Market Allocations in Big Data

Melika Abolhasani[1], Hossein Esfandiari[1], MohammadTaghi Hajiaghayi[1], and Hadi Yami[1]

Department of Computer Science, University of Maryland.
{mabolhas,hossein,hajiagha,hadiyami}@cs.umd.edu

Abstract. We consider the problem of envy-free assignment in matching markets when the edges of the graph is given in a streaming fashion. Our market is modeled as bipartite graph of unit demand buyers and distinct goods. Our goal is to find prices for goods and an assignment of goods to the buyers so that the assignment is envy free and we achieve other objectives such as maximizing social welfare or revenue. An Envy-free assignment is specially important to us since it assures stability and maximizes buyers' overall satifaction with their experience. We show that we can accomplish these goals when we only have $O(k^2)$ available memory which is a lowerbound for memory needed to solve this problem.

Keywords: market allocation, streaming, pricing, matching

1 Introduction

Streaming algorithms were first theoretically introduced in fields such as data mining and machine learning over 20 years ago in order to model problems in which the data cannot be accessed all at once. In other words, the data is arriving in an online fashion and limited memory is available for storing this data. These streams can only be read once (or a limited number of times). In this setting, even problems which have efficient algorithms when all the data is available offline might become challenging.

Over the past decade, there has been a significant demand for algorithms to process and handle the dynamic data coming from huge and growing graphs of social networks, webpages and their links, citations in academic work, etc, and therefore, research on graph streaming algorithms has become quite popular. The assumption in these models is that the edges of the graph (sometimes along with their weighted if the graph is weighted) are revealed one by one, and as with the case of other streaming models, limited memory for storage of these edges is available.

Bipartite matching is an inseparable model from many market design and economics problems when we have a set of items (goods, markets, etc) and a set of potential buyers (customers, companies, and so on), and the goal is to assign the items to the buyers, or maybe find what fraction of each market a company is interested in. Along with these models sometimes comes the criterions of envy-freeness, envy-free pricing and envy free allocations. An envy-free assignment is one in which every agent prefers his own assignment to other agents. We will formally define this notion later.

In this paper, we study the following market design problems. Suppose we are given a bipartite graph G with a set of n unit-demand buyers $b_1, \ldots, b_n$ on one side, and a set of k distinct items $v_1, \ldots, v_k$ on the other side. The utility of buyer b_j for item v_i is denoted by u_{v_i,b_j}, and is shown by a weighted edge between the corresponding two vertices. The price assigned to item i is denoted by p_i. The goal is to assign prices to items, and then items to buyers such that the market clears, the assignment is envy free, and the social welfare (or in the other problem that we consider, revenue) is maximized. This problem has been introduced and examined in the paper "On Profit-Maximizing Envy-free Pricing", however, we consider this problem in a streaming setting.

Definition 1. *An assignment of prices to items and items to buyers in our setting is called* envy-free*, if no buyer prefers his outcome to the outcome of other buyers. More formally, if we assign item v_{i_1} to buyer b_{j_1} and item v_{i_2} to buyer b_{j_2}, then we have $u_{v_{i_1},b_{j_1}} - p(i_1) \geq u_{v_{i_2},b_{j_1}} - p(i_2)$. In other words, buyer b_{j_1} prefers v_{i_1} (item assigned to him) rather than v_{i_2} (item assigned to another buyer).*

1.1 Our contributions

We design a streaming algorithm with $O(k^2)$ available memory which can perform as good as any algorithm which solves the same problem in the offline setting. Recall that k is the number of items. That is the memory that we need in streaming setting to achieve a result as good as the offline case is independent of the number of buyers.

1.2 Related Work

related work

2 Pricing problem: Maximizing Social Welfare

In this section, we consider the problem of assigning prices to items, and items to buyers in a streaming setting such that the assinment would be envy-free, and the *social welfare* is maximized. The social welfare would be sum of the weights (or utilites) of the assigned edges. As we explained before, we only use $O(k^2)$ memory for storage of the stream of edges. Our approach is to store the $k+1$ edges with maximum weight for each item, and to run the optimum algorithm to find the social welfare maximizing envy-free assignment in offline setting when the stream ends. In this section we show that

3 Pricing problem: Maximizing Revenue

Just like the previous section, we try to find an envy-free assignment of prices to items, and items to buyers in our described market when the input is observed in a streaming fashion. However, in this section, we aim to maximize *revenue* instead of social welfare. Again, we show that we can achieve the best possible solution to this problem when only $O(k^2)$ memory is available.

References