ARTICLE IN PRESS

Discrete Applied Mathematics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

A bounded-risk mechanism for the kidney exchange game

Hossein Esfandiari [a,*], Guy Kortsarz [b]

[a] Harvard University, School of Engineering and Applied Sciences, Cambridge, USA
[b] Rutgers University, Department of Computer Science, Camden, USA

ARTICLE INFO

Article history:
Received 4 December 2016
Received in revised form 18 December 2017
Accepted 19 December 2017
Available online xxxx

Keywords:
Kidney exchange game
Graph matching
Mechanism design
Low risk mechanism
Almost truthful

ABSTRACT

In this paper, we introduce and study the notion of *low risk mechanisms*. Intuitively, we say a mechanism is a low risk mechanism if the randomization of the mechanism does not affect the utility of agents by a lot. Specifically, we desire to design mechanisms in which the variances of the utility of agents are *small*. Inspired by this work, later, Procaccia et al. (0000) study the approximation–variance tradeoffs in mechanism design.

In particular, here we present a low risk mechanism for the *pairwise kidney exchange game*. This game naturally appears in situations that some service providers benefit from pairwise allocations on a network, such as the kidney exchanges between hospitals. Ashlagi et al. (2013) present a 2-approximation randomized truthful mechanism for this problem. This is the best-known result in this setting with multiple players. However, we note that the variance of the utility of an agent in this mechanism may be as large as $\Omega(n^2)$, where n is the number of vertices. Indeed, this is not desirable in a real application. In this paper, we resolve this issue by providing a 2-approximation randomized truthful mechanism in which the variance of the utility of each agent is at most $2+\epsilon$.

As a side result, we apply our technique to design a *deterministic* mechanism such that, if an agent deviates from the mechanism, she does not gain more than $2\lceil\log_2 m\rceil$, where m is the number of players.

1. Introduction

Kidney transplant is the only treatment for several types of kidney diseases. Since people have two kidneys and can survive with only one kidney, they can potentially donate one of their kidneys. It may be the case that a patient finds a family member or a friend willing to donate her kidney. Nevertheless, at times the kidney's donor is not compatible with the patient. These patient–donor pairs create a list of incompatible pairs. Consider two incompatible patient–donor pairs. If the donor of the first pair is compatible with the patient of the second pair and vice-versa, we can efficiently serve both patients without affecting the donors.

In this paper, we consider pairwise kidney exchange, even though there can be more complex combinations of transplantation of kidneys, that involves three or more pairs. Nevertheless, such chains are complicated to deal with in the real life applications since they need six or more simultaneous surgeries.

To make the pool of patient–donor pairs larger, hospitals combine their lists of pairs into one big pool, trying to increase the number of treated patients by exchanging pairs from different hospitals. This process is managed by some national supervisor. A centralized mechanism can look at all of the hospitals together and increase the total number of kidney exchanges. The problem is that for a hospital its key interest is to increase the number of its own served patients. Thus,

* Corresponding author.
E-mail address: esfandiari@seas.harvard.edu (H. Esfandiari).

https://doi.org/10.1016/j.dam.2017.12.029
0166-218X/

the hospital may not report some patient–donors pairs, namely, the hospital may report a partial list. This partial list is then matched by the national supervisors. Undisclosed set of pairs are matched by the hospitals locally, without the knowledge of the supervisor. This may have a negative effect on the number of served patients.

A challenging problem is to design a mechanism for the national supervisor, to convince the hospitals not to hide information, and report all of their pairs. In fact, if hiding any subset of vertices does not increase the utility of an agent, she has no intention to hide any vertex.

1.1. Notations and definitions

To model this and similar situations hospitals are called *agents*, and each patient–donor pair is modeled by a vertex. Let m be the number of agents. Each agent owns a disjoint set of vertices. We denote the vertex set of the ith agent by V_i and $\vec{V} = \{V_1, V_2, \ldots, V_m\}$ is called the vector of vertices of the agents. Denote an instance of the kidney exchange problem by $(G, \vec{V})$, where G is the underlying graph, $\vec{V}$ is the vector of vertices of the agents and E is the edge set. Each vertex in $G = (V, E)$ belongs to exactly one agent. Thus, $V = \cup_{i=1}^{m} V_i$ holds.

In this game, the utility of an agent i is the expected number of matched vertices in V_i and is denoted by u_i. Similarly, the utility of an agent i with respect to a matching M is the number of vertices of V_i matched by M and is denoted by $u_i(M)$. The social welfare of a mechanism is the size of the output matching.

A mechanism for the kidney-exchange game is the mechanism employed by the national supervisor to choose edges among the reported vertices. The process is a three step process. First the agent expose some of their vertices. Then the mechanism chooses a matching on the reported graph. Finally, each agent matches her unmatched vertices, including her non disclosed vertices, privately.

Formally, a kidney exchange mechanism F is a function from an instance of a kidney exchange problem $(G, \vec{V})$ to a matching M of G. The mechanism F may be randomized.

Given that some pairs are undisclosed, we say a kidney exchange mechanism F is α-approximation if for every graph G the number of matched vertices in the maximum matching of G is at most α times the expected number of matched vertices in $F(G)$. This means that for every graph G

$$\frac{|Opt(G)|}{E[|F(G)|]} \leq \alpha,$$

where $Opt(G)$ is the maximum matching in graph G, and the expectation is over the run of the mechanism F.

We define the notion of *bounded-risk* mechanisms as follows.

Definition 1.1. A mechanism is a bounded-risk mechanism if the variance of the utility of each agent is bounded by a constant.

Recall that agents are reporting the vertices, and they might not report (i.e. hide) a subset of their vertices if this increases their expected utility. We say a kidney exchange mechanism is truthful if there is no agent i who can increase her utility by hiding a subset of her vertices and matching them privately. In other words, for each agent i and any subset of vertices $V_i' \in V_i$ we have

$$u_i(F(G)) \geq u_i(F(G \setminus V_i')) + u_i(F(G \setminus V_i'), V_i'),$$

where $u_i(F(G \setminus V_i'), V_i')$ is the expected number of vertices that agent i matches privately if she hides V_i'. We also define *almost truthful* mechanisms, to use for our side result.

Definition 1.2. We say that a kidney exchange mechanism is almost truthful if by deviating from the mechanism, an agent have an additional gain of at most $O(log m)$ vertices in the utility, where m is the number of agents.

Indeed, in a real application finding the right subset of vertices to hide is costly. Roughly speaking, this cost involves extracting the information of m other agents, and hence, is an increasing function in m. Thus, in an almost truthful mechanism we hope that agents ignore gaining $O(\log m)$ vertices compared to the cost involved and report the true information. Remark that, in this paper we do not consider a cost for deviating from the mechanism, and thus, we use two different definitions of truthful mechanisms and almost truthful mechanisms.

1.2. Related work

The model considered in this paper was initiated by Sönmez and Ünver [14] and Ashlagi and Roth [5]. Sönmez and Ünver [14] show that there is no deterministic truthful mechanism that gets the maximum possible social welfare (see Fig. 1).

Achieving social welfare optimal mechanisms, which are truthful, is thus not possible. However, achieving approximate truthful mechanisms may be possible. Ashlagi et al. [3] used the same example as in Fig. 1 to show that there is no deterministic truthful mechanism for the kidney-exchange game, with approximation ratio better than 2. Moreover, they show that there is no randomized truthful mechanism with an approximation ratio better than 8/7. They also