Group Decision and Negotiation, 2:363–373 (1993)

Cryptographic Protocols for Vickrey Auctions

HANNU NURMI
Department of Political Science, University of Turku, FIN-20500 Turku, Finland

ARTO SALOMAA
Academy of Finland and Department of Mathematics, University of Turku, FIN-20500 Turku, Finland

Abstract

Although the sealed second-price or Vickrey auctions have some nice theoretical properties, they are fairly seldom utilized in practice. It has been suggested that they are vulnerable to bid-taker cheating and that the revelation of bids after the bidding makes the bidders reluctant to reveal their true valuations. We outline procedures based on modern mathematical cryptography that are instrumental in avoiding some of these difficulties and thereby will improve the properties of the Vickrey auctions.

Key words: auctions, cryptographic protocols, public-key cryptography

1. Introduction

Over the past decades the study of social institutions, arrangements, and norms has opened new vistas for the understanding of how human societies work. By so doing it has been instrumental in posing a new question of normative nature, e.g., how to improve the performance of existing institutions. Much of this literature takes welfare economics and social choice theory as its point of departure. A typical assumption is that society consists of rational agents pursuing their own interests and establishing institutions that they deem acceptable and necessary from their own viewpoint. The emergence of the institutions is not necessarily construed as a result of a purposeful design by individuals; it may indeed be a byproduct of activities striving for other goals.

Auctions are well-known institutions in all parts of the world. In auctions a number of buyers and sellers of some good try to determine how much of the good is to be sold by whom to whom and at what price. Interestingly enough, several types of auctions are known to exist. They all share the feature that both sellers and buyers make offers to sell and buy, respectively, at a given price. The offers to buy are usually called bids. The types of auctions differ in the way in which the bids and selling offers are allowed to be made. Thus, for example, in the English (or ascending-bid) auction, the price of a good is successively raised until only one buyer is willing to make a bid for the good at the most recently announced price (provided that only one good is to be sold). It is essential that at any point in the succession of bids the bidders know the value of the most recent bid. In Dutch (or descending-bid) auctions, on the other hand, the auctioneer starts from a high price for a given good and successively lowers it until one buyer is willing to make a bid at the given price. (See Engelbrecht-Wiggans 1980 for an overview of various types of auctions. For a more specific survey of game-theoretic models of auctions, see McAfee and McMillan 1987.)

These two types of auctions are typically oral ones, although their implementation in communication networks presents no major problems. A sealed version of the English auction or first-price sealed-bid auction is obtained by requiring the bidders to submit sealed offers and by awarding the good to the highest bidder at the price that he/she (hereinafter "she") offers in her bid.

In this article we focus on a relatively new type of auction, viz., the Vickrey auction or sealed second-price procedure. This procedure was devised by Vickrey (1961) to overcome the practical difficulties involved in arranging oral auctions while preserving their theoretically nice properties. In the following section we discuss the theoretical properties of Vickrey auctions. Thereafter, we outline ways of solving some of the problems related to those auctions.

2. The Vickrey auction

In the sealed second-price or Vickrey auction the sealed bids are submitted to the auctioneer within a predetermined period of time. After this period has elapsed, the auctioneer opens the bids and the good is sold to the bidder who has submitted the highest bid. The price, however, is not the one that she offers in her bid, but the second-highest bid.

Under conditions to be discussed shortly the Vickrey auction has some nice properties (see Vickrey 1961; Rothkopf et al. 1990). First, if all bidders bid truthfully, the result is a Nash equilibrium. In other words, no bidder has an incentive not to reveal her value of the good in her bid provided that the others bid truthfully. In fact, a stronger claim can be made: revealing one's true value is the dominant strategy, i.e., regardless of whether the others reveal their true values in their bids, the best thing a bidder can do is to bid truthfully.

Second, the outcome reached through the Vickrey auction is Pareto optimal, i.e., no reallocation leads to a situation that is weakly preferred by all to the outcome. The good is sold to the bidder with the highest private valuation of the good, and no bidder is willing to offer a higher price than the one obtained by the seller. Thus, the procedure is efficient in the technical sense. The conditions under which the above properties characterize the Vickrey auction are (i) that the bidders are risk-neutral, (ii) that the bidders are symmetric, (iii) that the private values assigned to the good by bidders are statistically independent, and (iv) that the price at which the good will be sold is a function of the bids only.

The first condition states that the bidders are indifferent between options that have the same monetary value. The second condition, in turn, requires that the bidders draw their values for the good from the same probability distribution, i.e, that all players are of the same "type". The third and fourth conditions are self-explanatory.

3. Obstac es for use

Although some of the conditions under which the Vickrey auctions have the above-mentioned nice properties are stringent, it is remarkable that they are very rarely utilized in practice. One possible explanation for this is that those who decide which auction type is to be adopted do not find it optimal. This explanation is, however, inadequate, because under the above conditions all four types, some of which are frequently used, lead to the same outcomes. Thus, whoever decides the auction form—be it the bidders or sellers—cannot expect a better outcome from the other basic auction forms. Rothkopf et al. (1990) mention five potential reasons for the reluctance of real-world actors to make use of this type of auction. All of them—the bidders' risk-aversion, bidder asymmetry, nonindependent private valuations of the bidders, institutional inertia, and the fact that multiple objects are usually offered for sale instead of a single good—are shown to be inadequate explanations for the rarity of Vickrey auctions. However, the authors argue that two hitherto undiscussed reasons might provide a more satisfactory explanation: (1) the bidders' fear of getting cheated by the seller, and (2) the unwillingness to disclose one's true valuations. Of course, it is also possible that the bidders cheat the seller by colluding and driving down the price of the good. According to Robinson (1985), this possibility can to some extent be avoided by using sealed instead of oral auctions. This suggestion rests on the observation that the agreements between bidders are easily enforceable in oral auctions, while in sealed auctions their stability is not self-evident.

In sealed auctions the seller may cheat the bidders by first looking at the submitted bids and then having her accomplice submit losing bids, thereby driving up the price that the winning bidder has to pay. The mere possibility of this happening may give the bidders an incentive to deviate from the truthful value revelation strategy. Vickrey (1961) suggests the use of trusted middlemen in the handling of submitted bids.

The unwillingness of the bidders to reveal their true valuations of the good may be due to conditions prevailing on the market. Thus, for example, the firms bidding for some contract may be reluctant to reveal their true valuations of the contract, because their competitors would thereby get valuable information about the firm's cost estimates, which, in turn, could reveal the level of technology at the firm's disposal. In the following, we shall outline a cryptographic version of

Vickrey auctions that would seem helpful in overcoming these obstacles to the use of Vickrey auctions.

4. Provisionally secret bidding

Cryptography is the study of secret writing. In classic cryptosystems the encryption (or ciphering) keys are related to the decryption (or deciphering) keys so that given one key the other one can easily be determined. In public-key cryptosystems, by contrast, this link between the keys is broken: if one knows the encryption key, one does not *ipso facto* know the decryption key (see Diffie and Hellman 1976). Public-key cryptosystems make use of one-way functions, i.e. functions with the property that given the value of the function it is computationally intractable to determine the argument of the function without some additional "trapdoor" information (see Salomaa 1985, 1990). One example of such a function is modular multiplication: given a product (modulo k) of two large prime numbers, it is computationally intractable to determine the primes in question. In cryptosystems that are based on the intractability of modular exponentiation (e.g., the well-known RSA system (see Rivest et al. 1978)), the sender A of a message announces publicly that messages to her can be sent by performing an easy computation, viz., by raising the numeric version of the message partitioned into blocks of suitable size to the tth power (mod n), where both t and n are fixed numbers. A chooses n so that it is a product of two large primes. The factorization of n is A's trapdoor information which enables her to easily decrypt messages sent to her. Anyone else but A faces a computationally intractable task in trying to determine the factors of n. It is noteworthy that the sender of the message is also ignorant about the factorization.

Cryptographic protocols are ways of systematically utilizing the results of mathematical cryptography in devising communication systems that serve the purposes of the parties involved with regard to the secrecy of messages. These purposes may, for example, call for a complete exclusion of third parties from the information transmitted between two parties, or for a partial revelation of secrets by all parties. Consider as an example a protocol that guarantees the secrecy of messages from party A to party B so that:

(i) A knows that nobody else but B can read the message;
(ii) B knows that the message came from A; and
(iii) B knows that A cannot afterwards claim not having sent the message to B.

Let us denote the numerical encoding of the message by w (the encoding is obtained by mapping the letters into numbers 0, . . . , 25), the A's (B's, respectively) public encryption key by e_A (e_B). Similarly, A's (B's) private decryption key is denoted by d_A (d_B). We shall make the assumption that for all messages and all keys, the following holds: $d_A(e_A(w)) = e_A(d_A(w)) = w$.

A protocol satisfying the above three conditions is the following:

Step 1 A decrypts w using d_A. Thus w is transformed into $d_A(w)$.
Step 2 A then encrypts the result using B's public encryption key. The result is $e_B(d_A(w))$. This is sent to B.
Step 3 B first decrypts what she received by using her private decryption key. The result is (by the above assumption) $d_A(w)$.
Step 4 B finally encrypts the result of *Step 3* by A's public encryption key to obtain w (Salomaa 1985).

It should be emphasized that this protocol—or any other protocol for that matter—can be no safer than the underlying cryptosystem, i.e., the way in which messages are encrypted and decrypted using the public and private keys. Without going into the details of these systems, we can assume that adequate cryptosystems are available (see Salomaa 1985, 1990 for details).

It can readily be seen that the above protocol fulfills the requirements (i) through (iii). Requirement (i) is satisfied because in deciphering $e_B(d_A(w))$ anyone else but B is facing a computationally intractable task. Requirement (ii) is also satisfied, since in decrypting the message received B will have to use A's public encryption key in *Step 4*. Requirement (iii), finally, follows from the fact that in *Step 3* B possesses the message w in a version where it is decrypted by A's private decryption key. As only A is supposed to know this key, *Steps 3* and *4* guarantee (iii).

Simple as it is, this protocol (call it protocol 1) is obviously applicable in sealed bidding in general and in Vickrey auctions in particular: A denotes the bidders and B the bid-taker or seller.The requirements (i) through (iii) are clearly desirable in bidding contexts. However, the problem of bid-taker cheating remains. After the bids have arrived, the bid-taker is the only one who knows their content. In case all bids are publicly posted and B's private decryption key is made public after the bidding period is over, the possibility of the bid-taker cheating by having "artificial" bids submitted is restricted. If one wants to exclude this possibility altogether, the best way to proceed is to ask the bidders to submit their offers in encrypted forms, as in the following protocol, which will be called protocol 2.

Step 1 A sends B the following message: $<e_A(w), e_B(d_A(j(A)))>$ where w is A's offer and $j(A)$ is A's identification, e.g., name and phone number.
Step 2 B posts all the received $e_A(w)$'s.
Step 3 After the bidding time is over, the bidders A send B their private decryption keys, which are made public.
Step 4 All the bidders and the seller can now decrypt the prices offered and be assured that the good is sold to the highest bidder at the second highest price.

Step 1 enables the bid-taker B to identify the bidders, although she cannot find out their bids because they are given in encrypted form. *Step 2,* in turn, excludes the kind of bid-taker cheating in which "artificial" bids are added to the genuine ones after the bid-taker finds out the prices offered. At *Step 3* the private keys are made public, whereupon everyone can recover the bids offered at *Step 1.* Although protocol 2 excludes the particular type of bid-taker cheating mentioned by Vickrey (1961) and Rothkopf et al. (1990), it requires the revelation of all bids. On the other hand, it is difficult to envisage an efficient procedure in which the bids would not be known even to the bid-taker. In protocol 2 the identification of the bidders does not have to be known to all bidders; the protocol only requires that the bids be recoverable by everyone. Thus, the bid-taker may publish the decryption keys only while keeping the identity of the bidder to herself.

5. Minimal bid-revelation

Suppose that one wishes to design a version of the Vickrey auction in which as few bids as possible are made known to the bidders and the bid-taker. Obviously, the bid-taker has to know at least one bid, viz., the second-largest one. Otherwise she could not collect the payment for the good. But does she need to know the content of other bids? In the following we shall outline a cryptographic protocol (let us call it protocol 3) in which she does not. The protocol is an application of Yao's solution to what he calls two millionaires problem (commonly also known as the age protocol), which is an arrangement enabling a group of people to determine who is the wealthiest (or oldest for that matter) without disclosing the exact wealth (or age) of any member of the group (Yao 1982). (See also Salomaa 1985, 1990.) While the protocol faces some incentive problems in wealth-revelation contexts, the fact that we are dealing with sealed (or rather secret) second-price auctions guarantees that the incentives for not bidding truthfully are absent. Let us start with a special case with two bidders, A and B, only (see Salomaa 1985; Nurmi 1989). Let the set of possible bids be representable by positive integers in the interval [1,100]. Let A's (B's, respectively) private bid be the *i*th (*j*th) value in the interval.

Step 1 B chooses randomly a large number x and encrypts it with A's public encryption key to obtain $e_A(x) = k$. Both A and B send to the bid-taker their bids in the form: $<e_C(e_A(p_A s_A)), e_C e_A(s_A)>$ and $<e_C(e_B(p_B s_B)), e_C e_B(s_B)>$, respectively. Here e_C is the bid-taker's public encryption key, p_A (p_B, respectively) A's (B's) bid, and s_A (s_B) a number privately chosen by A (B).

Step 2 B sends A the value $k-j$, that is, the difference between the encrypted version of the number that she randomly chose in *Step 1* and the ordinal number that identifies her own bid.

Step 3 A now computes a number sequence:

$y_u = d_A(k-j+u)$, where $u = 1, \ldots, 100$.

A does not reveal this sequence to B. A also computes for each u the following value:

$z_u = y_u \pmod q$ where q is a prime chosen by A.

Each z-value must be smaller than $q-1$. The difference between any two z-values has to be at least 2. Should either of these requirements not be the case, another value of q is chosen by A.

Step 4 A reports to B the following sequence:

$z_1, \ldots, z_i, z_{i+1}+1, z_{i+2}+1, \ldots, z_{100}+1, q.$

The condition that the difference between any two z-values be at least two guarantees that no number appears twice in this sequence.

Step 5 B determines whether the following condition is met:

$z_j = x \pmod q$.

If it is, then j, which identifies B's bid, is no larger than i, which identifies A's bid. If it is not, then B's bid is strictly larger than A's. This conclusion follows from the fact that

$z_j = y_j \pmod q = d_A(k-j+j) = d_A(k) = x$.

Here the second equality is a consequence of *Step 3* with $u=j$. If j is strictly larger than i, z_j+1 is received and the equations do not hold.

Step 6 B informs A about her conclusion in *Step 5*. The losing bidder sends her decryption key to the bid-taker, whereupon the latter can recover her bid (see *Step 1*) and thus the price to charge the winning bidder.

That the conclusion reached in *Step 5* is correct can be seen from the following (Salomaa 1990). Consider the jth number z_j' in the sequence of z_u ($u = 1, \ldots, 100$) generated by A in *Step 3*. If i is larger than or equal to j, then $z_j' = z_j = y_j = x \pmod q$. If, on the other hand, $j > i$, then $z_j' = z_j + 1$ which is *not* congruent with z_j, and thus z_j' is not congruent with $x \pmod q$. We notice that the requirements that the absolute value of the difference between any two z-values be at least 2 and that each value be smaller than $q-1$ are necessary to guarantee that each z-value appears only once in the sequence of *Step 4*. The requirements are easily met because the primes are large; consequently, the differences between two z-values are not likely to be 0 or 1. Another choice of q is seldom needed.

Clearly, B's incentives for not reporting the conclusion correctly to A are absent. Yet, regardless of the conclusion, B does not know the amount offered by

A. She only knows that the amount either exceeds her own offer or is at most as large as the latter. If A's bid is strictly higher than B's, then obviously A should be awarded the good at the price that B offers. If, on the other hand, A's bid is exactly as high as B's, then the price of the good can be immediately determined, viz., it is equal to A's bid. However, it is not known which one of the parties ought to get the good at that price. By repeating the procedure so that A's and B's roles are reversed, one finds out whether A and B happened to bid exactly the same amount.

Having determined which of the parties has a higher offer, we still have to be assured that the party with the losing bid reports her bid correctly to the bid-taker. If bidder collusion can be disregarded, the only possibility for misreporting is that the losing bidder reports that her bid was higher than it actually was. This possibility can be counteracted by the requirement that the losing bidder reveal those choices that she has made at various stages of protocol 3. Thus, if B is the losing bidder, she informs A of j and x, whereupon A can check that B is not inflating the price of the good. If A, in turn, is the losing bidder, she informs B about the y_u ($u = 1, \ldots, 100$) sequence that she computed. Thus B can check the correctness of the result reached. An additional protection against cheating is provided by the bid-taker who has been given the size of the losing bid in *Step 6*.

In the two bidder case it thus appears that only one bid has to be publicly announced to implement the cryptographic version of the Vickrey auction. But protocol 3 can be utilized in an n-bidder context as well. A straightforward way to do that is to perform all $n(n-1)/2$ pairwise comparisons of bids using protocol 3. One then forms an n-by-n matrix of 0's and 1's so that the cell $(k,m) = 1$ if the kth bidder has a higher bid than the mth bidder. Otherwise, $(k,m) = 0$. If a unique highest bid exists, then obviously there is a row in the matrix with 1's in all non-diagonal cells. If a unique second-highest bid exists, then there is a row with 1's in all non-diagonal cells except one, viz., the rth column (if the largest bid is that of bidder r).

In general one can do with much fewer than $n(n-1)/2$ pairwise auctions. From the point of view of sorting, we have to find the second largest among n numbers. Many sorting algorithms, linear in terms of n, are known for this task. However, we have to take the additional precaution that the largest number remains secret. The following procedure, where protocol 3 is applied $2n - 2$ times, seems appropriate.

Assume that the bidders are $A_1, A_2, \ldots, A_n$. Assume, further, that the bid-taker has informed each bidder about the bidders following her in this ordering. Thus, A_1 knows the whole list, but A_n knows only that she is the last bidder. As before, the bidders initially give the bid-taker their bids in encrypted form.

In the first round the highest bidder is found out by $n - 1$ applications of protocol 3. A_1 first compares her bid with A_2, then with A_3, until she finds someone, say A_5, with a bid at least as high as hers. A_5 then takes over, and so forth. When some bidder, say A_i, reaches A_n without losing, she reports to the bid-taker as the highest bidder. She only discloses her identity, not her bid. Observe that the identity of A_i remains secret for most of the bidders, although A_n knows it. If there are several bidders with the same highest bid, the rightmost bidder in the list wins. This is only one of the possible ways to handle this situation.

The second round proceeds exactly as the first, except that now A_i changes her bid to a surely losing one. Consequently, the second highest bidder A_j is found out by another $n-1$ applications of the protocol. A_i gives her secret decryption key to the bid-taker, who now has all the information she needs.

The above procedure of two rounds seems to disclose very little unnecessary information to the participants. Quite a different possibility is to modify the original protocol 3 to a multiparty protocol that gives the answer directly. The result will be much more involved than the procedure outlined above. Some sophisticated types of cheating are still possible. In particular, the bidders might use in the protocol different bid values from the ones they originally gave to the bid-taker. Such cheating will be found out if everything is disclosed at a later stage.

During the process of comparing bids to find out the two highest ones, the only available information about the bids is of an ordinal nature, i.e., after each comparison the bidders know which one of them has a higher offer. Once the two highest bids have been found, the size of the smaller one has to be made known at least to the bid-taker and the highest bidder. Preferably, the value of the second-highest bid should be public so as to eliminate some obvious forms of discrimination. As we pointed out above, an arrangement in which no bid would eventually be made known to at least the bid-taker would seem downright impossible. The procedure described above thus represents the best one can hope to achieve as far as bid-revelation is concerned.

One could argue, however, that there are circumstances in which protocol 3 encourages rather than discourages a particular type of cheating, viz., one in which the bid-taker solicits her accomplices to submit high bids to get higher price from a bidder—let us call her T—that the bid-taker knows desperately wants the item or right being auctioned. In this kind of situation the fact that the second highest bidder's identity is not revealed would seem an additional source of concern for T. Although protocol 3 does not eliminate this type of cheating altogether, it renders its success much less likely than in ordinary sealed bid auctions. In the latter it is conceivable that some kind of information about the magnitude of T's bid leaks to the bid-taker before the time for submitting the sealed bids has expired. Thus the bid-taker could call upon her collaborators to submit second highest bids to drive up the price that T has to pay. In protocol 3, in contrast, only T's identity (in case she happens to be the highest bidder) is known to all others. And even this information is revealed "too late," since by that time all bids have already been submitted in encrypted form. If the bid-taker knew T's identity before the bidding starts, any attempt to drive up the price by soliciting bids from accomplices runs the risk of backfiring, viz., it may well turn out that one of the solicited bids is the highest, whereupon protocol 3 calls for the identity of the bidder being revealed. Now, what about the second highest bidder's incentive to reveal her bid? Protocol 3 by itself does not contain any. The bidders and the bid-taker are assumed to commit themselves to these rules before the bidding begins. It is, however, worth pointing out that the losing bidder has no way of revising her bid once the fact that she is the losing bidder is known.

6. Practical considerations

Due to the rather extensive computations needed in going through the steps of protocol 3, it is impossible to utilize the cryptographic version of Vickrey auctions unless the bidders and the bid-taker have considerable computing capacity at their disposal. If the RSA cryptosystem is used, it has been suggested as a rule of thumb to choose a composite number *n* so that it consists of at least 200 digits. This would guarantee a quite safe system, as the computing time needed to factor arbitrary 100-digit numbers is huge even by the best presently known methods.

To perform modular exponentiations of 200-digit numbers requires special equipment—RSA chips that are currently available—even though from a theoretical point of view the computations are easy for a legitimate user. This requirement necessarily excludes the use of the above protocol in most present-day auctions. On the other hand, once the special equipment is available, the steps of protocol 3 can be taken quite swiftly.

Nothing in protocol 3 hinges upon the parties, i.e., the bidders and the bid-taker, to gather in the same place. Every step can be performed in an electronic communications network or by ordinary letter, for that matter.

7. Concluding remarks

In the preceding we have outlined a procedure for executing Vickrey auctions using cryptographic methods. The procedure seems to minimize, or even to avoid, the two problems that stand in the way of a general adoption of Vickrey auctions: the fear of bid-taker cheating, and the reluctance to reveal bids. The procedure does not, however, destroy the particularly nice property of the Vickrey auctions, viz., that truthful bidding is the dominant strategy for the bidders (assuming that they are risk-neutral and independent). In situations where the bid-taker knows that there is one bidder who desperately wants the item being auctioned, this property may be absent, but this is also the case when traditional sealed second-price auctions are considered. The implementation of cryptographic Vickrey auctions requires special computational equipment.

Acknowledgments

The authors gratefully acknowledge the comments of Professor K. Chatterjee and the referees on an earlier version of this article. The work reported here has been supported by the Academy of Finland grants 1071178 (Nurmi) and 1071041 (Salomaa).

References

Diffie, W., and M. Hellman. (1976). "New Directions in Cryptography," *IEEE Transactions on Information Theory* IT-22, 644–654.

Engelbrecht-Wiggans, R. (1980). "Auctions and Bidding Models: A Survey," *Management Science* 26, 119–142.

McAfee, R. P., and J. McMillan. (1987). "Auctions and Bidding," *Journal of Economic Literature* 25, 699–738.

Nurmi, H. (1989). "Computational Approaches to Bargaining and Choice," *Journal of Theoretical Politics* 1, 407–426.

Rivest, R., A. Shamir, and L. Adleman. (1978). "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," *ACM Communications* 21, 120–126.

Robinson, M.S. (1985). "Collusion and the Choice of Auction," *Rand Journal of Economics* 16, 141–145.

Rothkopf, M.H., Th.J. Teisberg, and E.P. Kahn. (1990). "Why Are Vickrey Auctions Rare?", *Journal of Political Economy* 98, 94–109.

Salomaa, A. (1985). *Computation and Automata*. Cambridge: Cambridge University Press.

Salomaa, A. (1990). *Public-Key Cryptography*. Berlin-Heidelberg-New York: Springer-Verlag.

Vickrey, W. (1961). "Counter Speculation, Auctions and Competitive Sealed Tenders," *Journal of Finance* 16, 8–37.

Yao, A.C. (1982). "Protocols for Secure Computation." In *23rd Annual Symposium on Foundations of Computer Science*, pp. 160–164. IEEE Computer Society Press.