My research interest primarily lies in theoretical computer science and more specifically in approximation algorithms and game theory. In the first half of my Ph.D. the focus of my research was mostly on approximation algorithms. After spending a summer at Yahoo! as an intern, I became interested in the sponsored search problems and as a result the new field of algorithmic game theory which is the intersection of algorithms and game theory. During the past two years, I studied various problems related to sponsored search and electronic commerce from both combinatorial aspects as well the game theoretic aspects.

Next, I provide a short summary of the problems that I would like to work on in future and a very short summary of the work that I have done during my Ph.D.

Recurring Auctions

A large number of sellers in eBay sell their goods through eBay's ascending proxy auction. These sellers usually put only a small number of copies for sale at any time although they may have a large inventory. Most of the times, a bidder interested in buying a copy keeps participating in subsequent auctions until she wins. In classical analysis of auctions, a bidder who does not win an auction, is assumed to have a utility of 0 and that is used as the basis for the argument that truthful bidding is a dominant strategy for second price auctions. However, that assumption fails to hold when dealing with a series of auctions for similar items considering that a loosing bidder can still participate in a subsequent auction and obtain a positive utility at a later time. It is easy to see that in such a setting, bidding one's true value is no longer a dominant strategy for bidders. Also, as a consequence of this non-truthful bidding, a seller should strategize on how many items to auction concurrently and for how long.

In a joint work with Alaei[3], we study the bidding strategy and equilibrium of such auctions. To account for the time that a bidder waits before winning an item we use a discount factor for the utility obtained at a later time. We assume the number of bidders and the distribution of their valuations are drawn from a publicly known distribution that has converged over time. We show how to compute the dominant bidding strategy in the symmetric case where the valuations of bidders are i.i.d random variables and discount factor is the same for everyone.

One interesting problem is to compute a seller's strategy in terms of the number of concurrent auctions and their duration to maximize his revenue. It is also interesting to consider the case when there are multiple sellers.

The problem of recurring eBay auctions can be generalized to a wider class of recurring auctions in which valuations of participants are drawn from a known distribution, but the set of participants change from one auction to the next one. For example in the case of eBay auctions, buyers usually participate only in one of the many concurrent auctions and when they loose, they choose from one of the many subsequent concurrent auctions to participate in. The constant changing of the set of competitors and incomplete information about others actions rule out the possibility of monitoring and therefore the classical analysis of repeated games and folk theorem are not applicable.

Prior Free Optimal Mechanism Design

One of the interesting problems in the mechanism design is to find a truthful auction that would maximize the seller's revenue assuming no information about the bidders valuations' which is referred to as *prior free optimal truthful mechanism design*. In the context of auctions for digital goods, an interesting truthful Random Sampling Auction (RSOP) has been proposed by Goldberg et al. [8]; This method has been analyzed further by Feige et al. [6] considering the seller's revenue, who have shown that it is 15-competitive in the worst case which was still far from the conjectured competitive ratio of 4. In a joint work with Alaei, and Srinivasan [4], we showed that if the winner's index is greater than 6 in the optimal solution, the competitive ratio of RSOP is better than 4 and in the rest of the cases, the competitive ratio is no worse than 4.68. We employed a mix of probabilistic techniques and computer-aided analysis. We also showed that even if the winner's index gets arbitrarily large, one can construct instances for which the competitive ratio is no better than 2.65. One interesting question left is to find a tight lower bound for each winner index and to find a matching upper bound. In [4], we present a combinatorial solution for the case when we have only two type of bids. However it is not clear if the solution can be generalized. Another interesting problem is to find a combinatorial approach for analyzing the competitive ratio of RSOP.

Another major question in optimal prior free mechanisms, is to find an optimal auction that will match the lower bound given in [7]. Hartline et al. showed in [9] that for 3 bidders there is an auction with the competitive ratio 13/6 which matches the lower bound, however the problem is still open for more than 3 bidders.

Online problems with known input distributions

In the context of online algorithms, one approach is to design algorithms with a good worst case performance. Another approach is to design algorithms with good expected performance. Computer scientists have traditionally adopted the former approach since it does not depend on the probability distribution of the input and worst case analysis is generally much simpler than expected performance analysis. Nevertheless, many algorithms with theoretically not so good worst case performance do better in practice as they have a better performance in expectation.

The offline version of a good number of online algorithms like online resource allocation [5], use linear programming to find an intermediate fractional solution. A similar LP approach usually cannot be used for the online version since solving the LP requires all the information to be known in advance. Still it might be possible to find certain forms of LP for which an approximation to the optimal solution could be constructed incrementally. Consider for example the following class of linear programs:

Definition 0.1 (Online Packing LP:)**.** Consider the linear program of the following form:

$$\text{maximzie:} \quad \mathbf{c}.\mathbf{x} \tag{0.1}$$
$$\mathbf{A}\mathbf{x} \leq \mathbf{b} \tag{0.2}$$
$$x_i \in [0,1] \tag{0.3}$$

such that the matrix $\mathbf{A}$ and the vectors $\mathbf{b}$ and $\mathbf{c}$ are non-negative. We say each variable x_i is either *active* or *inactive*. If inactive then $x_i = 0$, otherwise x_i could take any value in $[0, 1]$. We learn the status (i.e., active/inactive) of variable x_i at time t_i and as soon as we find out its status, we are required to specify its value (if inactive it should be set to 0). Once initialized, the value of a variable cannot be changed at a later time. As for distribution information, we know that each variable x_i will be active with probability p_i independent of other variables. The goal is to design an online algorithm that maximizes the expected value of the objective function.

As an example, the online resource allocation and online matching problem could be easily modeled as an online packing LP.

If we could efficiently compute the expected value of the objective function for the optimal offline solver, then we could use the following online algorithm. At each time t_i when the status of variable x_i is learned, if x_i is active then compute the expected value of the objective for the rest of the LP once for $x_i = 0$ and once for $x_i = 1$ and then choose the one with the highest total expected value. Consequently, it is interesting to see if there is an efficient way to compute the expected value of the objective function. It is worth mentioning that we can efficiently compute an upper bound on the expected value of the objective by adding the constraints $x_i \leq p_i$ and solving the offline LP. It might also be possible to compute an approximation to the optimal online algorithm if we could compute a constant factor lower bound for the expected value of the objective function and then use that instead of the exact expected value in the online algorithm.

Uniform Price Reverse Auctions

This problem is motivated by the troubled assets reverse auction proposed by [1] which was meant to be used by the government as an efficient way of infusing liquidity into the market in places where it could be most useful. This problem can be modeled as a uniform price auction for a single divisible good in which the good to be auctioned is the budget of the government. Bids and payments are made in terms of shares of securities that bidders hold weighted by the reference price of each security. Such a model not only makes it much easier to analyze the equilibrium and the bidding strategies but also makes it possible to apply more sophisticated mechanisms like Ausubel's clinching in the context of reverse auctions.

In a joint with alaei [1], we proposed a more general framework that applies to a wider class of games including the reverse auction. In this class of games, each participant i chooses a real number a_i through some process such that at the end the utility of each participant can be written as $u_i(a_i, A)$ where $A = \sum_i a_i$. That means the utility of each participant should depend only on her own choice and the aggregate choice of all of the participants (including herself). The aforementioned reverse auction problem as well as some other classical problems including uniform price auctions, cartels, etc. could be expressed in this model. We showed in [1] that if the utility functions u_i have certain properties then there is always an equilibrium or a continuum of equilibria with unique payoffs. It is interesting to study other characteristics of such game like whether the equilibrium point of such games can always be computed using a simple mechanism that uses an ascending/descending clock. It is also interesting to see if the same generalization can be applied to the class of games in which the participants choose a vector instead of a single number and their utilities then depend to their own vector as well as the aggregate vector of all participants' choices. Such a framework could be useful for modeling package auctions (e.g., each bidder submits her demand from each item). Then, the interesting question is to see under what conditions a mechanism using multiple clocks exists that can find the equilibrium of such a game.

Previous Work

I have worked on other combinatorial problems with application to sponsored search as well as some other fields. In a joint work with Alaei [2], we generalized the notion of submodularity to sequences (in contrast to sets) and we showed that the greedy approach will find a solution that is at least $1 - \frac{1}{e}$ of the optimal solution for maximization problems that satisfy monotonicity as well as sequence submodularity. As an application, one can easily prove that greedy solution is a $1 - \frac{1}{e}$ approximation for online ad allocation problem when the probability distribution of query keywords does not change over time (but it does not need to be known). Another problem that I studied during my internship at Yahoo, was the query rewriting problem. Query rewriting has been studied before in the field of Information Retrieval. In a joint work with Chang, Kumar, and Wang [13], we defined the challenges in query rewriting for advertisement and we showed that greedy approach is a promising solution from both theoretical and experimental perspective.

Also, I looked at problems with applications to other fields. In [10], we improved approximation factor for the data migration problem. The previous bound for the half duplex model was 9.5 [11].

We improved it to 6.5. We also gave a bound of 4 for the full duplex model problem. In [12] we studied a new variant of TSP problem in which the salesman is traveling around with a car that has a limited tank capacity and there are some specified gas station locations where it can be refilled. We gave a combinatorial solution for the problem of going from one city to another and we gave an approximation algorithm for the problem in which the salesman wants to visit a set of cities.

Conclusion

In short, I would like to pursue my research in the field of algorithmic game theory. There are plenty of research problems in this area that are interesting both practically and theoretically. My plan is to mainly focus my research on the problems in this area, while continuing to work on combinatorial optimization as well.

References

[1] S. Alaei and A. Malekian. An analysis of troubled assets reversed auction. Working Paper, 2008.

[2] S. Alaei and A. Malekian. Maximizing sequence-submodular functions and its application to online advertising. Working Paper, 2008.

[3] S. Alaei and A. Malekian. Strategic bidding in ebay recurring auctions. Working Paper, 2008.

[4] S. Alaei, A. Malekian, and A. Srinivasan. On random sampling auctions for digital goods. 2008.

[5] G. Calinescu, A. Chakrabarti, H. J. Karloff, and Y. Rabani. Improved approximation algorithms for resource allocation. In *IPCO*, pages 401–414, 2002.

[6] U. Feige, A. Flaxman, J. D. Hartline, and R. D. Kleinberg. On the competitive ratio of the random sampling auction. In *WINE*, pages 878–886, 2005.

[7] A. V. Goldberg, J. D. Hartline, A. R. Karlin, and M. E. Saks. A lower bound on the competitive ratio of truthful auctions. In *STACS*, pages 644–655, 2004.

[8] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *SODA*, pages 735–744, 2001.

[9] J. D. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *ACM Conference on Electronic Commerce*, pages 175–182, 2005.

[10] S. Khuller, Y. A. Kim, and A. Malekian. Improved algorithms for data migration. In *APPROX-RANDOM*, pages 164–175, 2006.

[11] S. Khuller, Y.-A. Kim, and Y.-C. J. Wan. Algorithms for data migration with cloning. In *PODS '03: Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 27–36, 2003.

[12] S. Khuller, A. Malekian, and J. Mestre. To fill or not to fill: The gas station problem. In *ESA*, pages 534–545, 2007.

[13] A. Malekian, C. C. Chang, R. Kumar, and G. Wang. Query rewrites for keyword-based advertising. In *ACM Conference on Electronic Commerce*, 2008.