Budget-Constrained Auctions

Deeparnab Chakrabarty
Mentor: Vijay V. Vazirani

We are interested in the following revenue maximization problem for an auctioneer: Given a set of indivisible items on which bidders bid on, how to allocate these items to the bidders so as to maximize the total revenue obtained. Apart from the bids, the bidders also specify a budget constraint which is the maximum the auctioneer can charge the bidder.

Examples of such budget-constrained auctions include those used for the privatization of public assets in western Europe, or those for the distribution of radio spectra in the US, where the magnitude of the transactions involved put financial or liquidity constraints on bidders. More recently, the ad-auctions hosted by search engine companies are notable examples of budget-constrained auctions.

Budget constraints of agents open up many interesting avenues of research in auctions. Firstly, note that the above problem of revenue maximization is trivial with no budget constraints: allocate an item to its highest bidder. However, the algorithmic problem with budget constraints is NP-hard. Secondly, apart from the algorithmic aspect, there is an information eliciting aspect in auctions. After all, the agents are free to bid whatever they like on the items and nothing, apart from the risk of not winning the item, stops them from bidding lower than what they truly value the item. Thus a desirable property of an auction (which is nothing but a mechanism which takes bids as input and returns an allocations and demands payment) is to be truth-telling: that is, bidding the true value of an item is a weakly dominant strategy for a bidder. One famous such mechanism which has been used extensively in auctions is the VCG mechanism. However, it is known that with budget constraints, VCG no longer retains the truthfulness property leaving open the question of a suitable auction design.

In recent work with Gagan Goel, we tackled the algorithmic problem mentioned in the first paragraph and obtained approximation algorithms for the problem. In fact, our algorithms guarantee a revenue within 75% of the optimal revenue possible.

In this proposal we suggest two lines of work. Firstly, we believe that 75% is more of a theoretical threshold and the algorithm should do much better in practice. Moreover, we have heuristics in mind which could outperform the algorithm in question in certain instances. We propose a detailed experimental study of the algorithms and heuristics. To analyze the performance, we will test our algorithms on synthetic as well as real data. Synthetic data will be various random models of bidders and budgets. Real data, on the other hand, might be more difficult to arrange. The main reason for this being that budgets are mostly private information which is hard to get ones hand on to. However, many auctions (including the FCC auctions) have bidding data available and certain accepted methods exist to estimate budgets.

Secondly, we believe one of the algorithms presented in work with Goel can be modified to give an auction mechanism. We propose to formally devise this mechanism and study its properties. We do not believe it will be truthful, however, we wish to study the property of the equilibria reached in these auctions under various assumptions on information available. Our hope is to come with an auction design where maybe truth-telling may be an α-dominant strategy: deviating from truth telling can only give you at most α times more pay-off.