

# Making VCG More Robust in Combinatorial Auctions via Submodular Approximation

Makoto Yokoo and Atsushi Iwasaki

Dept. of Intelligent Systems  
Kyushu University  
744 Motoooka, Nishi-ku, Fukuoka, 819-0395 Japan  
yokoo/iwasaki@is.kyushu-u.ac.jp

## Abstract

The Vickrey-Clarke-Groves (VCG) protocol is a theoretically well-founded protocol that can be used for combinatorial auctions. However, the VCG has several limitations such as (a) vulnerability to false-name bids, (b) vulnerability to loser collusion, and (c) the outcome is not in the core. Yokoo, Matsutani, & Iwasaki (2006) presented a new combinatorial auction protocol called the Groves Mechanism with SubModular Approximation (GM-SMA). This protocol satisfies the following characteristics: (1) it is false-name-proof, (2) each winner is included in a Pareto efficient allocation, and (3) as long as a Pareto efficient allocation is achieved, the protocol is robust against the collusion of losers and the outcome is in the core. The GM-SMA is the first protocol that satisfies all three of these characteristics. The basic ideas of the GM-SMA are as follows: (i) it is based on the VCG protocol, i.e., the payment of a winner in this protocol is identical to the payment in one instance of the Groves mechanism, which is a class of protocols that includes the VCG, (ii) when calculating the payment of a bidder, we approximate the valuations of other bidders by using a submodular valuation function (submodular approximation).

This paper shows a high-level presentation of the GM-SMA protocol, and discusses open problems and the relationship to other works in AI.

## Introduction

Internet auctions have become an integral part of Electronic Commerce and a promising field for applying AI technologies. Among various studies related to Internet auctions, those on combinatorial auctions have lately attracted considerable attention (an extensive survey is presented in (de Vries & Vohra 2003)). Although conventional auctions sell a single item at a time, combinatorial auctions sell multiple items with interdependent values simultaneously and allow the bidders to bid on any combination of items. In a combinatorial auction, a bidder can express complementary/substitutable preferences over multiple bids. By taking into account such preferences, we can increase the participants' utilities and the revenue of the seller.

One important characteristic of an auction protocol is that it is *strategy-proof*. A protocol is strategy-proof if, for each

bidder, declaring his/her true valuation is a *dominant strategy*, i.e., an optimal strategy regardless of the actions of other bidders. In theory, the revelation principle states that in the design of an auction protocol, we can restrict our attention to strategy-proof protocols without loss of generality.

The Vickrey-Clarke-Groves (VCG) protocol (Vickrey 1961; Clarke 1971; Groves 1973) is a strategy-proof protocol that can be applied to combinatorial auctions. The VCG protocol satisfies Pareto efficiency. Under several natural assumptions, the VCG protocol is the only strategy-proof, Pareto efficient protocol (Holmstrom 1979).

However, the VCG protocol has several limitations, including the problems described below.

**Vulnerability to false-name bids:** The authors have pointed out the possibility of a new type of fraud called false-name bids, which utilizes the anonymity available on the Internet (Yokoo, Sakurai, & Matsubara 2004; 2001). False-name bids are bids submitted under fictitious names, e.g., multiple e-mail addresses. Such a dishonest action is very difficult to detect, since identifying each participant on the Internet is virtually impossible. As shown in (Yokoo, Sakurai, & Matsubara 2004; 2001), the VCG is not robust against false-name manipulations.

**Vulnerability to loser collusion:** As shown in (Ausubel & Milgrom 2002), in the VCG protocol, there is a chance that some bidders, who would be losers if they bid their true valuations, becomes winners and increase their utility if they collude and adjust their bids.

**The outcome is not in the core:** As shown in (Ausubel & Milgrom 2002), in the VCG protocol, there is a chance that the outcome of the auction is not in the core. This means that a seller has an incentive to deviate from the protocol and to sell goods to the bidders who are not winners.

Yokoo, Matsutani, & Iwasaki (2006) presented a new combinatorial auction protocol called the *Groves Mechanism with SubModular Approximation* (GM-SMA), which removes the most of above mentioned problems. In the rest of this paper, we show a high-level presentation of the GM-SMA protocol. Furthermore, we discuss open problems and the relationship to other works in AI.

## Model

Assume there are a set of bidders  $N = \{1, 2, \dots, n\}$  and a set of goods  $M = \{1, 2, \dots, m\}$ . Each bidder  $i$  has his/her preferences over  $B \subseteq M$ . Formally, we model this by supposing that bidder  $i$  privately observes a parameter, or signal,  $\theta_i$ , which determines his/her preferences. We refer to  $\theta_i$  as the *type* of bidder  $i$ . We assume  $\theta_i$  is drawn from a set  $\Theta$ . We also assume a *quasi-linear, private value* model with *no allocative externality*, defined as follows.

### Definition 1 utility of a bidder

The utility of bidder  $i$ , when  $i$  obtains a bundle, i.e., a subset of goods  $B \subseteq M$  and pays  $p_{B,i}$ , is represented as  $v(B, \theta_i) - p_{B,i}$ .

We assume the valuation  $v$  is normalized by  $v(\emptyset, \theta_i) = 0$ . Furthermore, we assume *free disposal*, i.e.,  $v(B', \theta_i) \geq v(B, \theta_i)$  for all  $B' \supseteq B$ .

In a traditional definition (Mas-Colell, Whinston, & Green 1995), an auction protocol is (dominant-strategy) *incentive compatible* (or *strategy-proof*) if declaring the true type/valuation is a dominant strategy for each bidder, i.e., an optimal strategy regardless of the actions of other bidders. In this paper, we use an extended definition of incentive compatibility so that it can address false-name bid manipulations, i.e., we define that an auction protocol is (dominant-strategy) *incentive compatible* if declaring the true type *by using a single identifier* is a dominant strategy for each bidder. To distinguish between the traditional and extended definitions of incentive compatibility, we refer to the traditional definition as *strategy-proof* and to the extended definition as *false-name-proof*.

As in the case of strategy-proof protocols, the revelation principle holds for false-name-proof protocols (Yokoo, Sakurai, & Matsubara 2004). More specifically, if a certain property (e.g., Pareto efficiency) can be achieved using some auction protocol in a dominant-strategy equilibrium, i.e., a combination of dominant strategies of bidders, the property can also be achieved using a false-name-proof protocol. Thus, we can restrict our attention to false-name-proof protocols without loss of generality.

We say an auction protocol is *Pareto efficient* when the sum of all participants' utilities (including that of the auctioneer), i.e., the social surplus, is maximized in a dominant-strategy equilibrium. We have proved that there exists no false-name-proof protocol that satisfies Pareto efficiency (Yokoo, Sakurai, & Matsubara 2004). Therefore, we need to sacrifice efficiency to some extent when false-name bids are possible.

## VCG Protocol

Next, we briefly outline the VCG protocol. To simplify the protocol description, we introduce the following notation. For a set of goods  $B$  and a set of bidders  $S$ , where  $\Theta_S$  is a set of types of bidders in  $S$ , we define  $V^*(B, \Theta_S)$  as the sum of the valuations of  $S$  when  $B$  is allocated optimally among  $S$ .

To be precise, for a feasible allocation  $g = (B_1, B_2, \dots)$ , where  $\bigcup_{i \in S} B_i \subseteq B$  and for all  $i \neq i', B_i \cap B_{i'} = \emptyset$ ,

$V^*(B, \Theta_S)$  is defined as  $\max_g \sum_{i \in S} v(B_i, \theta_i)$ , where  $\theta_i$  is the type of bidder  $i$ .

The VCG protocol can be described as follows. For simplicity, we describe the protocol assuming each bidder  $i$  declares his/her true type  $\theta_i$ .

1. The auctioneer finds a Pareto efficient allocation. The auctioneer can choose an arbitrary one if there exist multiple Pareto efficient allocation.
2. When  $B$  is allocated to bidder  $i$  in that allocation, bidder  $i$  pays the following VCG price  $p_{B,i}$ :

$$p_{B,i} = V^*(M, \Theta_{N \setminus \{i\}}) - V^*(M \setminus B, \Theta_{N \setminus \{i\}}).$$

Let us describe how this protocol works.

**Example 1** Assume there are two goods 1 and 2, and three bidders, bidder 1, 2, and 3, whose types are  $\theta_1, \theta_2$ , and  $\theta_3$ , respectively. The valuation for a bundle  $v(B, \theta_i)$  is determined as follows.

	{1}	{2}	{1, 2}
$\theta_1$	6	0	6
$\theta_2$	0	0	8
$\theta_3$	0	5	5

The Pareto efficient allocation is allocating good 1 to bidder 1 and good 2 to bidder 3, respectively. The price for bidder 1 is calculated as  $8 - 5 = 3$ . Similarly, the price for bidder 3 is calculated as  $8 - 6 = 2$ .

Now, we can see that this outcome is not in the core. Let us assume another allocation where both goods are sold to bidder 2 at price 7. Then, both the seller and bidder 2 prefer this allocation to the outcome of the VCG protocol. Such a coalition, in this case, the coalition of the seller and bidder 2, is called a *blocking coalition*.

Let us consider another situation.

**Example 2** Assume there are two goods 1 and 2, and two bidders, bidder 1 and 2, whose types are  $\theta_1$  and  $\theta_2$ , respectively. The valuation for a bundle  $v(B, \theta_i)$  is determined as follows.

	{1}	{2}	{1, 2}
$\theta_1$	4	3	7
$\theta_2$	0	0	8

The Pareto efficient allocation is allocating both goods to bidder 2. The price for bidder 2 is calculated as  $7 - 0 = 7$ .

Now, we can see that the VCG protocol is not false-name-proof. Bidder 1 can split his/her bid using another identifier, i.e., bidder 3, so that the situation becomes identical to Example 1. Then, bidder 1 can obtain both goods by paying  $2 + 3 = 5$ , which is less than his/her valuation. Thus, bidder 1 can increase his/her utility by using a false-name bid.

Let us consider another situation.

**Example 3** Assume there are two goods 1 and 2, and three bidders, bidder 1, 2, and 3, whose types are  $\theta_1, \theta_2$ , and  $\theta_3$ , respectively. The valuation for a bundle  $v(B, \theta_i)$  is determined as follows.

	{1}	{2}	{1, 2}
$\theta_1$	4	0	4
$\theta_2$	0	0	8
$\theta_3$	0	3	3

The Pareto efficient allocation is allocating both goods to bidder 2. The price for bidder 2 is calculated as  $7 - 0 = 7$ .

Now, we can see that the VCG protocol is vulnerable to the collusion of losers. If bidder 1 and bidder 3 collude and increase their valuations, they can create a situation identical to Example 1. Then, bidder 1 obtains good 1 at price 3, and bidder 3 obtains good 2 at price 2. For each bidder, this price is less than his/her valuation, thus bidders can increase their utilities by engaging in collusion.

## Groves Mechanism with SubModular Approximation (GM-SMA)

### Basic Idea

The basic idea behind the GM-SMA protocol can be described as follows. It is well-known that the valuations of bidders satisfy a property called *submodularity*, the above mentioned limitations of the VCG protocol disappear (Ausubel & Milgrom 2002). The valuations of a set of bidders  $N$  is called submodular if the following condition holds: for all  $S \subseteq N, B', B'' \subseteq M, V^*(B', \Theta_S) + V^*(B'', \Theta_S) \geq V^*(B' \cup B'', \Theta_S) + V^*(B' \cap B'', \Theta_S)$ .

When  $B'$  and  $B''$  are mutually exclusive, this condition means that the sum of the valuations of having  $B'$  and  $B''$  separately is always larger than (or at least equal to) the valuation of having  $B'$  and  $B''$  together, i.e., there is no complementarity.

However, the submodularity condition is not satisfied in general. As a matter of fact, the main motivation of using a combinatorial auction is the existence of complementary goods.

The basic idea of the GM-SMA protocol is, if the valuations of bidders are not submodular, we approximate the valuations as a submodular function and use that function to calculate the payments in the VCG. As a result, the protocol has a similar property to the VCG when the valuations are submodular.

### Protocol Description

To describe the GM-SMA protocol, we first define a concept called *submodular approximation*.

**Definition 2** *submodular approximation*  $U^*$

The function  $U^*$  defined below is called a *submodular approximation* of  $V^*$ .

For a feasible allocation  $g = (B_1, B_2, \dots)$ , where  $\bigcup_{i \in S} B_i \subseteq B$  and for all  $i \neq i', B_i \cap B_{i'} = \emptyset, U^*(B, \Theta_S)$  is defined as  $\max_g \sum_{i \in S} v'(B_i, \theta_i)$ .  $v'$  is a function that satisfies  $v'(B, \theta_i) \geq v(B, \theta_i)$  for all  $i, B$ .

Furthermore, we assume  $v'$  is chosen so that  $U^*$  becomes submodular, i.e., for all  $S \subseteq N, B', B'' \subseteq M, U^*(B', \Theta_S) + U^*(B'', \Theta_S) \geq U^*(B' \cup B'', \Theta_S) + U^*(B' \cap B'', \Theta_S)$  holds.

From this definition,  $U^*$  satisfies the following conditions:

- For all  $S \subseteq N, B \subseteq M, U^*(B, \Theta_S) \geq V^*(B, \Theta_S)$  holds.
- For all  $S \subseteq N, B \subseteq M, U^*(M, \Theta_N) \geq U^*(B, \Theta_S) + U^*(M \setminus B, \Theta_{N \setminus S})$  holds.

The Groves Mechanism with SubModular Approximation (GM-SMA) is described as follows.

1. The auctioneer finds a Pareto efficient allocation. The auctioneer can choose an arbitrary one if there exist multiple Pareto efficient allocation.
2. When  $B$  is allocated to bidder  $i$  in that allocation, the auctioneer calculates the GM-SMA price  $p_{B,i}$  defined as follows:
  - When  $B = \emptyset, p_{B,i} = 0$ .
  - Otherwise,  $p_{B,i} = U^*(M, \Theta_{N \setminus \{i\}}) - V^*(M \setminus B, \Theta_{N \setminus \{i\}})$ .

If  $v(B, \theta_i) \geq p_{B,i}$ , then  $B$  is allocated to bidder  $i$  at price  $p_{B,i}$ . If  $v(B, \theta_i) < p_{B,i}$ ,  $B$  is not allocated to any bidder.

Please note that we use  $U^*$  only for the first term. This protocol is quite similar to the Groves mechanism (Groves 1973), since the payment of a winner in this protocol is identical to the payment in one instance of the Groves mechanism. However, it is not an instance of the Groves mechanism because it does not guarantee to achieve a Pareto efficient allocation. This protocol can be considered as a modification of the Groves mechanism, in which a winner has an option not to obtain the goods if the payment becomes too high.

### Examples

Here, we describe how the GM-SMA works in the situation of Example 1. In this example, we use  $v'(\{1\}, \theta_2) = v'(\{2\}, \theta_2) = 8/2 = 4$ , and we use  $v'(\cdot, \cdot) = v(\cdot, \cdot)$  for other bundles and bidders. By choosing  $v'$  in this way,  $U^*$  satisfies the submodular condition.

The Pareto efficient allocation is allocating good 1 to bidder 1 and good 2 to bidder 3, respectively. The price for bidder 1 is calculated as  $(4+5) - 5 = 4$ , where 4 corresponds to  $v'(\{1\}, \theta_2)$  and 5 corresponds to  $v'(\{2\}, \theta_3)$ . Similarly, the price for bidder 3 is calculated as  $(4+6) - 6 = 4$ , where 4 corresponds to  $v'(\{2\}, \theta_2)$  and 6 corresponds to  $v'(\{1\}, \theta_1)$ .

This outcome is in the core. Since the revenue of the seller is 8, the seller does not have an incentive to deviate from the protocol with bidder 2 (who is willing to pay at most 8).

Next, we describe how the protocol works in the situation of Example 2. Here,  $v'$  is defined in a similar way to the previous example. The Pareto efficient allocation is allocating both goods to bidder 2. The price for bidder 2 is calculated as  $7 - 0 = 7$ .

In this case, even if bidder 1 uses a false-name and creates a situation identical to Example 1, his/her utility cannot be positive since bidder 1 must pay  $4 + 4 = 8$  in that situation.

Finally, we describe how the protocol works in the situation of Example 3. Here,  $v'$  is defined in a similar way to

the previous example. The Pareto efficient allocation is allocating both goods to bidder 2. The price for bidder 2 is calculated as  $7 - 0 = 7$ .

In this case, even if bidder 1 and bidder 3 collude and create a situation identical to Example 1, their utilities cannot be positive since each must pay 4 in that situation.

### Characteristics of GM-SMA

The GM-SMA satisfies the following characteristics.

- (1) The GM-SMA is false-name-proof.
- (2) In the GM-SMA, each winner is a bidder who is included in a Pareto efficient allocation<sup>1</sup>.
- (3) In the GM-SMA, as long as the allocation is Pareto efficient, the protocol is robust against the collusion of the losers and the outcome is in the core.

The GM-SMA is the first protocol that satisfies all three of these characteristics. The VCG is not false-name-proof. In addition, even without false-name bids, the VCG protocol does not satisfy characteristic (3).

Furthermore, in existing false-name-proof protocols (Yokoo, Sakurai, & Matsubara 2001), characteristic (2) is not satisfied, i.e., a bidder who is not included in any Pareto efficient allocation becomes a winner quite often. In such a case, the outcome is not in the core. Also, losers who are in a Pareto efficient allocation will have a strong incentive to collude and act as a single bidder.

### Discussions

The efficiency and the revenue of the GM-SMA critically depend on the accuracy of the submodular approximation. Although several simple methods for submodular approximation is introduced in (Yokoo, Matsutani, & Iwasaki 2006), the accuracy of these methods is not satisfactory. Developing better approximation methods is an open problem. Another open problem is to develop protocols that can be applied to similar, but different situations, such as combinatorial procurement auctions.

The winner determination problem in combinatorial auctions has been a popular research topic in AI after pioneering works presented in (Sandholm 1999; Fujishima, Leyton-Brown, & Shoham 1999). Solving the winner determination problem is NP-complete. Various heuristic search and optimization techniques have been introduced for solving the winner determination problem efficiently. When valuations of bidders are submodular, the winner determination problem becomes polynomial time solvable (Lehmann, O'Callaghan, & Shoham 2002). It might be possible to develop a false-name-proof protocol that is computationally efficient by extending the GM-SMA protocol.

Mechanism design is also a popular research topic in AI and multi-agent systems. The VCG protocol is a general

and powerful method in mechanism design. It has been considered as a normative protocol, which might be too complicated to be used in practice, but it is theoretically well-founded. However, as discussed in this paper, the VCG has several limitations. Besides the issues discussed in this paper, many researchers pointed out other limitations of the VCG. For example, Karlin, Kempe, & Tamir (2005) show that the payment can be very high when the VCG is applied to procurement auctions.

### References

- Ausubel, L. M., and Milgrom, P. R. 2002. Ascending auctions with package bidding. *Frontiers of Theoretical Economics* 1(1).
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* 2:19–33.
- de Vries, S., and Vohra, R. V. 2003. Combinatorial auctions: A survey. *INFORMS Journal on Computing* 15.
- Fujishima, Y.; Leyton-Brown, K.; and Shoham, Y. 1999. Taming the computation complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*, 548–553.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41:617–631.
- Holmstrom, B. 1979. Groves' scheme on restricted domains. *Econometrica* 47(5):1137–1144.
- Karlin, A.; Kempe, D.; and Tamir, T. 2005. Beyond VCG: Frugality of truthful mechanisms. In *Proceedings of 46th IEEE Symp. on Foundations of Computer Science*.
- Lehmann, D.; O'Callaghan, L. I.; and Shoham, Y. 2002. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM* 49(5):577–602.
- Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*. Oxford University Press.
- Sandholm, T. 1999. An algorithm for optimal winner determination in combinatorial auction. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*, 542–547.
- Vickrey, W. 1961. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.
- Yokoo, M.; Matsutani, T.; and Iwasaki, A. 2006. False-name-proof combinatorial auction protocol: Groves mechanism with submodular approximation. In *Proceedings of the Fifth International joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2006)*.
- Yokoo, M.; Sakurai, Y.; and Matsubara, S. 2001. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence* 130(2):167–181.
- Yokoo, M.; Sakurai, Y.; and Matsubara, S. 2004. The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior* 46(1):174–188.

<sup>1</sup>Note that the converse is not true, i.e., there is a chance that the GM-SMA could fail to achieve a Pareto efficient allocation. This is inevitable since the GM-SMA is false-name-proof.