Research plan

Paolo Penna

We plan to study the algorithmic properties that characterize incentive compatible mechanisms, under various notions of "incentive compatibility" that are driven by the applications. Following the (algorithmic) mechanism design approach of Nisan and Ronen [NR01], one can formulate in a mathematically precise way the fact that it is not *convenient* to manipulate a certain "protocol", and thus the protocol is safely implemented in a distributed setting. A mechanism consists of an *algorithm* and a suitable *payment* function whose combination should induce the players to the behave according to the "prescribed rules". Such mechanisms exist only if the underlying algorithm obeys certain requirements, in addition to the usual optimality of the solution.

(In)approximability results that incorporate these requirements can be regarded as the loss of performance due to a selfish behavior (the mechanism cannot give optimal solutions because otherwise it would not be incentive compatible). Our study should contribute to this important line of research by investigating different incentive compatibility conditions and by deriving new algorithmic/game theoretic techniques. We describe some of these issues more in detail below.

Truthfulness and approximability. Truthfulness prescribes that truth-telling is a dominant (utility maximizing) strategy for the players. Truthfulness is in fact an *algorithmic* property since it is equivalent to certain "monotonicity" conditions on the algorithm (how the output changes when some part of the input changes). These general conditions, however, are difficult to understand and to use. Indeed, depending on the problem under consideration, the result can be very different:

- For certain problems, one must give up optimality of the algorithm and look for suitable *approximations*. The most notable example is the scheduling problem introduced by Nisan and Ronen [NR01].

- For other problems, though existing *computationally efficient* algorithms violate the required monotonicity condition, either they can be "adapted" or new algorithms with the same approximation guarantee can be designed "from scratch".

Typically, problems in the first class have a "multidimensional domain", while those in the latter have a simpler "one-dimensional domain". We would like to investigate new mechanism design techniques (and characterizations of truthfulness) together with the approximability of algorithms that obey these conditions.

Concrete starting questions. An interesting line of research consists in characterizing truthfulness for certain domains and, from this, obtain optimal mechanism for *non-utilitarian* problems (i.e., when the goal is not the minimization of the sum of the players' costs). Of a particular interest are the min-max problems (including makespan minimization). Is it possible to extend the simple domains studied in [AT01, LS09, APP09] and get positive results? Does randomization help in such domains?

Collusion. The effect of *colluding* players is a central (and sometimes unavoidable) question. Schummer [Sch00] proves that, if players can *exchange compensations*, then the only mechanisms that can resist to collusions are the trivial (useless) mechanisms that output a fixed solution. Mechanisms that resist to collusion if players do *not exchange compensations* have been successfully developed in the context of cost-sharing problems (see the recent characterization [PV10] and references therein).

Concrete starting questions. It is not clear how central is the continuity assumption on the domain used by Schummer [Sch00]. Is there a "discrete" domain for which his impossibility result does not apply? A possible candidate might be the two-values domains in [LS09].

False identities. This is a typical scenario for Internet applications. The issue of false identities in mechanism design has been raised by Yokoo et al. [YSM04] in the context of combinatorial auctions. Recently, the same issue has been considered for "budget constrained" problems like cost-sharing [PSSW09] or mechanisms without money [TIY10]. In all such cases, existing truthful mechanisms can be manipulated using false identities.

Concrete starting questions. Characterizations of incentive compatible mechanisms that obey this additional requirement are of great interest. For instance, which cost-sharing mechanisms can be "easily" adapted so to deal with false identities? What property characterizes mechanisms that resist to false identities? What is the "price" of false identities? Do reputation mechanisms help?

Repeated games, mechanism design, protocol analysis. A recent work by Nisan et al. [NSVZ11] showed that *repeated games* provide a unifying approach to many existing positive results: cost sharing, mechanisms without money, convergence and incentive compatibility of BGP and of certain TCP/IP games.

Roughly, the desired incentive compatibility condition is guaranteed if the game of interest can be expressed as a repeated base game possessing a *stronger* incentive compatibility condition. Here 'repeated game' means that players keep playing best response (in the base game).

Concrete starting questions. We would like to investigate the limitations of this approach and the possibility to extend certain convergence properties (instead of looking at "worse case dynamics" we might content ourself with convergence for dynamics in which a randomly chosen player moves). Also, is it possible to extend incentive compatibility if players do not always play a best-response? Can we exploit these results to design new "greedy-type" truthful mechanisms?

References

[APP09] V. Auletta, P. Penna, and G. Persiano. Private capacities in mechanism design. In *34th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 5734 of *LNCS*, pages 112–123, 2009.

[AT01] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 482–491, 2001.

[LS09] R. Lavi and C. Swamy. Truthful mechanism design for multidimensional scheduling via cycle monotonicity. *Games and Economic Behavior*, 67(1):99 – 124, 2009.

[NR01] N. Nisan and A. Ronen. Algorithmic Mechanism Design. *Games and Economic Behavior*, 35:166–196, 2001.

[NSVZ11] N. Nisan, M. Schapira, G. Valiant, and A. Zohar. Best-response mechanisms. In *Proc. of 2nd Symp. on Innovations in Computer Science (ICS)*, 2011.

[PSSW09] P. Penna, F. Schoppmann, R. Silvestri, and P. Widmayer. Pseudonyms in cost-sharing games. In *Proc. of the 5th International Workshop on Internet and Network Economics (WINE)*, volume 5929 of *LNCS*, pages 256–267, 2009.

[PV10] E. Pountourakis and A. Vidali. A complete characterization of group-strategyproof mechanisms of cost-sharing. In *Proc. of European Symposium on Algorithms (ESA)*, volume 6346 of *LNCS*, pages 146–157, 2010.

[Sch00] J. Schummer. Manipulation through bribes. *Journal of Economic Theory*, 91(3):180–198, 2000.

[TIY10] T. Todo, A. Iwasaki, and M. Yokoo. False-name-proofness in facility location problem on the real line. In *Proc. of the 6th Internet and Network Economics (WINE)*, pages 559–562. Springer, 2010.

[YSM04] M. Yokoo, Y. Sakurai, and S. Matsubara. The Effect of False-name Bids in Combinatorial Auctions: New Fraud in Internet Auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.