Games and Economic Behavior 86 (2014) 491–509

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Optimal collusion-resistant mechanisms with verification ☆

Paolo Penna [a], Carmine Ventre [b,*]

[a] Dipartimento di Informatica ed Applicazioni "R.M. Capocelli", Università di Salerno, via Ponte Don Melillo, I-84084 Fisciano (SA), Italy
[b] School of Computing, Teesside University, Borough Road, Middlesbrough, TS1 3BA, UK

ARTICLE INFO

Article history:
Received 27 June 2011
Available online 17 September 2012

JEL classification:
C70
D44
D82

Keywords:
Game theory
Algorithmic mechanism design
Transferable utilities

ABSTRACT

We present the first general positive result on the construction of *collusion-resistant* mechanisms, that is, mechanisms that guarantee dominant strategies even when agents can form arbitrary coalitions and exchange compensations (sometimes referred to as *transferable utilities* or *side payments*). This is a much stronger solution concept as compared to truthful or even group strategyproof mechanisms, and only impossibility results were known for this type of mechanisms in the "classical" model.

We describe collusion-resistant mechanisms *with verification* that return *optimal solutions* for a wide class of mechanism design problems (which includes utilitarian ones as a special case). Note that every collusion-resistant mechanism *without verification* must have an *unbounded* approximation factor and, in general, optimal solutions cannot be obtained even if we content ourselves with truthful ("non-collusion-resistant") mechanisms. All these results apply to problems that have been extensively studied in the algorithmic mechanism design literature like, for instance, task scheduling and inter-domain routing.

1. Introduction

Computations over the Internet often involve self-interested parties which may manipulate the protocol by misreporting some information. For instance, the protocol may be designed to route packets so that no router is congested. The protocol takes its decision based on the information reported by the owners of the routers (in the Internet protocol these are the so-called Autonomous Systems, Feigenbaum et al., 2005) who might have an incentive in exaggerating their costs for forwarding packets (the protocol gives them less traffic or a higher compensation). In a *mechanism design* approach (Nisan and Ronen, 2001) the algorithm used by the protocol is augmented with suitable payments which are supposed to motivate the parties involved (called *selfish agents*) to report the true information (their private costs or *types*).

A mechanism is said *truthful* if it guarantees that each agent can maximize her utility by truthfully reporting her type. Unfortunately, truthful mechanisms can be manipulated by colluding agents (several agents acting as a single one). This behavior is very common and thus one needs *collusion-resistant* mechanisms: no coalition of agents can raise the utilities of its members even if they can exchange compensations (Schummer, 2000; Goldberg and Hartline, 2005).

The design of truthful mechanisms is already a challenging problem; general constructions are known only for the class of so-called *utilitarian* problems (roughly speaking, the goal is to minimize the *sum* of all agents' costs). Many problems of practical interest are *not* utilitarian and truthful mechanisms for them do not exist, even if we consider approximate

☆ Work partially supported by the European Union under IST FET Integrated Project AEOLUS (IST-015964). A preliminary and partial version of this paper appeared in ACM EC'09.

* Corresponding author.
E-mail addresses: penna@dia.unisa.it (P. Penna), C.Ventre@tees.ac.uk (C. Ventre).

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2012.09.002

solutions and mechanisms running in time exponential in the bit length of the input (these inapproximability results stem from "incentive compatibility" considerations and not from computational ones).

The class of collusion-resistant mechanisms is even more limited as it essentially consists of only trivial mechanisms that output a fixed solution (Schummer, 2000; Goldberg and Hartline, 2005): these mechanisms ignore the input and thus are of no use (their approximation guarantee is *unbounded*).

1.1. Our contribution

This work presents the first general positive result on the construction of collusion-resistant mechanisms.[1] This becomes possible by considering so-called mechanisms *with verification*, that is, mechanisms that gain a limited information on the agents' types based on the observation of the computed solution. Mechanisms with verification are introduced in Nisan and Ronen (2001) and further developed in Auletta et al. (2009, 2006), Ventre (2006), Penna and Ventre (2008) where a generalization of the original notion is considered. The model of verification we consider obeys this more general definition (the interested reader may refer to Appendix C where we highlight the differences and show how the new model is more general than the original one) and is presented in Section 2 where we give examples and formal definitions.

The main result of this work is a general construction of mechanisms with verification which are *collusion-resistant* and return *optimal solutions* for a *generalization of utilitarian problems* (see Theorem 9). Under a mild assumption on the domain (an upper and a lower bound on the costs), these new mechanisms "kill two birds with one stone" because we can deal with non-utilitarian problems and guarantee the strongest solution concept simultaneously.

To get a feel of the implications of this result, we compare the performance of these new mechanisms with verification against the "classical" mechanism design approach, which in the sequel we denote to as mechanisms *without verification*. To this aim, we consider one of the most studied problems in (algorithmic) mechanism design, that is, scheduling selfish unrelated machines (Nisan and Ronen, 2001; Koutsoupias and Vidali, 2007; Mu'alem and Schapira, 2007; Gamzu, 2007):

Without verification, collusion-resistant mechanisms must have an *unbounded* approximation ratio even for two machines. The best known truthful ("non-collusion-resistant") mechanism is only $O(n)$-approximate, n being the number of machines.

With verification, there exists a collusion-resistant mechanism whose approximation factor is $(1+\varepsilon)$, for any $\varepsilon > 0$.

We give several other applications of our result to problems studied in the literature. The results are similar to those above with the notable fact that, in some cases, we obtain *polynomial-time* mechanisms, while the impossibility results (e.g., unbounded approximation ratio for collusion-resistant mechanisms) holds for mechanisms that are computationally unbounded (see Table 1).

Our construction relies on payments similar to those used in VCG mechanisms (Vickrey, 1961; Clarke, 1971; Groves, 1973), and it gives an explicit and computationally-efficient way of constructing the whole mechanism (see Corollary 12). A well-known drawback of VCG mechanisms is that of being far from frugal (see, e.g., Archer and Tardos, 2007) and therefore our mechanisms are not frugal as well. In this work we do not consider frugality issues directly (that is, how much the mechanism overpays the agents) since our aim is that of giving a *general* technique for the construction of collusion-resistant mechanisms. The optimality of the payments is an important issue *in general* since even truthful (not necessarily VCG) mechanisms must have large payments for rather simple problems (Elkind et al., 2004; Karlin et al., 2005).

The conditions for obtaining our mechanisms are stated in terms of certain *algorithmic properties* so that the design of the entire mechanism reduces to the design of an algorithm that fulfills these conditions. Theorem 9 says that every algorithm that is optimal for what we call a *weakly utilitarian* cost function (see Definition 7) can be turned into a collusion-resistant mechanism with verification. The same result can be obtained if the algorithm satisfies a condition independent from the optimality (see Theorem 22). This poses a new algorithmic challenge: designing *computationally-efficient* algorithms that obey these conditions. This work provides the first positive results in this direction.

1.2. Related work

Every utilitarian problem admits an exact truthful mechanism (without verification) via VCG mechanisms (Vickrey, 1961; Clarke, 1971; Groves, 1973). The inapproximability of non-utilitarian problems has been extensively studied (Nisan and Ronen, 2001; Archer and Tardos, 2001; Christodoulou et al., 2007; Mu'alem and Schapira, 2007; Gamzu, 2007; Lavi and Swamy, 2009; Christodoulou et al., 2010). These works prove lower bounds on the approximation ratio that *any* truthful mechanism without verification must have on the problem(s). These lower bounds hold for exponential-time mechanisms and for finite domains, a special case of the bounded domains studied here.

Mechanisms with verification are introduced in Nisan and Ronen (2001) in order to overcome these impossibility results and the limitations of VCG mechanisms. The model of verification is then generalized and truthful constructions are

[1] Goldberg and Hartline (2005) describe several mechanisms that achieve a weaker "collusion-resistant with high probability" condition for a simple "unstructured" problem in which identical goods are sold to the users.