Searching for the Possibility – Impossibility Border of Truthful Mechanism Design

RON LAVI
Faculty of Industrial Engineering and Management, The Technion, Israel

One of the first results to merge auction theory and algorithmic theory, [Lehmann et al. 2002], considers a combinatorial auction setting, and describes a computationally efficient and truthful auction for "single-minded" players, i.e. players that desire only one specific subset of items. The natural continuation would have been the case of multi-minded players. Indeed, the literature has managed to extend the results to the more general cases, by using randomness (e.g. [Lavi and Swamy 2005; Dobzinski et al. 2006]). However, in the deterministic case, virtually no positive advancement was made ever since.[1] An even worse state of affairs exists in the algorithmic domain of job scheduling. The seminal paper on algorithmic mechanism design [Nisan and Ronen 2001] describes a basic impossibility for the multi-dimensional version of the problem, while [Archer and Tardos 2001] observe the possibility in the single-dimensional case. Here, the situation is even worse, as the transition from the possibility to the impossibility does not depend on any computational assumptions, and since we do not even know if randomness can make a significant difference. Obviously, it seems that there is a specific, inherent difficulty that prevents the design of truthful mechanisms for multi-dimensional domains. But, to date, we are not able to give a formal statement of this informal conclusion, and to exactly characterize the difficulties. *What structural properties turn plain difficulties into exact mathematical impossibilities?* This short note aims to raise the awareness to this important question.

Recall that in a combinatorial auction we wish to find an allocation $(S_1, ..., S_n)$ of m items to n players, where the usual goal is to maximize the social welfare, i.e. the term $\sum_i v_i(S_i)$. The key assumption is that the valuation function $v_i(\cdot)$ of player i is private information of the player, hence we desire *truthful* mechanisms, in which the dominant strategy of each player is to reveal her true valuation. The classic VCG mechanism (due to [Vickrey 1961; Clarke 1971; Groves 1973]) obtains exactly that: it is truthful, and it maximizes the welfare. Unfortunately, it is

[1] Though there have been a few advances for several special cases of the problem, e.g. by [Bartal et al. 2003].

Author's address: ronlavi@ie.technion.ac.il

NP-hard to compute. Can we design a polynomial-time truthful mechanism that *approximates* the optimal welfare? If we allow randomization, then the answer is positive. [Lavi and Swamy 2005], [Dobzinski et al. 2006], and [Dobzinski 2007] give polynomial-time mechanisms that obtain $O(\sqrt{m})$-approximations. The first mechanism is truthful-in-expectation, while the latter two are universally truthful, i.e. truthful for any realization of the coin tosses.[2] These solutions do not *guarantee* the approximation ratio – there is some probability of achieving an outcome with a much lower welfare. This is problematic, especially since we cannot use standard amplification methods to reduce this probability to become as small as we wish, as the repetition will destroy the incentive properties.

Motivated by this question, [Lavi et al. 2003] ask if there even *exist* truthful deterministic mechanisms for general combinatorial auctions, other than VCG, and show an impossibility[3]: every truthful combinatorial auction for general bidders, with a dense range, that satisfies unanimity, decisiveness, and weak-IIA, must be a weighted welfare maximizer. The range of alternatives is "dense" if the auction does *not* simply bundle all items, always allocating them together to one of the players, thereby bringing back single-dimensionality. Unanimity states that if every player is single-minded in a way that enables all players to simultaneously win (and furthermore this allocation is in the range), then all players win. This unanimity condition is common in the social choice literature, and was first considered by [Arrow 1951]. Decisiveness states that, for every type declaration of the other players, there always exists a type declaration for player i that will award her all items, and another type declaration in which she will get nothing. This condition seems very natural in the context of auctions, as most common "real life" auctions satisfy it. Weak-IIA states that a player cannot cause the overall allocation to change (by changing her own declaration), while her values for the original bundle as well as for the new bundle remain unchanged. This is perhaps the most subtle condition of the four, and is motivated by its similarity to Arrow's well-known IIA condition.

While [Lavi et al. 2003] show that adding *some* requirements is a must, as there do exist degenerate examples of truthful combinatorial auctions that are not weighted welfare maximizers, their result is clearly not the end of the story, as it does not answer our basic question. The need to approximately maximize the welfare does not imply these four properties, and hence, hypothetically, it may well be that by violating some of the requirements, we could come up with a truthful and polynomial-time auction that achieves the $\sqrt{m}$ approximation. We are still missing a complete understanding of the limitations of truthfulness. This search is relevant to many other algorithmic problems (as we continue to detail below), but in the context of auctions perhaps the most relevant question concerns the design of truthful auctions for restricted valuations. [Feige 2006] shows an algorithmic possibility: with sub-additive valuations, one can achieve a 2-approximation in polynomial-time. However, the incentives issue has not been settled yet, and *no truthful mechanism with a constant*

[2]However, the mechanism of [Lavi and Swamy 2005] optimally solves also the case where we have multiple copies of each item.
[3]We actually describe a somewhat stronger result that is given in [Lavi et al. 2007]

approximation ratio is known to date, be it deterministic, or randomized. In this respect, further understanding the limitations of truthfulness is extremely vital.

An interesting corollary of [Lavi et al. 2003] is for the case of "multi-unit" auctions (i.e. when all items are identical). Their theorem holds for this case, and furthermore, they show that all the additional assumptions can be replaced by the requirement that all items are always allocated, if we have an auction with two players that obtains an approximation ratio lower than 2. By further claiming that weighted welfare maximizers are as hard as exact welfare maximizers, they get: Every polynomial-time truthful multi-unit auction among two players, that always allocates all items, has an approximation ratio of at least 2.[4] In contrast, without the truthfulness requirement, one can obtain an FPTAS for the problem by reducing it to knapsack. Interestingly, [Dobzinski and Nisan 2007] give a truthful 2-approximation for every number of players (not only two), that uses a polynomial number of value queries. Hence a concrete open question emerges: does there exist a better-than-2 approximation mechanism that is both polynomial-time and truthful. By the result of [Lavi et al. 2003], it will somehow need to rely on the fact that items are not always allocated, but it is not clear how such a relaxation can help. We should also note that [Dobzinski and Nisan 2007] additionally give a truthful PTAS for k-minded players, for every *fixed k*. Thus, the question is open mainly for *general* valuations.

The second domain that we discuss is job scheduling: n jobs are to be assigned to m machines, where machine i incurs a cost of p_{ij} from executing job j. Importantly, this cost is private information of that machine. The machines are assumed to be strategic, each one selfishly trying to minimize its own cost. The load of machine i is the sum of costs of the jobs assigned to i, and the maximal load over all machines (in a given schedule) is termed the "makespan" of the schedule. We wish to design a truthful mechanism that minimizes the makespan. This goal is inherently different than welfare-maximization, and, while we can still use VCG here, its outcome may be far from optimal. Indeed, [Nisan and Ronen 2001], who have first studied this problem in the context of mechanism design, observed that VCG provides only an m-approximation to the optimal makespan. More importantly, they have shown that *no truthful deterministic mechanism can obtain an approximation ratio better than 2*, regardless of computational considerations. To date, we do not know of any truthful mechanism, deterministic or randomized, that achieves an approximation ratio that is significantly[5] lower than m. [Archer and Tardos 2001], on the other hand, considered a natural restriction of this domain, that makes it single-dimensional, and showed how to construct many possibilities (namely a truthful optimal mechanism, and a polynomial-time and truthful approximation). Thus, here too we see the contrast between single-dimensionality and multi-dimensionality. In fact, our knowledge here is significantly thinner: we do not know if randomization can help in any way, and the gap between the proven impossibility and the

[4]A bidding language that generalizes "OR bids" is needed for this result. [Dobzinski and Nisan 2007] show how to replace this, in the communication model, with "value queries".

[5][Mu'alem and Schapira 2007] give a randomized truthful mechanism with approximation ratio $\frac{7}{8}m$.

proven possibility is extremely and unacceptably high. Recently, [Mu'alem and Schapira 2007] and [Christodoulou et al. 2007] extend the lower bound to randomized and fractional mechanisms, and [Koutsoupias and Vidali 2007] slightly increase the lower bound to 2.61. These new results are technically non-trivial, and this fact only emphasizes the large gaps that still remain.

In this context, an interesting demonstration of a possibility result to a multi-dimensional scheduling domain was given by [Lavi and Swamy 2007], who give a truthful 2-approximation in case the processing time of each job is either "low" or "high". They do not rely on explicit price constructions, but rather on a cycle-monotonicity condition (defined in [Rochet 1987]) that completely characterizes truthfulness. This condition reduces to a "weak monotonicity" condition (see [Lavi et al. 2003; Saks and Yu 2005; Bikhchandani et al. 2006]) if assuming a convex domain, and in fact *all* the impossibilities mentioned in this note actually rely on the monotonicity condition in the proof itself. Thus, the bottom line of all these questions can be summarized by asking to characterize the algorithmic implications of weak monotonicity, in multi-dimensional problem domains. One should note here that in the extreme case of a completely unrestricted domain, [Roberts 1979] gives a full answer (namely that nothing but VCG is possible) – the point is to understand the implications for restricted multi-dimensional cases, e.g. for the two important problem domains that were discussed here.

REFERENCES

Archer, A. and Tardos, É. 2001. Truthful mechanisms for one-parameter agents. In *The Proc. 42nd Annual Symposium on Foundations of Computer Science (FOCS)*. 482–491.

Arrow, K. 1951. *Social Choice and Individual Values*. Wiley.

Bartal, Y., Gonen, R., and Nisan, N. 2003. Incentive compatible multi-unit combinatorial auctions. In *The Proc. of the 9th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'03)*.

Bikhchandani, S., Chatterjee, S., Lavi, R., Mu'alem, A., Nisan, N., and Sen, A. 2006. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica 74(4)*, 1109–1132.

Christodoulou, G., Koutsoupias, E., and Kovács, A. 2007. Mechanism design for fractional scheduling on unrelated machines. In *The Proc. of the 34th International Colloquium on Automata, Languages and Programming (ICALP)*.

Clarke, E. 1971. Multipart pricing of public goods. *Public Choice 8*, 17–33.

Dobzinski, S. 2007. Two randomized mechanisms for combinatorial auctions. In *The Proc. of the 10th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*.

Dobzinski, S. and Nisan, N. 2007. Mechanisms for multi-unit auctions. In *The Proc. of the 8th ACM Conference on Electronic Commerce (ACM-EC)*.

Dobzinski, S., Nisan, N., and Schapira, M. 2006. Truthful randomized mechanisms for combinatorial auctions. In *The Proc. of the 38th ACM Symposium on Theory of Computing (STOC)*.

Feige, U. 2006. On maximizing welfare when utility functions are subadditive. In *The Proc. of the 38th ACM Symposium on Theory of Computing (STOC)*.

Groves, T. 1973. Incentives in teams. *Econometrica 41(4)*, 617–631.

Koutsoupias, E. and Vidali, A. 2007. A lower bound of $1+\phi$ for truthful scheduling mechanisms. In *The Proc. of the 32nd International Symposium on Mathematical Foundations of Computer Science (MFCS)*.

Lavi, R., Mu'alem, A., and Nisan, N. 2003. Towards a characterization of truthful combinatorial auctions. In *The Proc. 44th Annual Symposium on Foundations of Computer Science (FOCS)*. 574–583.

Lavi, R., Mu'alem, A., and Nisan, N. 2007. An impossibility result for ex-post implementable multi-item auctions with private values. Manuscript.

Lavi, R. and Swamy, C. 2005. Truthful and near-optimal mechanism design via linear programming. In *The Proc. of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*.

Lavi, R. and Swamy, C. 2007. Truthful mechanism design for multidimensional scheduling. In *The Proc. of the 8th ACM Conference on Electronic Commerce (ACM-EC)*.

Lehmann, D., O'Callaghan, L., and Shoham, Y. 2002. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM 49(5)*, 577–602.

Mu'alem, A. and Schapira, M. 2007. Setting lower bounds on truthfulness. In *The Proc. of the 18th Symposium on Discrete Algorithms (SODA)*.

Nisan, N. and Ronen, A. 2001. Algorithmic mechanism design. *Games and Economic Behavior 35*, 166–196.

Roberts, K. 1979. The characterization of implementable choice rules. In *Aggregation and Revelation of Preferences*, J.-J. Laffont, Ed. North-Holland, 321–349.

Rochet, J. C. 1987. A necessary and sufficient condition for rationalizability in a quasilinear context. *Journal of Mathematical Economics 16*, 191–200.

Saks, M. and Yu, L. 2005. Weak monotonicity suffices for truthfulness on convex domains. In *The Proc. of the 6th ACM Conference on Electronic Commerce (ACM-EC)*. 286–293.

Vickrey, W. 1961. Counterspeculations, auctions, and competitive sealed tenders. *Journal of Finance 16*, 8–37.