Approximating the Qualitative Vickrey Auction by a Negotiation Protocol[1]

Koen V. Hindriks Dmytro Tykhonov Mathijs de Weerdt

Delft University of Technology, Delft 2628CD, The Netherlands

Abstract

Some auction mechanisms can be shown to be efficient and strategy-proof. However, they generally require that the preferences of at least one side of the auction are publicly known. However, sometimes it is very costly, impossible, or undesirable to publicly announce such preferences. It thus is interesting to find methods that do not impose this constraint but still approximate the outcome of the auction. In this paper we show that a multi-round multi-party negotiation protocol may be used to this end if the negotiating agents are capable of learning opponent preferences. The latter condition can be met by current state of the art negotiation technology. We show that this protocol approximates the theoretical outcome predicted by a so-called Qualitative Vickrey auction mechanism (even) on a complex multi-issue domain.

1 Introduction

There are many different types of auctions, for which it is possible to show that they have nice theoretical properties. However, some of the mechanisms proposed also require constraints to be in place which are not easy to meet in practice. This is particularly true for auctions which may result in complex, multi-issue outcomes. It thus becomes interesting to look for alternative methods that may be used that guarantee outcomes that approximate the auction mechanism. In this paper, we study a particular auction mechanism called a Qualitative Vickrey Auction (QVA) [1]. This is a generalization of the well-known Vickrey auction to a general complex multi-issue setting where payments are not essential.

The QVA studied here is particularly useful in a context where a single buyer tries to obtain a *complex deal* with one out of many sellers that are interested in making such a deal. An example is a buyer that is interested in buying a supercomputer. A range of potential suppliers is available that may provide a supercomputer. Apart from price, which may be fixed by a budget and therefore less interesting, supercomputers have many features (processing speed, memory, etc.) that need to be settled. Such a deal thus is complex as many issues have to be agreed upon. The QVA mechanism obtains a Pareto-efficient outcome that involves the seller that can make the best deal still acceptable to him. The mechanism requires that the buyer publicly announces her preferences. This is, however, hard to realize in practice for various reasons: the buyer may not know the complete domain of possible outcomes, it may be hard to specify all preferences over a complex set of outcomes, or the buyer may not want to publicly reveal all her preferences in complete detail.

In this paper, we show that the theoretical outcome predicted by the QVA mechanism can be approximated by a specific negotiation protocol. The only assumption that we need to make to obtain this result is that the negotiating agents are able to learn the preferences of their opponents during a single negotiation session. Techniques to do so are available [2], making our proposal one that can be implemented given the current state of the art in negotiation. The idea introduced here is that a protocol that consists of multiple negotiation rounds in which sellers are provided an opportunity to outbid the winner of the previous round may be used to approximate the QVA.

[1]Originally appeared in: K. V. Hindriks, D. Tykhonov and M. de Weerdt, "Approximating the Qualitative Vickrey Auction by a Negotiation Protocol", In: Proc. of The Eleventh Int. Workshop on Agent Mediated Electronic Commerce (AMEC 2009), May 2009

2 Negotiation Protocols

Using bilateral negotiation based on a simple alternating offers protocol between the buyer and each of the sellers does not by itself take the fact into account that multiple sellers are contending for a deal with the seller. To incorporate this fact, we introduce a protocol that consists of multiple rounds of bilateral negotiations between the buyer and the sellers. After each round, the buyer communicates the *winning agreement* of that round to the sellers that did not win. All of the sellers then are provided with the opportunity to improve the agreement they reached with the buyer in the last round in a next round of negotiation sessions. A seller will do so if he can make an offer that has a utility value above his reservation value, and that he thinks has a higher utility to the buyer than the winning agreement of the last round. Negotiation is therefore assumed to resume for the seller in a next round, starting with the agreement reached in the last round. This process continues until no seller is prepared to negotiate in a next round to improve their last offer. The winning agreement of the last round then is the final agreement of the negotiation process.

It is advantageous for a seller to learn the buyer's preferences in this process, because this can be used to reach an agreement that satisfies the buyer as best as possible while at the same time maximizing the utility for the seller itself. Such an opponent model can be used to assess if an offer can be made that has the same utility value as the winning agreement from the point of view of the seller but that has a higher utility for the buyer. Only if such an offer cannot be made, an additional concession has to be made. The size of the negotiation space is decreased in every next round due to the fact that the buyer will only accept offers that improve the winning agreement reached in previous round. This process forces the final agreement closer to that of the reservation value of the sellers, in line with the dominant strategy sellers have in the QVA.

The same protocol can also be applied without informing sellers about intermediate agreements. In this case, the buyer only indicates to a seller that it did not win in the last round. The winning agreement of the previous round thus can no longer be used as a reference point that needs to be improved upon from the buyer's point of view, and a seller instead continues negotiation in the next round with the agreement it reached itself in the previous round. As a result, the sellers have less information on how to outbid the winning seller of the previous round. Still, the buyer does have this information as it knows the winning agreement of the previous round and, therefore, would only accept offers of a seller that improve the winning agreement of the previous round. The agreement reached without revealing the winning agreement in each round converges to that of the theoretical outcome of the QVA, assuming the negotiating parties are able to learn the preferences of their opponent. As the sellers have less information in this second setup, they will have more difficulty in proposing offers that improve the winning agreement of previous rounds and more rounds may be needed to explore options to find such offers. When the negotiation protocol terminates and a final agreement is reached, this agreement is made public in order to allow sellers to verify that the buyer has not manipulated the process. This is sufficient for sellers that have a reasonable opponent model to assess whether the process has been fair, as they can check whether they believe they could have improved this final agreement to obtain a deal.

3 Conclusion

In this paper we proposed and experimentally validated two setups based on multi-round negotiations that are capable of approximating the Qualitative Vickrey Auction (QVA), obtaining a (near) Pareto-efficient outcome where the best seller wins (in the eyes of the buyer). In both setups the buyer as well as the sellers can use any reasonable negotiation strategy. The protocol proposed introduces multiple negotiation rounds in which sellers that lost in the previous round are given an opportunity to improve their offers and possibly outbid the winner. We showed experimentally that both setups converge to the results of the QVA.

The results of the second experiment indicate that even if no information is made public until the end of the negotiation the protocol converges to the results of the QVA. The number of rounds needed to find the winning contract is, however, significantly higher than in the first setup. This can be explained by the fact that sellers have no information about the winning agreement of the previous negotiation round and would have to make several offers before they can outbid the winner.

References

[1] B. Paul Harrenstein, Mathijs M. de Weerdt, and Vincent Conitzer. A qualitative vickrey auction. In *Proceedings of the ACM Conference on Electronic Commerce*. ACM, 2009.

[2] Koen V. Hindriks and Dmytro Tykhonov. Opponent modelling in automated multi-issue negotiation using bayesian learning. In *Proc. of the 7th Int. Conference on Autonomous Agents and Multiagent Systems*, pages 331–338, 2008.