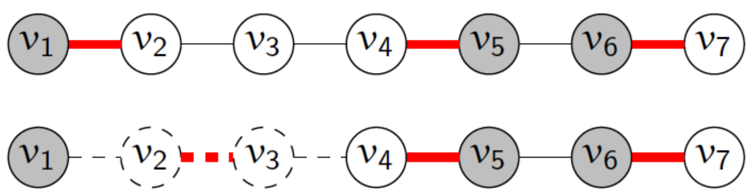

Theorem 3.1. *Maximum cardinality matching is not truthful.*

Proof. In the graph returned by the maximum cardinality matching (Fig. 1, graph 1), both Hospital 1 (grey) and Hospital 2 (white) have a gain of 3. However, if Hospital 2 matches two of their nodes internally (Fig. 1, graph 2) instead of submitting them to the matching algorithm, then their gain increases to 4. Thus, the maximum cardinality matching is not truthful, since there is at least one scenario in which gain(withholding patients) > gain(reporting all patients) for one of the hospitals. □

3.2. **Mix-and-Match.** After it became clear that maximum cardinality matching was not truthful because it didn't align with the incentives

Figure 2. Example of Mix-and-Match matching

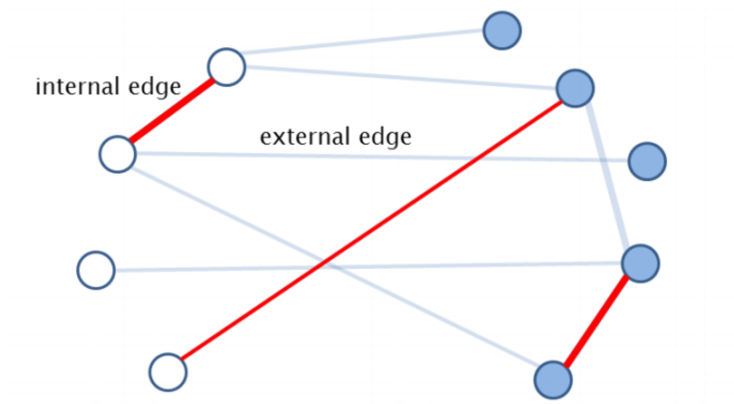

of hospitals, Ashlagi et. al presented a new deterministic mechanism called Mix-and-Match that was universally truthful.

This algorithm first determines the matchings that maximize the number of internal edges for each hospital, and then breaks ties by returning the matching with maximum cardinality. In other words, Mix-and-Match selects the matching with maximum number of external edges from the set of matchings with maximum number of internal edges (see Fig. 2). While Mix-and-Match is truthful because it stimulates the hospitals' incentives to report all their nodes to the system, it has an approximation ratio of 2, meaning that it could produce as few as 1/2 as many edges as the maximum cardinality matching. Thus, it is not a satisfactory solution to the problem, and we would like to find a more efficient truthful mechanism.

4. Flip-and-Match

In an attempt to create a mechanism that is truthful in expectation while being more efficient than Mix-and-Match, Ashlagi et. al then presented a new theoretical mechanism titled Flip-and-Match. In contrast to the deterministic algorithms of the maximum cardinality matching and Mix-and-Match, this mechanism is randomized, which means that it runs a probability distribution over multiple deterministic mechanisms. Specifically, Flip-and-Match selects with equal probability between the matching returned by Mix-and-Match and the maximum cardinality matching. From analyzing the results of this algorithm, they prove that this mechanism has an approximation ratio of 4/3, meaning that it is far more efficient than any truthful deterministic mechanism. However, this question still remained: is this mechanism truthful in expectation? Ashlagi et. al conjectured that it was, but I will demonstrate that FLIP-AND-MATCH is in fact not truthful.

FIGURE 3. Examples of FLIP-AND-MATCH matching discussed in proof

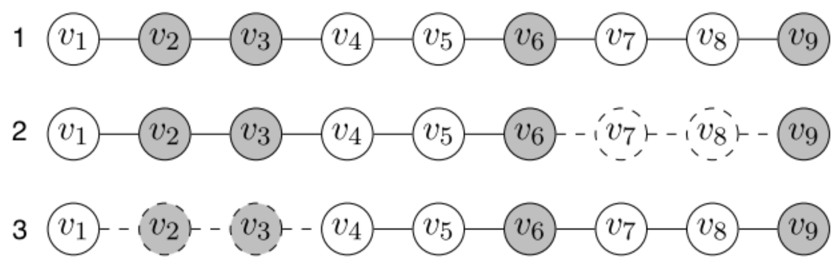

Theorem 4.1. FLIP-AND-MATCH *is not truthful.*

Proof. We will prove that FLIP-AND-MATCH is not truthful in expectation by using the three example graphs shown in Fig. 3. When applied to the initial compatibility graph (Fig. 3-1), MIX-AND-MATCH returns the matching M1 = (v2, v3),(v4, v5),(v7, v8). The gain of Hospital 1 (white) is 4 while the gain of Hospital 2 (grey) is 2. A maximum cardinality matching M2 would match all nodes except one, and the unmatched node could be either grey or white. In other words, M2 matches either 4 white nodes and 4 gray nodes or 5 white nodes and 3 gray nodes.

We distinguish between these two cases:

Case 1: In this maximum cardinality matching, M2 matches 4 white nodes and 4 gray nodes, so the expected gain of Hospital 1 from the application of FLIP-AND-MATCH on Graph 1 is 4. Consider Graph 2 in which Hospital 1 hides the white nodes v7 and v8 and matches them internally. In the new graph, MIX-AND-MATCH returns the matching (v2, v3),(v4, v5) and contains 2 matched white nodes while the maximum cardinality matching is (v1, v2),(v3, v4),(v5, v6) that contains 3 matched white nodes. The expected gain of Hospital 1 (including the hidden nodes) is 4.5.

Case 2: M2 matches 5 white nodes and 3 gray nodes and hence the expected gain of Hospital 2 from the application of FLIP-AND-MATCH to the original Graph 1 is 2.5. Consider Graph 3 in which

Hospital 2 hides nodes v2 and v3 (and matches them internally). In the new instance, MIX-AND-MATCH returns the matching (v4, v5),(v7, v8) that contains no matched gray nodes while the maximum cardinality matching is (v4, v5),(v6, v7),(v8, v9) that contains 2 matched gray nodes. The expected gain of Hospital 2 (including the hidden nodes) is 3. In both cases, one of the hospitals has an incentive to deviate from truth-telling and withhold nodes. □

5. WEIGHT-AND-MATCH

In order to address the shortcomings of MIX-AND-MATCH and FLIP-AND-MATCH, Caragiannis et al. presented a new mechanism called WEIGHT-AND-MATCH. Like FLIP-AND-MATCH, this is a randomized mechanism that selects with equal probability between two deterministic mechanisms. However, this one makes several distinct changes to ensure that the mechanism is truthful in expectation.

FIGURE 4. Weights assigned in WEIGHT-AND-MATCH

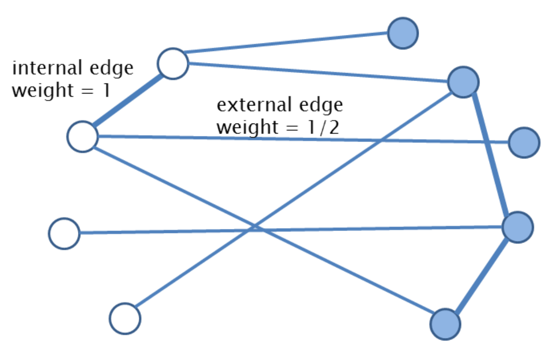

It starts out by assigning weights to the edges, so that internal edges are assigned weight = 1 while external edges are assigned weight = 1/2 (see Fig. 4). The mechanism then calculates all maximum-weight matchings on the weighted graph, and finds the maximum-cardinality and the minimum-cardinality graphs out of this set (see Fig. 5). Finally, WEIGHT-AND-MATCH selects equiprobably between the maximum-weight maximum-cardinality matching and the maximum-weight minimum-cardinality matching. In their paper, Caragiannis et al. demonstrated that this mechanism is truthful in expectation, achieves an approximation ratio of 3/2, and runs in polynomial time. I will go through a proof that WEIGHT-AND-MATCH is truthful, unlike FLIP-AND-MATCH.

FIGURE 5. Two maximum-weight matchings in WEIGHT-AND-MATCH (red and green)

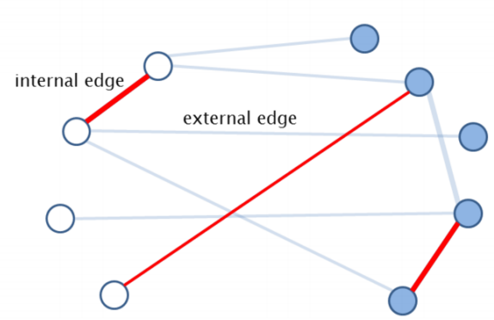

Lemma 5.1. WEIGHT-AND-MATCH *is truthful in expectation.*

Proof. In order to prove that WEIGHT-AND-MATCH is truthful, we need to demonstrate that no hospital will ever have the incentive to be untruthful under this mechanism. Let M1 and M2 be the two matchings returned by the mechanism on initial compatibility graph G. Assume that Hospital 1 hides some nodes and matches them internally. Let G' be the induced graph that does not contain the nodes that the hospital hid, nor the edges between them. The mechanism returns two matchings on input graph G', which we will call M3 and M4, representing those two matchings augmented by the edges Hospital 1 uses to match the hidden nodes internally. In order to see whether either hospital would ever have an incentive to deviate from being truthful, we will compare M1 to M3 and M2 to M4. We will denote by gain(M) the gain of Hospital 1 from matching M and by wgt(M) the weight of matching M. Our proof will follow from the next two lemmas. □

FIGURE 6. Example demonstrating $n_{gg}(M1) = n_{gg}(M3)$

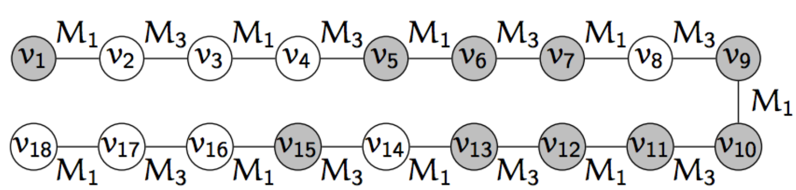

Lemma 5.2. *gain(M3) = gain(M1) - 2(wgt(M1) - wgt(M3)).*

Proof. We will denote by $n_{ww}(M), n_{wg}(M)$, and $n_{gg}(M)$ the number of edges in matching M connecting two white nodes, two nodes belonging to different hospitals, and two gray nodes, respectively. We will first show that $n_{gg}(M1) = n_{gg}(M3)$. Consider the quantities $n_{ww}(M), n_{wg}(M)$ and $n_{gg}(M)$. In addition , consider the symmetric difference $M1\Delta M3 = (M13) \cup (M31)$. We will call the connected components C, which could be either either cycles or paths. It suffices to prove $gain(C3) = gain(C1) - 2(wgt(C1) - wgt(C3))$ for each component. It holds that: $gain(M) = 2n_{ww}(M) + n_{wg}(M)$
$wgt(M) = n_{ww}(M) + 1/2n_{wg}(M) + n_{gg}(M)$
We prove that: $n_{gg}(M1) = n_{gg}(M3)$. C contains a block of t consecutive gray nodes $b_1, ..., b_t$. The number t of nodes in the connected component cannot be even, for the following reason: Assume that t is even: M1 contains $t/2-1$ edges $((b_2, b_3), (b_4, b_5), ..., (b_{t-2}, b_{t-1}))$ and M3 contains t/2 edges $((b_1, b_2), (b_3, b_4), ..., (b_{t-1}, b_t))$. By replacing the edges of M1 with those of M3 we acquire a matching with greater weight or equal weight and minimum cardinality. Summing over all components we prove $n_{gg}(M1) = n_{gg}(M3)$.
$gain(M3) = 2n_{ww}(M3) + n_{wg}(M3)$
$= 2n_{ww}(M3) + n_{wg}(M3) + 2n_{gg}(M3) - 2n_{gg}(M1)$
$= gain(M1) - 2(wgt(M1) - wgt(M3))$. □

Figure 7. Examples of connected components of $M2\Delta M4$

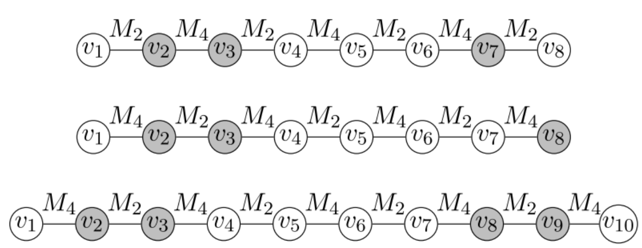

Lemma 5.3. $gain(M4) \leq gain(M2) + 2(wgt(M2) - wgt(M4))$.

Proof. Consider the quantities $n_{ww}(M)$, $n_{wg}(M)$ and $n_{gg}(M)$, and consider the symmetric difference $M2\Delta M4 = (M24) \cup (M42)$. Given that M2 is a maximum-weight matching in G, it holds that $wgt(C2) \geq wgt(C4)$. We now classify connected components C of subgraph G induced by $M_2\Delta M_4$ that are paths into different types: the first and last letters are w or g and denote whether the left and right endpoint of the connected component is a white or gray node, respectively. The second and third numbers are either 2 or 4 and denote whether the first and the last edge of the connected component belong to matching M2 or M4, respectively (see examples of w22w, w44g, and w44w in Fig.7). We will prove the inequality through case analysis.

Case 1. If C is a cycle, or a path of type w22w, w24w, w42w, w22g, w24g, g22g, g24g, g42g, or g44g, we have $gain(C4) \leq gain(C2)$ and the inequality follows easily since wgt(C2) $\geq wgt(C4)$.

Case 2. If C is a path of type w42g or w44g, we claim that $wgt(C2)+wgt(C4)$ is non-integer. Since the first and the last node in the path belong to different hospitals, there is an odd number of external edges in C, and each such edge contributes 1/2 to the sum $wgt(C2)+wgt(C4)$. Since $wgt(C2) \geq wgt(C4)$, we know that $wgt(C2) - wgt(C4) \geq 1/2$. The original inequality then follows by observing that $gain(C2) = gain(C4) - 1$ in this case.

Case 3. If C is of type w44w, observe that C4 contains one more edge than C2 and, hence, $wgt(C2) > wgt(C4)$. Also, we know that the number of external edges in C is even, and that wgt(C2) + wgt(C4) is an integer. Thus, we also know that $wgt(C2) \geq wgt(C4) + 1$. Then, the original follows by further observing that $gain(C2) = gain(C4)-2$.

Since $wgt(M1) = wgt(M2)$ and $wgt(M3) = wgt(M4)$, by Lemmas 5.2 and 5.3 we have that the expected gain $1/2(gain(M3) + gain(M4))$ of Hospital 1 when it hides some white nodes and matches them internally is upper-bounded by the expected gain $1/2(gain(M1)+gain(M2))$ when it acts truthfully. Thus, since hiding nodes will never result in an expected gain that is greater than the one that Weight-and-Match produces, we conclude that this matching is truthful in expectation. □

6. Discussion and Future Directions

The Weight-and-Match mechanism discussed in my paper is a promising approach to solving the pairwise kidney matching problem, and is more efficient than previous truthful mechanisms. In their paper, Caragiannis et al. perform additional analysis to determine upper and lower bounds of the approximation ratio that a randomized truthful matching can achieve, along with proving the approximation ratio of Weight-and-Match and the fact that it can run in polynomial time. The results in their paper will likely guide future work in finding an optimal mechanism for pairwise kidney exchanges among larger numbers of hospitals. This problem is an interesting example of applied graph theory, and demonstrates just how impactful this type of analysis can be for real-world problems. With further research, hopefully an optimal mechanism will be implemented to maximize the number of kidney patients who are able to get live-saving transplants.

7. Acknowledgments

I want to thank Professor Postnikov for helping me learn how to improve my math presentation skills, along with helping me learn how to write a math paper. I'd also like to thank all of the students in the class for exposing me to so many cool new topics and consistently giving fascinating presentations.

References

1. I. Caragiannis et al. / Theoretical Computer Science 589 (2015) 5360
2. A.E. Roth et al. / Journal of Economic Theory 125 (2005) 151 188
3. I. Ashlagi, F. Fischer, I.A. Kash, A.D. Procaccia, Mix and match, in: Proceedings of the 11th ACM Conference on Electronic Commerce, EC, 2010, pp. 305314.
4. A.E. Roth, T. Snmez, M.U. nver, Kidney exchange, Q. J. Econ. 119 (2004) 457488

*Figures in paper from Caragiannis et al.