

Combinatorial Auctions with Externalities

(Extended Abstract)

Piotr Krysta* Tomasz Michalak⁺ Tuomas Sandholm[†] Michael Wooldridge*

*Computer Science
University of Liverpool
Liverpool L69 3BX, UK

⁺Electronics & Computer Science
University of Southampton
Southampton SO17 1BJ, UK

[†]Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA

ABSTRACT

Although combinatorial auctions have received a great deal of attention from the computer science community over the past decade, research in this domain has focussed on settings in which a bidder only has preferences over the bundles of goods they themselves receive, and is indifferent about how other goods are allocated to other bidders. In general, however, bidders in combinatorial auctions will be subject to *externalities*: they care about how the goods they are not themselves allocated are allocated to others. Our aim in the present work is to study such *combinatorial auctions with externalities* from a computational perspective. We first present our formal model, and then develop a classification scheme for the types of externalities that may be exhibited in a bidder's valuation function. We develop a bidding language for combinatorial auctions with externalities, which uses weighted logical formulae to represent bidder valuation functions. We then investigate the properties of this representation: we study the complexity of the winner determination problem, and characterise the complexity of classifying the properties of valuation functions. Finally, we consider approximation methods for winner determination.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

General Terms

Theory

Keywords

auctions, externalities, complexity, winner determination

1. INTRODUCTION

Combinatorial auctions have been closely studied over the past decade [1, 3]. In a combinatorial auction, a number of goods are simultaneously put to auction, and agents can submit bids for bundles of goods. Within the computer science/AI literature, four main aspects of combinatorial auctions have been considered: *bidding languages*, where the goal is to design compact, expressive, natural, and computationally tractable languages for defining bidder valuation functions; *mechanism design*, where the goal is typically to design bidding, allocation, and payment schemes so that bidders

Cite as: Combinatorial Auctions with Externalities (Extended Abstract), P. Krysta, T. Michalak, T. Sandholm, and M. Wooldridge, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

are incentivised to truthfully report their valuation function; *winner determination*, where the goal is typically to compute efficiently a social welfare-maximising allocation of goods to bidders, given a representation of bids/preferences [4]; and *preference elicitation*, where the goal is to elicit efficiently a valuation function from a potential bidder.

Although details differ, a common model for such combinatorial auctions is the following. We have a set \mathcal{Z} of goods to be auctioned to agents $\mathcal{N} = \{a_1, \dots, a_n\}$, and each agent $a_i \in \mathcal{N}$ has preferences represented by a *valuation function*, $v_i : 2^{\mathcal{Z}} \rightarrow \mathbb{R}$, assigning a numeric value to every possible bundle of goods. Implicit within this framework is a rather significant (and arguably rather unrealistic) assumption: that *bidders only have preferences over the allocation of goods that they receive, and are indifferent about how other goods are allocated to other agents*. This point is very well-known in the economics literature, where the term *externality* is used to describe the effect that a transaction has on an individual that is not directly involved in the transaction. If the individual is adversely affected by the transaction, then the externality is said to be *negative*, while if the individual benefits from the transaction, then the externality is *positive*. In a combinatorial auction with externalities, bidders have preferences not just over the bundles of goods *they* receive, but also over the way in which other goods are allocated to *others*. This holds even in the extreme case where a bidder is allocated *no* goods: they may still have preferences over the way in which goods are allocated to others, and if the externalities are sufficiently severe, such a bidder may even be motivated to pay the auctioneer to prevent another agent being allocated some good, even though they themselves are allocated nothing. In this work, we begin to consider the computational aspects of combinatorial auctions with externalities.

2. THE FRAMEWORK

We start by assuming a finite, non-empty set $\mathcal{Z} = \{z_1, \dots, z_m\}$ of *atomic goods*. We assume these goods are indivisible and that each good is unique. Next, we assume a finite, non-empty set $\mathcal{N} = \{a_0, a_1, \dots, a_n\}$ of *agents* (a.k.a. bidders). An *allocation* is a function $\alpha : \mathcal{N} \rightarrow 2^{\mathcal{Z}}$ s. t. $\alpha(a_1), \dots, \alpha(a_n)$ partitions \mathcal{Z} . The intended interpretation is that $\alpha(a_i)$ is the set of goods allocated to agent a_i under allocation α . Let \mathcal{A} denote the set of all possible allocations over \mathcal{N}, \mathcal{Z} .

In the literature on combinatorial auctions, a *valuation function* for an agent $a_i \in \mathcal{N}$ is usually understood as a function $v_i : 2^{\mathcal{Z}} \rightarrow \mathbb{R}$, i.e., a function that gives the value $v_i(Z)$ to agent $a_i \in \mathcal{N}$ of the bundle of goods $Z \subseteq \mathcal{Z}$. Implicit in such a definition of valuation functions is the idea that a valuation depends *only* on the goods that are allocated to a_i , and not on the way that goods are allocated to other agents. In the present paper, we will be concerned with valu-

ation functions for agents that take into account not just the goods allocated to that agent, but also the way that goods are allocated to others. Thus, for our purposes, a valuation for agent $a_i \in \mathcal{N}$ is a function $v_i : \mathcal{A} \rightarrow \mathbb{R}$.

Bringing the above components together, we say a combinatorial auction with externalities is a tuple

$$\langle \mathcal{Z}, \mathcal{N}, v_1, \dots, v_n \rangle$$

where \mathcal{Z} is the set of goods, \mathcal{N} is the set of agents, and $v_i \in \mathcal{V}$ is the valuation function for agent $a_i \in \mathcal{N}$.

The WINNER DETERMINATION problem in this setting is analogous to conventional combinatorial auctions: the aim is to find an allocation α^* that maximizes *social welfare*:

$$\alpha^* = \arg \max_{\alpha \in \mathcal{A}} \sum_{a_i \in \mathcal{N}} v_i(\alpha).$$

3. OUR RESULTS: AN OVERVIEW

Classifying Valuation Functions: The first issue we considered was the development of a classification scheme for assignments and valuation functions. For example, a key property we considered was whether or not a valuation function v_i for agent i could be considered *free* of externalities. Intuitively, a valuation function v_i is externality free if it assigns the same value to all allocations in which agent i receives the same goods. Formally, $v_i : \mathcal{A} \rightarrow \mathbb{R}$ is externality free iff:

$$\forall \alpha_1, \alpha_2 \in \mathcal{A} : \alpha_1(a_i) = \alpha_2(a_i) \text{ implies } v_i(\alpha_1) = v_i(\alpha_2).$$

Our classification scheme considers a wide range of different properties of valuation functions, relating to the types of externalities that might be present. For example, we considered valuation functions that depend only on the goods that are allocated to a particular group of agents, but not about how goods are allocated within that group; and similarly, the idea of a valuation function depending on who a particular set of goods is allocated to, and so on.

A Bidding Language: The next issue we considered was how to succinctly represent valuation functions $v_i : \mathcal{A} \rightarrow \mathbb{R}$. We developed a *weighted rule* bidding language for combinatorial auctions with externalities, which derives inspiration from the *weighted formula* representations that have been used in other areas of AI (cf. [2]). We specify agent a_i 's valuation function v_i as a set of rules \mathcal{R}_i , with each rule taking the form (*condition, value*), where *condition* is a logical predicate over allocations \mathcal{A} , and *value* $\in \mathbb{R}_+$. To obtain the value of an allocation α given a set of rules \mathcal{R} , we sum the values x of all the rules (φ, x) in \mathcal{R} whose condition φ is satisfied by α .

The conditions of our rules are essentially conventional propositional logic, with primitive propositions replaced by expressions for referring to allocations. These expressions are of the form $a_i : z$, where $a_i \in \mathcal{N}$ is an agent and $z \in \mathcal{Z}$ is a good. The intended interpretation of the expression $a_i : z$ is, naturally enough, that agent a_i is allocated good z . These atomic expressions may be combined with classical Boolean connectives to form complex conditions. We write $\alpha \models \varphi$ to mean that the allocation α satisfies the condition φ . A *rule* is a pair (φ, x) where φ is a condition and $x \in \mathbb{R}_+$. A set of rules \mathcal{R} induces a valuation function $v_{\mathcal{R}}$:

$$v_{\mathcal{R}}(\alpha) = \sum_{(\varphi, x_i) \in \mathcal{R} \& \alpha \models \varphi} x_i$$

Classifying Valuation Functions: Given that we have a classification of properties of valuation functions for combinatorial auctions with externalities, and a concrete representation for such valuation

functions, it is natural to ask how hard it is to check whether a particular valuation function exhibits a particular property. For example, consider the decision problem EXTERNALITY FREENESS, where we are given a rule set \mathcal{R} (over \mathcal{Z}, \mathcal{N}), and we are asked whether $v_{\mathcal{R}}$ is free of externalities, as defined above. We showed that in general, it is co-NP-complete to check whether or not a valuation function is externality free. We then went on to consider the other properties of valuation functions as considered in our classification scheme: it transpires that co-NP-completeness is the characteristic complexity class of such classification problems.

Complexity of Winner Determination: Recall the winner determination problem, as defined above. The computational complexity of this problem depends upon the choice of representation for valuation functions. Assuming the weighted rule based bidding language that we describe above, we proved that WINNER DETERMINATION can be solved in polynomial time with a polynomial number of queries to an NP-oracle, and moreover that the problem is as hard as any function problem that can be solved in polynomial time with a polynomial number of queries to an NP-oracle. We also identified a case where the number of NP-oracle queries required to solve the problem is “small” compared to the input size ($\log_2 n$ queries, where n is the total number of rules).

Winner Determination with ILPs: Since winner determination for the weighted rule representation is an NP-hard optimisation problem, it is natural to consider approaches to solving such problems. *Integer Linear Programming* (ILP) is one of the most successful and widely-used practical approaches to solving computationally complex optimization problems. We developed an ILP approach to exact winner determination.

Approximation Algorithms: An obvious question is whether it is possible to find a polynomial time approximation algorithm for winner determination, assuming our rule-based representation of valuation functions. We showed that, in general, such approximation algorithms are not possible for the weighted rule representation. However, we identified sub-classes of weighted rules for which approximations are possible.

Acknowledgements: Sandholm was supported by US National Science Foundation grant IIS-0905390. Krysta was supported by EPSRC grant EP/F069502/1. Michalak and Wooldridge were supported by EPSRC grant GR/S62727/01.

4. REFERENCES

- [1] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. The MIT Press: Cambridge, MA, 2006.
- [2] S. Jeong and Y. Shoham. Marginal contribution nets: A compact representation scheme for coalitional games. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC'05)*, Vancouver, Canada, 2005.
- [3] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 209–242. Cambridge University Press: Cambridge, England, 2007.
- [4] T. Sandholm. Optimal winner determination algorithms. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, pages 337–368. The MIT Press: Cambridge, MA, 2006.