Journal of Computer and System Sciences 78 (2012) 2–14

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Computing optimal outcomes under an expressive representation of settings with externalities ☆

Vincent Conitzer [a,*], Tuomas Sandholm [b]

[a] Department of Computer Science, Duke University, Durham, NC 27708, USA
[b] Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA

ARTICLE INFO

Article history:
Received 3 August 2009
Received in revised form 29 August 2010
Accepted 10 February 2011
Available online 15 March 2011

Keywords:
Expressive markets
Externalities
Representation

ABSTRACT

When a decision must be made based on the preferences of multiple agents, and the space of possible outcomes is combinatorial in nature, it becomes necessary to think about how preferences should be represented, and how this affects the complexity of finding an optimal (or at least a good) outcome. We study settings with *externalities*, where each agent controls one or more variables, and how these variables are set affects not only the agent herself, but also potentially the other agents. For example, one agent may decide to reduce her pollution, which will come at a cost to herself, but will result in a benefit for all other agents. We formalize how to represent such domains and show that in a number of key special cases, it is NP-complete to determine whether there exists a nontrivial feasible solution (and therefore the maximum social welfare is completely inapproximable). However, for one important special case, we give an algorithm that converges to the solution with the maximal concession by each agent (in a linear number of rounds for utility functions that additively decompose into piecewise constant functions). Maximizing social welfare, however, remains NP-hard even in this setting. We also demonstrate a special case that can be solved in polynomial time using linear programming.

1. Introduction

In many domains, a decision needs to be made based on the preferences of multiple agents. Often, the space of possible outcomes is combinatorial in nature, so that it becomes necessary to consider how preferences should be represented, as well as to design algorithms for finding an optimal (or at least a good) outcome. (For a recent overview of such work, see [7].)

Combinatorial auctions (for an overview, see [12]) are a common example. In such an auction, there are multiple items to be allocated among the agents, so an outcome is defined by a specification of which bundle of items each agent gets (plus, perhaps, payments to be made by or to the agents). Variants such as *combinatorial reverse auctions* (where the auctioneer seeks to *procure* a set of items) and *combinatorial exchanges* (where the agents trade items among themselves) have also received attention [30,38,29,3,36].

A pervasive assumption in this work (with a very recent exception [25]) has been that there are *no allocative externalities*: that is, no agent cares what happens to an item unless that agent herself receives the item. This is insufficient to model

☆ A short, early conference version appeared in the proceedings of the 2005 AAAI conference. This work has been supported by the National Science Foundation under grants IIS-0121678, IIS-0427858, IIS-0812113, IIS-0905390, IIS-0964579, and CAREER IIS-0953756; it has also been supported by two Sloan Fellowships.

* Corresponding author.
E-mail addresses: conitzer@cs.duke.edu (V. Conitzer), sandholm@cs.cmu.edu (T. Sandholm).

0022-0000/$ – see front matter
doi:10.1016/j.jcss.2011.02.009

situations where a bidder who does not win an item still cares which other bidder wins it—for example, this may be the case if the item is a nuclear weapon [22]. Recently, some work in *sponsored-search auctions*, where multiple advertisement slots on a page are for sale, has started to consider the role of externalities [31,1,15,17,24,16,18,33,11]. Here, the idea is that the attention that the user pays to one ad can depend on which other ads are shown simultaneously.

More generally, and more closely related to this paper, there are many important domains where actions taken by one agent affect many other agents. For example, if one agent takes on a task, such as building a bridge, many other agents may benefit from this (and the extent of their benefit in general depends on how the bridge is built, for example, on how heavy a load it can support). Similarly, if a company reduces its pollution, many individuals may benefit, even if they have nothing to do with the goods that the company produces. An action's effect on an otherwise uninvolved agent's utility is commonly known as an *externality* (for a discussion, see [26]). When making decisions based on the preferences of multiple agents, externalities must be taken into account, so that (potentially complex) arrangements can be made that are truly to every agent's benefit.

One domain in which externalities play a fundamental role, and that fits in the framework described in this paper, is the design of expressive mechanisms for donating to (say) charitable causes [9,14]. The basic idea here is as follows. If one agent donates to a charity, then another agent who also cares about this charity benefits. For that reason, it may happen that, even if each individual agent does not care enough about the charity to give money to it by herself, it is nevertheless possible that all of the agents prefer a joint arrangement in which each agent gives a certain amount to the charity. This is because, thanks to the arrangement, each individual agent's donation is effectively multiplied by the number of agents. This opens up the possibility of mechanisms that take everyone's preferences over the charities as input, and then determine an arrangement for how much each agent should pay. Externalities play a fundamental role here: an agent giving to a charity imposes an externality on the other agents who care about this charity, and this is why the agents can benefit from a joint arrangement.

In this paper, we study whether optimal (or at least good) outcomes can be efficiently computed, under a quite general representation of settings with externalities. To our knowledge, this is the first such study of a general representation of settings with externalities. A common objective is to maximize *social welfare*, which is the sum of the agents' utilities. However, in most settings, there are constraints that must be satisfied. Typically, there are *voluntary participation* constraints, meaning that no agent is made worse off by participating in the mechanism. Additionally, if only the agents themselves know their preferences, and the agents are self-interested (the setting of *mechanism design*), then there may be *incentive compatibility* constraints, meaning that no agent should be able to make herself better off by misreporting her preferences.[1]

After introducing our basic representation scheme for settings with externalities, we study the computational complexity of the following problem: given the agents' preferences, find a good (if possible, an optimal) outcome that satisfies the voluntary participation constraints. This problem is analogous to the *winner determination* problem in combinatorial auctions and exchanges, which consists of finding an optimal allocation of the items, given the bids. The winner determination problem in combinatorial auctions and exchanges has received a tremendous amount of previous attention (for example [34, 13,35,38,5,8,39,10,20,19,4]). In this paper, we will mostly focus on restricted settings that cannot model, *e.g.*, fully general combinatorial auctions and exchanges, so that we do not inherit all of the complexities from those settings (which would trivialize our results). Also, in this first research on the topic, we do not consider any incentive compatibility constraints—that is, we take the agents' reported preferences at face value. This is reasonable when the agents' preferences are common knowledge; when there are other reasons to believe that the agents report their preferences truthfully (for example, for ethical reasons, or because the party reporting the preferences is concerned with the global welfare rather than the agent's individual utility)[2]; or when we are simply interested in finding outcomes that are good relative to the reported preferences (for example, because we are an optimization company that gets rewarded based on how good the outcomes that we produce are relative to the reported preferences). Nevertheless, we believe that incentive compatibility is an important topic for future research, and we will discuss it at the end of the paper in Section 9. As we noted, we do impose voluntary participation constraints.

The rest of this paper is organized as follows. In Section 2, we define our representation and the basic problems that we study under this representation. We show that the problem of finding a nontrivial feasible solution is hard in a number of special cases, including when each agent controls only one variable (Section 3); when there are only negative externalities and each agent controls at most two variables (Section 4); and when there are only negative externalities and there are only two agents, but there is no constraint on how many variables they control (Section 7). In Section 5, we give an algorithm for the case where there are only negative externalities and each agent controls only one variable. Under minimal assumptions, this algorithm finds or converges to the feasible outcome with the "maximal concessions" by each agent; moreover, given some additional assumptions (under which the hardness results proven in other sections still hold), the algorithm requires only a linear number of rounds. Nevertheless, in Section 6, we show that finding the social welfare maximizing outcome remains hard even in this setting. Finally, in Section 8, we show that the social welfare maximizing outcome can be found in polynomial time using linear programming if all the utility functions are piecewise linear and concave.

[1] Very recent research has studied the relationship between expressiveness and social welfare [2].

[2] For example, in a large organization, when a representative of a department within the organization is asked what the department's needs are, it is possible that this representative acts in the organization's best interest, rather than the department's.