

Auction in Dynamic Environments: Incorporating the Cost Caused by Re-allocation

Shigeo Matsubara
NTT Communication Science Laboratories
NTT Corporation
2-4 Hikaridai, Seika-cho, "Keihanna Science City"
Kyoto 619-0237, Japan
matsubara@cslab.kecl.ntt.co.jp

ABSTRACT

This paper proposes an auction protocol for solving a resource allocation problem in dynamic environments. In such environments, the valuation of resources has uncertainty for each bidder, i.e., this valuation depends on the situation not only at the point when the auction is held but also at the point when the allocated resources are actually used. For example, a bidder's valuation in fine weather may be different from that in rainy weather. A solution for dealing with this problem is to execute auctions whenever an event occurs and then to re-allocate resources. Re-allocating resources, however, may cause disutility. Moreover, it does not always provide an equilibrium strategy because it can be viewed as a sequential auction, which means that we cannot accurately predict what outcome will be obtained. To solve this problem, we propose an auction protocol that allows bidders to declare the cost due to re-allocation and then decides an allocation based on this cost of re-allocation as well as the surplus obtained from the allocated resources themselves in the realized situation. We prove that a bidder's truth telling is in equilibrium and that a socially efficient allocation is obtained in the proposed protocol.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*; K.4.4 [Computers and Society]: Electronic Commerce

General Terms

Design, Economics

Keywords

mechanism design

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.
Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.

1. INTRODUCTION

Auctions have been widely studied in the field of multi-agent research to efficiently allocate resources [1]. A major focus in auction studies is asymmetric information among sellers and buyers, i.e., how to elicit private information from participants and find an efficient allocation. Internet and agent-mediated electronic commerce involve uncertainty, which causes various types of problems, e.g., uncertainty of a participant's identity [14], uncertainty of auctioned goods [4], execution failure [9], and so on.

This paper focuses on auctions in dynamic environments. Dynamic environments are another source of uncertainty. A computational agent's behavior in dynamic environments has been actively studied, e.g., how to control an autonomous robot in real-world environments. These planning studies have examined whether an agent should continue to execute an existing plan or behave reactively [11, 5]. A similar problem occurs in resource allocation in auctions.

Consider a meeting room assignment problem. There are multiple seminar organizers contending for particular time-slots, with an auction determining the winner. Here, a valuation of resources often depends on environmental conditions. The number of participants depends on weather, e.g., fine or rainy. Therefore, the seminar organizer who anticipates that the weather will be fine and has won the particular time-slot of the large meeting room may discover after winning that it is going to rain on the day of the seminar, thus lowering the expected turnout. In this case, the valuation of the large meeting room becomes very low.

Re-allocation seems to mitigate this problem. However, re-allocation significantly affects the bidder incentive problem. If we held auctions whenever some environmental conditions change in order to re-allocate resources whose valuations change, a bidder could not always have an equilibrium strategy because this could be viewed as a sequential auction. This causes the problem of a bidder having to examine many strategies to obtain better utility, and the system designer cannot accurately predict what outcome he or she can obtain through an auction.

Another problem of re-allocation is that changing from an existing allocation to another one incurs some cost. For example, if the seminar organizer had announced that the seminar will be held in room 3 in building A but then changes the room on the morning of the seminar day, it may cause some trouble. Here, note that seminar participants have a preference for the change in the room. Their disutility is

small if the newly assigned room is located on the same floor as the previously assigned room. Their disutility becomes larger if the newly assigned room is located on a different floor in the same building, and it becomes much larger if the newly assigned room is located in a different building. Such a participant's preference affects the number of participants, which leads to affecting the organizer's profit.

A solution for reducing the re-allocation cost is to give up the idea of deciding an allocation in advance and obtain a resource allocation on that day. However, this does not allow the organizer to announce the seminar venue in advance, making it difficult to gather many participants.

To solve resource allocation problems in dynamic environments, we have developed a new auction protocol. This protocol allows bidders to declare their preference for re-allocation and make an allocation plan that specifies the initial allocation and a set of transition rules describing when and how to change the allocation. An allocation plan enables a bidder to understand how the allocated resources will be changed in advance.

The contributions of this paper are (1) the proposition of an auction framework that accounts for re-allocation in dynamic environments, (2) the introduction of a new auction protocol that can induce bidders to tell their true valuation and disutility caused by re-allocation by extending the well-established Vickrey-Clarke-Groves (VCG) protocol, and (3) the proof that the protocol can obtain an ex ante Pareto efficient allocation.

Section 2 describes the model of an auction in dynamic environments. In section 3, we propose a new protocol that can avoid the inefficiency caused by environmental changes, and in section 4, we illustrate how our protocol works. Section 5 describes theoretical analysis, section 6 discusses related works, and section 7 concludes the paper.

2. MODEL

This section presents a formal model to enable rigorous discussion. In a trading environment, there exist a seller, i.e., an auctioneer in this paper, and multiple bidders *bidder i* ($i = 1, \dots, n$). The seller may have multiple goods to sell, g_1, g_2, \dots, g_m . This paper assumes the following situation.

ASSUMPTION 1. *There is an interval between the point when the auction is held and the point when the allocated goods are actually used by the winners.*

We assume a private value model, i.e., *bidder i*'s valuation of goods is independent of those of the other bidders. As discussed earlier, the benefit from having a good derives not only from using the good but also from whether it is allocated far in advance or just before its use. In other words, the former can be viewed as the primary benefit obtained as a result of using the good, while the latter can be viewed as the secondary benefit obtained during the interval between an allocation determination and the use of the allocated good.

ASSUMPTION 2. *For bidder i, the valuation is given as follows.*

$$v_i() = v_i^P() + v_i^S()$$

v_i^P represents the benefit obtained as a result of using the allocated good, while v_i^S represents the benefit obtained during the interval between an allocation determination and the use of the allocated good.

v_i^S might represent disutility, since we examine the possibility of re-allocation of goods in this paper. For example, if a seminar organizer wins a room reservation auction but then unfortunately the place of the seminar has changed, the organizer has to inform the seminar participants of the change, which incurs some cost. Thus, we call $v_i^S()$ disutility caused by re-allocation. Note that the organizer can hand a room reservation over to another bidder if the reserved date has not come yet, while the cost of announcing the seminar is a sunk cost.

In addition, we assume that v_i^P depends on the environmental conditions when the allocated goods are actually used by the winner.

ASSUMPTION 3. *The environmental conditions are represented as a set of random variables, $\{cond\} = \{cond_1, cond_2, \dots, cond_l\}$.*

Moreover, we assume the following.

ASSUMPTION 4. *The domain of a random variable is given and an auctioneer and all bidders know the probability distribution p of each random variable.*

For example, the domain of a random variable of weather is {fine, rainy} and a weather report is shared among an auctioneer and all bidders.

The valuation v_i^P of *bidder i* is represented as follows.

ASSUMPTION 5. *Let G denote an allocation of goods and G_i denote allocated goods to bidder i in G .*

For bidder i , the valuation for a bundle of goods G_i is denoted by $v_i^P = v_i^P(G_i; cond_1, cond_2, \dots, cond_l)$.

We use a notation of $v_i^P(G_i)$ as $v_i^P(G_i; cond_1, cond_2, \dots, cond_l)$ if there is no confusion.

We focus on the effect of re-allocation in this paper. Thus, disutility v_i^S of *bidder i* is represented as follows.

ASSUMPTION 6. *For bidder i , v_i^S is denoted by $v_i^S(G_i^{prev}, G_i^{current})$. G_i^{prev} represents a bundle of goods previously allocated to bidder i and $G_i^{current}$ represents bidder i 's current allocation of goods.*

v_i^S might be a function of time, e.g., an elapsed time or a remaining time.

Figure 1 shows an example of a bidder's valuation v_i^P . Bidders are seminar organizers contending for particular time-slots of meeting rooms. g_1 and g_2 represent time-slots of a large meeting room and a small meeting room, respectively. (g_1, g_2) indicates obtaining the time-slots of both rooms simultaneously. In this case, obtaining g_1 or g_2 is sufficient for each bidder, i.e., each bidder has a substitutable preference over the two goods. Here, $\{cond\} = \{weather\}$ and its domain is {fine, rainy}.

Figure 2 shows an example of disutility v_i^S of *bidder 1* caused by the re-allocation itself. Different bidders have different values of disutility. *null* means that no good is allocated. Disutility caused by changing *null* to g_1 represents the disutility that g_1 is not allocated in advance but just before its use, and thus there is not enough time for preparing something to use the allocated good. The disutility caused by changing g_1 to *null* represents disutility, e.g., the cost for announcing that the seminar has to be canceled because a seminar room of g_1 is no longer available. Introducing these disutilities into the discussion is a characteristic of this paper's approach.

	$g1$	$g2$	$(g1, g2)$
<i>bidder 1</i>	10	2	10
<i>bidder 2</i>	4	6	6
<i>bidder 3</i>	3	4	4

(a) fine

	$g1$	$g2$	$(g1, g2)$
<i>bidder 1</i>	5	6	6
<i>bidder 2</i>	7	5	7
<i>bidder 3</i>	6	1	6

(b) rainy

Figure 1: Valuations of goods conditioned on weather

		TO			
		<i>null</i>	$g1$	$g2$	$(g1, g2)$
FROM	<i>null</i>	0	-1	-1	-1
	$g1$	-3	0	-1	0
	$g2$	-3	-1	0	0
	$(g1, g2)$	-3	0	0	0

Figure 2: Disutility of *bidder 1* caused by re-allocation itself

This representation means that a total valuation of goods is determined by the final allocation of the goods, the environmental conditions, and (dis-)utility accumulated corresponding to the re-allocation from the initial allocation to the final allocation. This enables us to easily describe various problem domains, although it might be possible to assume another functional form of valuations.

To maximize social surplus, i.e., the sum of an auctioneer’s and bidders’ utilities, we need to examine the transition process from one allocation to another as well as the allocation itself, i.e., we have to deal with an allocation plan. An allocation plan consists of an initial allocation and a set of transition rules that specify which allocation is to be changed when some values of random variables occur. Figure 3 shows an example of an allocation plan. This means that $(g1, g2) = (\textit{bidder 1}, \textit{bidder 2})$ is an initial allocation and it changes to $(g1, g2) = (\textit{bidder 2}, \textit{bidder 1})$, i.e., *bidder 1* and *bidder 2* swap the rooms, if it rains. If it is fine, the initial allocation is not changed.

initial allocation: $(g1, g2) = (\textit{bidder 1}, \textit{bidder 2})$

final allocation:

$$\begin{aligned} (g1, g2) &= (\textit{bidder 1}, \textit{bidder 2}) \text{ if fine} \\ (g1, g2) &= (\textit{bidder 2}, \textit{bidder 1}) \text{ if rainy} \end{aligned}$$

Figure 3: Example of an allocation plan

Next, we give an expression of utility.

ASSUMPTION 7. *bidder i*’s utility, $u_i()$, is represented as follows.

$$\begin{aligned} u_i() &= v_i(AP) - \textit{payment}_i(AP) \\ &= v_i^P(G_i^{\textit{final}}; \{\textit{cond}\}) + v_i^S(AP) \\ &\quad - \textit{payment}_i(AP) \end{aligned}$$

This is called a quasi-linear utility. AP represents an allocation plan. $G_i^{\textit{final}}$ represents a bundle of goods allocated to *bidder i* in the final allocation.

A method for determining the amount of payment $\textit{payment}_i$ is described in the next section. We assume that a primary benefit depends only on the final allocation and environmental conditions. The disutility of v_i^S is summed up along the transition path of allocations.

If an allocation plan of AP is given, we can calculate the expected utility because we assume that the probability distribution of random variables, p , is known.

3. AUCTION PROTOCOL

In this section, we propose a new protocol that determines an allocation plan. First, we describe desirable properties of an auction protocol and point out that a simple protocol does not satisfy these properties, then we present our protocol.

3.1 Desirable properties of auction protocols

Desirable properties of auction protocols include individual rationality and Pareto efficiency. Individual rationality means that a bidder does not suffer any loss by participating in an auction if he or she is rational. Pareto efficiency means a bidder’s utility cannot be increased without reducing the auctioneer’s or other bidders’ utilities. In an auction, if Pareto efficiency is satisfied, social surplus, i.e., the sum of an auctioneer’s and all bidders’ utilities, is maximized because an auctioneer and bidders can make a money transfer and we assume a quasi-linear utility.

In designing a protocol, we impose the constraint of incentive compatibility, which means that truth telling is a best policy for each bidder. The revelation principle states that imposing an incentive compatibility constraint does not lose any generality of discussion [7]. This contributes to reducing the search space of protocols. In addition, in computational environments, if this constraint holds, a bidder does not have to spy on other bidders’ valuations, and the bidder benefits from telling his or her valuation to the auctioneer, which leads to attaining system stability and making the implementation of a computational agent easier.

3.2 Failure of simple auction protocols

In the following, we explain why simple auction protocols do not work well. A simple method for allocating goods is to employ the Vickrey-Clarke-Groves (VCG) protocol and to stick to it, i.e., behave in a conservative manner. The VCG protocol chooses an optimal allocation, i.e., an allocation that maximizes the expected social surplus and imposes the following payment on *bidder i*.

$$\textit{payment}_i = \sum_j v_j(G_{-i}^*) - \sum_{j \neq i} v_j(G^*)$$

Here, G^* represents the optimal allocation and G_{-i}^* represents the optimal allocation when *bidder i* does not exist. As a reference, we call this method the allocation-fixed method. However, this method fails to flexibly change an allocation in response to environmental changes, which reduces the social surplus.

EXAMPLE 1. Suppose that there are *bidder 1*, *bidder 2*, *bidder 3* and $g1, g2$. Bidders’ valuations are given in Figure 1 and the disutility caused by re-allocation is -0.5 for a

	$g1$	$g2$	$(g1, g2)$
bidder 1	7.5	4	8
bidder 2	5.5	5.5	6.5
bidder 3	4.5	2.5	5

Figure 4: Expected valuations of goods

bidder if its allocation changes and 0 if its allocation does not change. In addition, we assume $p(\text{weather} = \text{fine}) = 0.5$ and $p(\text{weather} = \text{rainy}) = 0.5$.

Figure 4 shows the expected valuations at the point when the auction is held. Based on these values, $g1$ and $g2$ are awarded to bidder 1 and bidder 2, respectively. Payments for bidder 1 and bidder 2 are 4.5 and 2.5, respectively.

Here, if bidder 1 and bidder 2 swap $g1$ and $g2$ when weather is known to be rainy, the social surplus increases by $6 + 7 - 5 - 5 - 0.5 - 0.5 = 2$. Note that bidder 1 and bidder 2 each suffer the disutility of -0.5 due to re-allocation. However, the allocation-fixed method misses this opportunity to increase the social surplus.

The second simple method for allocating goods in dynamic environments is to do nothing before the results of environmental changes are known, i.e., to behave reactively. However, if the disutility caused by an allocation change from null to some allocation is quite large, it reduces social surplus. For example, we cannot expect many participants at a seminar if it is announced just before it starts.

Another reactive method is to employ re-allocation whenever environmental changes occur. However, this brings about an incentive problem, i.e., a participant worries about whether the goods previously allocated to him or her should be released and thus seeks different goods because there is no equilibrium strategy.

EXAMPLE 2. Suppose the same situation as Example 1. An allocation and payments are determined based on Figure 4, i.e., bidder 1 wins $g1$ and pays 4.5 and bidder 2 wins $g2$ and pays 2.5.

Here, if bidder 3 overstates its valuation of $g1$, i.e., bidder 3 declares 5 instead of 3 if the weather is fine, bidder 1's payment becomes 5.5. In this case, if the weather turns out to be rainy, bidder 1's utility becomes less than zero, i.e., $u_1 = v_1^P(g1; \text{rainy}) - \text{payment}_1 = 5 - 5.5 = -0.5$. Therefore, bidder 1 tries to sell $g1$ and bidder 3 has a chance to buy $g1$. Thus, bidder 3 can manipulate the allocation of goods.

To change an allocation in response to environmental changes and prevent bringing about an incentive problem, we have to develop a method for carefully selecting the part of an allocation to be changed. In the next subsection, we explain how to construct an allocation plan and select a desirable allocation plan among a number of possible plans. In our method, the payment is determined based on the allocation plan, i.e., incorporating the effect of disutility caused by re-allocation into its calculation, which can avoid an incentive problem.

3.3 Protocol for determining an allocation plan

We propose an auction protocol in dynamic environments that re-allocates the goods in response to environmental

change. An extension of the message space and construction of an allocation plan are novel, although the payment calculation is based on the VCG protocol.

First, we explain a method for constructing an allocation plan that is used to determine an allocation and a payment. It requires enormous computation if we simply enumerate all cases. Thus, we take a method based on dynamic programming.

Here, we assume that there is an order in which values of random variables turn out, specifically such a sequence $cond_1 < cond_2 < \dots < cond_l$, where first the value of $cond_1$ is produced.

The construction method of an allocation plan is as follows.

1. Bidders declare valuations of any bundle of goods in any combination of values of random variables to the auctioneer. These valuation values may be true or false.
2. The auctioneer enumerates a combination of bidders' bids so that the allocation feasibility is satisfied, i.e., the same good is not allocated to different bidders simultaneously.
3. The auctioneer calculates social surplus in each case.
4. Next, the auctioneer examines the state before a value of the random variable, $cond_l$, turns out. For each possible allocation, find the optimal transition rule. Here, optimal means that it maximizes the expected social surplus. Note that possible allocations include allocations that some of goods are not allocated to any bidders.
5. Next, the auctioneer considers the state before the values of random variables, $cond_{l-1}$ and $cond_l$ turn out. For each possible allocation, the auctioneer finds the optimal transition rule.
6. The auctioneer continues the above steps until reaching the state before all values of the random variables turn out. At this point, a set of allocation plans, each of which consists of an initial allocation and a set of transition rules, is constructed.
7. The auctioneer finds the allocation plan that maximizes the expected social surplus. If there are more than one allocation plans that maximize the expected social surplus, the tie is randomly broken.

Next, we explain how to execute an allocation plan. The execution method of an allocation plan is as follows.

1. Given an allocation plan, the auctioneer announces it to bidders and then sets the initial allocation in an allocation plan to the current allocation.
2. If the value of random variable $cond_1$ turns out, the auctioneer changes the current allocation to another allocation specified in the allocation plan and announces it to bidders.
3. After values of all random variables turn out, the current allocation (a final allocation) is chosen. The auctioneer imposes the following payment on bidders.

$$\text{payment}_i = \sum_j v_j(AP_{-i}^*) - \sum_{j \neq i} v_j(AP^*)$$

Here, AP^* represents the optimal allocation plan and AP_{-i}^* represents the optimal allocation plan when $bidder i$ does not exist. The amount of $payment_i$ is equal to the other bidders' decreases in expected valuations due to $bidder i$'s participation.

4. EXAMPLES

This section illustrates how to construct an allocation plan and how to determine an allocation and payments by using an example.

Suppose the same problem setting as Example 2. In the proposed protocol, the auctioneer enumerates the combination of bidders' bids in step 2 and calculates the social surplus in step 3, which are shown in Figure 5

$g1$	$g2$	s.s.	s.s.
		fine	rainy
$bidder 1$	$bidder 1$	10	6
$bidder 1$	$bidder 2$	16	10
$bidder 1$	$bidder 3$	14	6
$bidder 2$	$bidder 1$	6	13
$bidder 2$	$bidder 2$	6	7
$bidder 2$	$bidder 3$	8	8
$bidder 3$	$bidder 1$	5	12
$bidder 3$	$bidder 2$	9	11
$bidder 3$	$bidder 3$	4	7

s.s. stands for social surplus.

Figure 5: Results obtained in steps 2 and 3

In step 4, each pair of transition rules are examined. There are 81 pairs of transition rules.

- 1: change an allocation to
 - ($bidder 1, bidder 1$) if fine,
 - ($bidder 1, bidder 1$) if rainy.
- 2: change an allocation to
 - ($bidder 1, bidder 2$) if fine,
 - ($bidder 1, bidder 1$) if rainy.
- (3-81: We omit the other pairs of transition rules due to space limitation.)

This example includes only one random variable. Thus steps 5 and 6 are skipped. In step 7, the auctioneer finds the allocation plan that maximizes the expected social surplus. The social surplus of the allocation plan in Figure 3 is calculated as follows. If it is fine, the disutility caused by re-allocation is 0 and the social surplus of the final allocation is 16, while if it rains, the disutility caused by re-allocation is $-0.5 - 0.5 = -1$ and the social surplus of the final allocation is 13. Therefore, the expected social surplus is $0.5 \times 16 + 0.5 \times (-1 + 13) = 14$ because $p(weather = fine) = 0.5$ and $p(weather = rainy) = 0.5$. The allocation plan in Figure 3 maximizes the expected social surplus. Another one is shown in Figure 6.

There are two possible allocation plans to be selected. A tie is randomly broken. Here, we choose the allocation plan in Figure 3. The payment of $bidder 1$ is calculated as follows. First, suppose $bidder 1$ does not participate in this auction.

initial allocation: $(g1, g2) = (bidder 2, bidder 1)$

final allocation:

$(g1, g2) = (bidder 1, bidder 2)$ if fine
 $(g1, g2) = (bidder 2, bidder 1)$ if rainy

Figure 6: Another allocation plan that maximizes the expected social surplus

initial allocation: $(g1, g2) = (bidder 3, bidder 2)$

final allocation:

$(g1, g2) = (bidder 3, bidder 2)$ if fine
 $(g1, g2) = (bidder 3, bidder 2)$ if rainy

Figure 7: Allocation plan when $bidder 1$ does not exist

In this case, the allocation plan that maximizes the expected social surplus is shown in Figure 7.

Therefore, $bidder 1$'s payment is calculated as follows.

$$\begin{aligned}
 payment_1 &= 0.5 \times (6 + 3) + 0.5 \times (5 + 6) \\
 &\quad - (0.5 \times 6 + 0.5 \times (7 - 0.5)) \\
 &= 3.75
 \end{aligned}$$

The above expression can be written as follows by making a calculation separately in the fine case and in the rainy case.

$$\begin{aligned}
 payment_1 &= 0.5 \times (6 + 3) + 0.5 \times (5 + 6) \\
 &\quad - (0.5 \times 6 + 0.5 \times (7 - 0.5)) \\
 &= 0.5 \times (6 + 3 - 6) \\
 &\quad + 0.5 \times (5 + 6 - (7 - 0.5)) \\
 &= 0.5 \times 3 + 0.5 \times 4.5
 \end{aligned}$$

Therefore, the auctioneer can impose the payment of 3.75 on $bidder 1$ in any case or impose the payment of 3 in the fine case and the payment of 4.5 in the rainy case on $bidder 1$. $payment_2$ is calculated to be 4.75 in a similar manner.

In Example 2, we pointed out that $bidder 3$ can manipulate the allocation. Suppose $bidder 3$ tells a lie and declares $10 - \epsilon$ for $g1$ in the fine case and $7 - \epsilon$ for $g1$ in the rainy case. If the proposed protocol is employed, the payment of $bidder 1$ becomes 7.75 ($10 - \epsilon$ in the fine case and $5.5 - \epsilon$) instead of 3.75. Here, $bidder 1$'s utility is $10 - (10 - \epsilon) = \epsilon$ in the fine case and $6 - 0.5 - (5.5 - \epsilon) = \epsilon$, thus $bidder 1$ does not suffer a loss. If $bidder 3$ further overstates his or her valuations, it changes the allocation plan to a different one. It can be easily verified that $bidder 3$'s utility becomes less than zero in such a case. Thus, $bidder 3$ does not benefit from his or her lie.

5. PROPERTIES OF PROPOSED PROTOCOL

This section proves that the proposed protocol satisfies desirable properties.

PROPOSITION 1. *For each bidder, truth telling is a best strategy.*

PROOF. The utility of each bidder is calculated as follows.

$$\begin{aligned}
u_i &= v_i(AP_i^*) - \text{payment}_i(AP^*) \\
&= v_i^P(AP^*) + v_i^S(AP^*) \\
&\quad - \left(\sum_j v_j(AP_{-i}^*) - \sum_{j \neq i} v_j(AP^*) \right) \\
&= v_i^P(G_i^{\text{final}}; \{\text{cond}\}) + v_i^S(AP^*) \\
&\quad + \sum_{j \neq i} v_j^P(G_j^{\text{final}}; \{\text{cond}\}) + \sum_{j \neq i} v_j^S(AP^*) \\
&\quad - \sum_j v_j(AP_{-i}^*)
\end{aligned}$$

By inspecting the above expression, the fifth term is independent from *bidder i*'s declaration, i.e., *bidder i* cannot manipulate the fifth term. The first to fourth terms are equal to the objective function that an auctioneer tries to maximize. Therefore, for *bidder i*, truth telling is a best strategy. \square

PROPOSITION 2. *The proposed protocol can attain an allocation that maximizes the expected social surplus.*

PROOF. The proposed protocol can induce bidders to declare their true valuations and find a path that maximizes the expected social surplus. Thus, the proposed protocol attains an allocation that maximizes the expected social surplus. \square

PROPOSITION 3. *The proposed protocol satisfies the individual rationality constraint on expected values.*

PROOF. If a bid of *bidder i* does not have any effect on the final allocation, it is not included in the series of tentative allocations, which means its utility is 0. On the other hand, if a bid of *bidder i* is included in the final allocation, *bidder i*'s utility becomes equal to the increase in social surplus by *bidder i*'s participation. In addition, a possibility exists that nothing is allocated. Thus, even if disutility caused by re-allocation is incurred by bidders, individual rationality on the expected value is still satisfied. \square

Note that proposition holds on expected values. It may happen that a bidder's utility becomes less than zero in some cases, although the expected utility is larger than or equal to 0. However, as shown in section 4, if we impose a conditional payment on bidders, a bidder's utility is more than or equal to 0.

PROPOSITION 4. *The proposed protocol can increase the expected social surplus to higher than or equal to that by the allocation-fixed method.*

PROOF. An allocation obtained by the allocation-fixed method is included in a candidate of an allocation plan, and the selected allocation plan gives the highest expected social surplus among all possible allocation plans. \square

Lastly, we discuss the computational burden. The VCG protocol needs to solve a combinatorial optimization problem in calculating an allocation and a payment. Many winner determination algorithms have been proposed [2, 12]. In our proposed protocol, the computation cost is more serious because a combination of allocations is examined to obtain an optimal transition rule, although dynamic programming

mitigates this problem. Developing an efficient way of finding an allocation plan is part of our future work.

An efficient way of finding an allocation plan might include an approximation method. However, comparing the result obtained by approximation methods to the optimal case is useful for evaluating how effective these approximation methods are. Therefore, we believe that this paper contributes to the field of auctions by setting a reference point for auctions in dynamic environments.

6. DISCUSSION

Dynamic games with incomplete information have been studied in game theory [10]. In this paper, we focused on the effect of re-allocation in the context of an auction. One way of attaining re-allocation is resale. In the economic literature, the effect of resale opportunities in auctions has been discussed [3]. Resale may improve inefficiency in some cases.

Milgrom, however, pointed out that there exists no mechanism that can reliably untangle an initial misallocation in any two-sided negotiation [8]. That is, at the point of resale, the seller has an incentive to exaggerate its value and the buyer has an incentive to pretend its value is lower. These misrepresentations can delay or scuttle a trade.

Moreover, in combinatorial cases, if we consider dynamic environments and the fact that a bidder's utility depends on the environmental conditions when the allocated goods are actually used, the loser can strategically manipulate an allocation, which results in the failure of an efficient allocation, as shown in Example 2.

Therefore, we examine a way to induce bidders to tell their true valuations. To do so, our model explicitly incorporates the cost due to re-allocation and the proposed auction protocol considers the disutility caused by the re-allocation as well as the valuations of goods.

Sandholm and Lesser proposed leveled commitment contracts [13]. Although a contract may be profitable to an agent when viewed ex ante, it may not be profitable when viewed after some future events have occurred, i.e., it comes to be viewed ex post. Leveled commitment contracts are a method of taking advantage of the possibilities provided by probabilistically known future events. Compared to the leveled commitment contracts method, Sandholm and Lesser pointed out the problems of contingency contracts; there is a potential combinatorial explosion of goods to be conditioned on and it is often impossible to enumerate all possible relevant future events in advance.

Our primary concern is to examine whether an auctioneer can induce a bidder's truth telling in the case where re-allocation is possible. In addition, we introduce the dynamic programming method of constructing an allocation plan to reduce computational cost.

Larson and Sandholm [6] studied deliberation and bidding strategies of bidders with unlimited but costly computation who are participating in auctions. A bidder does not a priori know his or her valuations for auctioned goods. Here, strategic computation becomes a problem. In their problem setting, a signal, i.e., a factor that determines bidders' valuations, is not shared, while in our setting, a signal, e.g., the result of weather, is shared among bidders. Therefore, we can attain an ex ante efficient allocation.

Ito *et al.* [4] dealt with cases where the valuations of goods depend on their qualities and experts can learn these

qualities but amateurs have no idea about the qualities. They have proposed a series of protocols that can elicit experts' information about the quality of the goods and attain an efficient allocation. In their problem setting, an amateur's valuation of the auctioned good is determined if the quality of the good is revealed. This rather resembles the problem setting in this paper. However, this paper assumes that nobody knows the value of a random variable in advance, which is a big difference between their studies and this paper from the viewpoint of mechanism design.

Conditional planning conditioned on environmental change and re-planning caused by environment changes have been also studied in planning research [11]. These research efforts have not dealt with an agent incentive problem, and thus this paper is different from these planning studies.

This paper assumes that probabilistic distribution of random variables is common knowledge. Consider the weather. It is easy to obtain weather information from a weather report and it is not likely that one bidder is considerably better informed of the weather forecast than the other bidders. In some cases, an assumption of common knowledge could be too restrictive and we may have to deal with cases in which bidders are differently informed. Extending the proposed protocol and dealing with such cases are our future work.

7. CONCLUSIONS

This paper proposed an auction protocol in dynamic environments. A bidder's valuation of goods depends on environmental conditions when the allocated goods are actually used. If the environmental conditions are represented by some random variables and the probabilistic distributions of the values of these random variables are given, a solution is to ask bidders to declare their expected valuations under the probabilistic distribution and find an efficient allocation of the goods. Flexibly re-allocating goods whenever a value of random variables turns out may increase social surplus; however, if we introduce re-allocation in a simple manner, the existence of an equilibrium is no longer guaranteed. This would cause bidders to worry about what bids to submit and it would make it difficult for a system designer to predict what outcome is obtained.

To solve this problem, this paper proposed an auction protocol that asks bidders to declare disutility caused by re-allocation as well as valuations of goods in each case and makes an allocation plan that specifies an initial plan and transition rules. The proposed protocol can find an allocation that maximizes the expected social surplus including the disutility caused by re-allocation and benefits obtained from the allocated goods in a final allocation. In this paper, we (1) proposed an auction framework that accounts for re-allocation in dynamic environments, (2) introduced a new auction protocol that can induce bidders to tell their true valuation and disutility caused by re-allocation by extending the well-established Vickrey-Clarke-Groves (VCG) protocol, and (3) proved that the protocol can obtain an ex ante Pareto efficient allocation.

If environments become more complicated and more random variables are required to represent the environment, the problem of computational cost again becomes serious. Our future work includes conducting a complexity analysis and/or a computer simulation to assess the performance of the proposed protocol in terms of the size of problems.

Moreover, we would need to find a way that first makes an

abstract allocation plan by using the proposed protocol and then makes a detailed plan in a reactive manner. Examining how it would affect a bidder's strategy and social surplus is our future work.

8. REFERENCES

- [1] R. K. Dash, N. R. Jennings, and D. C. Parkes. Computational-mechanism design: a call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.
- [2] Y. Fujishima, D. McAdams, and Y. Shoham. Speeding up ascending-bid auctions. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 554–559, 1999.
- [3] P. A. Haile. Auctions with private uncertainty and resale opportunities. *Journal of Economic Theory*, 108(1):72–110, 2003.
- [4] T. Ito, M. Yokoo, and S. Matsubara. A combinatorial auction among versatile experts and amateurs. In *Proceedings of the third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2004)*, pages 378–385, 2004.
- [5] D. N. Kinny and M. P. Georgeff. Commitment and effectiveness of situated agents. In *Proceedings of the Twelfth International Conference on Artificial Intelligence (IJCAI-91)*, pages 82–88, 1991.
- [6] K. Larson and T. Sandholm. Costly valuation computation in auctions. *Proceedings of Eighth conference of Theoretical Aspects on Rationality and Knowledge*, pages 161–182, 2001.
- [7] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [8] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.
- [9] R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz. Mechanism design with execution uncertainty. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI-02)*, 2002.
- [10] E. Rasmusen. *Games and Information: an introduction to game theory*. Blackwell Publishers, 3rd edition, 2001.
- [11] S. Russell and P. Norvig. *Artificial Intelligence A Modern Approach*. Prentice Hall, 1995.
- [12] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence Journal*, 135(1-2):1–54, 2002.
- [13] T. W. Sandholm and V. R. Lesser. Advantages of a leveled commitment contracting protocol. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, pages 126–133, 1996.
- [14] 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.