Designing Mechanisms for E-Commerce Security: An Example from Sealed-Bid Auctions

Wenli Wang, Zoltán Hidvégi, and Andrew B. Whinston

ABSTRACT: The use of computing technology is a necessary but not sufficient approach to enhancing e-commerce security. It should be supplemented by the application of economic mechanisms to design e-processes that discourage the exploitation of security weaknesses. Mechanisms that function well in traditional commerce should be updated to accommodate new problems raised in e-commerce, such as the lack of authentication. As an illustration, thiezezs research focuses on on-line auction mechanisms and develops a Leveled Partition Set (LPS) protocol for multi-unit sealed-bid auctions. LPS prevents on-line false-name bidding whereby bidders increase their utility by splitting their bids for a large bundle to several smaller bundles under false identities.

KEY WORDS AND PHRASES: Economic mechanisms, electronic commerce, information system, multi-unit auction, on-line auction, security.

E-commerce security has become a serious concern of firms and individuals that rely on distributed digital processing in their daily operations. Security breaches and frauds cost business and consumers millions of dollars. Although technical imperfections contribute to security breaches, fraud is often made possible by the fact that many trading processes viable in traditional commerce are fundamentally flawed when conducted over the Internet. Although computer scientists and practitioners have provided technical tools (e.g., firewalls, cryptographic protocols) that enhance computing and networking security, issues related to Internet-based e-business processes remain unaddressed [10]. There is a growing need for powerful tools and rigorous methods in the design and verification of correct e-process systems that operate over the Internet.

The authors are conducting a series of studies aimed at providing a package of such tools and methods. Researchers in the engineering and business domains most frequently seek technical solutions to the problem of improving on-line security. But no program of reasonable size, whether providing computing services or running business applications, is bug-free. Wang, Hidvégi, and Bailey have suggested the use of model checking, a modern formal verification tool, to mathematically verify the correctness of e-process technical implementations [10]. Since this method is quite expensive, a different perspective is explored in this paper—the rationale behind security breaches and fraud. A great many security problems are caused by the dishonest intentions of human users, and not by technical factors. Unscrupulous

A preliminary version of this paper was presented at the International Conference on Electronic Commerce 2000 (ICEC2000) held in Seoul, Korea, on August 21–24, 2000. Jae Kyu Lee served as guest editor for the paper.

International Journal of Electronic Commerce / Winter 2001–2002, Vol. 6, No. 2, pp. 139–156.

1086-4415/2002 $9.50 + 0.00.

hackers and fraud perpetrators exploit design flaws and implementation bugs to obtain economic benefits or simply for the thrill of it.

Given this understanding of the rationale for security breaches and fraud, the appropriate approach is to search for methods of modifying human intentions. Mechanism design in economics deals with the design of human incentives and trading rules that lead to optimal social welfare when all participants follow their own best interests. The revelation principle of mechanism design says that to any equilibrium in a game of incomplete information, there corresponds an associated revelation mechanism that has an equilibrium with the same outcome where the participating agents truthfully report their types. In other words, there always exists an incentive rule where the optimal strategy for participating agents is to truthfully reveal their private information. In light of this, mechanism design can be used to modify the incentives of e-process users by making their honest activities at least as profitable as their fraudulent ones. Although mechanisms that achieve this will not make fraud totally disappear (e.g., some users may not be rational), they can make it impractical by eliminating its monetary damage. Other than a few exceptions, there has been very little research on applying mechanism design to enhance e-commerce security [11, 12, 15].

This paper investigates a specific type of fraud, false-name bidding in on-line auctions, and develops a secure multi-unit sealed-bid on-line auction protocol to demonstrate why and how mechanism design can be applied to enhance e-commerce security.

Popularity of On-line Auctions and Fraud

Auctions were held in a remarkable range of situations in pre-Internet commerce, but no one ever predicted how vast their scope would be in the digital economy. The quick, even unexpected growth of eBay is a window showing the popularity of on-line auctions. As of June 2001, five years after its launch, eBay already had more than 24 million users and more than 6 million items for sale on any given day. On-line auctions have not only proven to attract consumer-to-consumer sales, but also attract business-to-consumer and business-to-business auction transactions. Numerous small businesses, using the Internet to reach a large audience without any physical constraints, apply auction mechanisms as transaction platforms. An examination of the customer base of on-line auction sites indicates that even many large and mid-sized corporations (e.g., Sun Microsystems) are expanding toward on-line auctions for their overstocks, discontinued items, or even general merchandise. Trading through on-line auctions is not a fad but a trend. Forrest Research estimates that on-line auction sales will reach $19 billion in 2003. Jupiter Communications predicts it will top $26 billion in 2004.

Auctions provide an effective medium whenever there is a need to establish a unique price or each individual item for sale. The Internet makes on-line auctions accessible to the general public and therefore introduces an unpredictable scope of demand and supply. Sellers of used, discontinued, or overstocked items prefer on-line auctions because it is difficult to predict the demand for these

items and hence their prices. In addition, the operational efficiency of the Internet allows auction sellers and buyers to interact at a very low participation and transaction cost, which is unimaginable in the traditional economy.

Auctions often lead to efficient and stable outcomes, regardless of whether the seller is in a weak or strong bargaining position [3]. This helps to explain the popularity of on-line auctions and the steadily growing number of sellers. Moreover, on-line auctioneers charge low listing and commission fees, attracting even individual sellers of inexpensive goods, who would otherwise have no access to conventional auctions. On-line auctions enable buyers to discover prices and participate in a full range of buying activities. In addition, since on-line auctions stimulate psychological desires to win, bidders sometimes find themselves making higher bids or buying more units than they initially planned. Overall, cost-efficiency and active participation of both sellers and buyers increase the liquidity of on-line auction markets, and this liquidity, in turn, attracts more sellers and buyers.

Despite the aforementioned trend and the reasons behind it, the increased popularity of on-line auctions has been accompanied by inevitable growing pains. According to Internet Fraud Watch, on-line auction fraud has become the number-one type of Internet fraud over the last three years [1]. In 1999, on-line auction fraud accounted for 87 percent of reported incidents, up from 68 percent in 1998. Indeed, on-line auctions seemed to attract fraud. Between 1998 and 1999, fraud related to on-line auctions soared by 76 percent, while fraud related to other types of on-line transactions plummeted by 44 percent. In the year 2000, the volume of auction-related fraud increased 23 percent and accounted for 78 percent of the reported incidents. Up to 2001, 31 percent of on-line Americans, or approximately 35 million people, participated in on-line auctions, and among them, 41 percent of on-line auction buyers reported having a problem [4]. In addition to these statistics, many cases of on-line auction fraud have become headline news, such as the malicious million-dollar bid over a Y2K related domain name and shill bidding on more than 1,100 art items in eBay [7, 14].

Uncollectable payments, undelivered products, inaccurate descriptions of goods, and fake comments undermine the trust between auction sellers and buyers and affect the credibility of on-line auction houses. These forms of fraud are not limited to auctions. They exist for fixed-priced on-line sales as well. However, the present research focuses on fraud rooted in the uniqueness of on-line auctions, that is, variable pricing executed in a distributed computing environment.

Variable pricing aims to discover the efficient and stable price for a product when buyer valuations of the product are unknown to the seller. Such flexibility leaves room for illegitimate price manipulation. The distributed computing environment of the Internet provides new ways to illegitimately manipulate prices that are often exploited by dishonest sellers and bidders.

Bidding under a false identity is a popular scam that exploits the lack of authentication on the Internet. If a real-world identity creates multiple on-line login names that are simultaneously active on one Web site, only one should be considered legitimate for a given transaction and all the others are regarded as false identities for the same transaction. In other words, only one-to-one

mapping is allowed between an on-line identity and a real-world counterpart for a particular Web trade. This requirement is needed because the design of a Web transaction often assumes, just as in a face-to-face physical transaction, that each trader can be identified with a unique driver's license, social security number, or business license. In the virtual world, it is difficult to verify whether such a requirement is satisfied because currently there is no rigorous method to tie an on-line screen name to its real-world counterpart. In practice, Internet users sometimes create more than one on-line login name / password combination even to a single Web business. This is done for one of two reasons: a benign cause such as the loss of the password or a malicious attempt to perpetrate a fraud.

The discrepancy between the reality and the assumption by e-business designers of one-to-one on-line / off-line identity mapping creates room for fraud. In on-line auctions, sellers can disguise themselves as bidders for their own goods in order to drive up the bidding price in ascending-bid auctions. This practice is known as shill bidding [12]. On the other hand, buyers can create several identities and "collude" with themselves in an auction trade in order to maximize their own surplus. This practice is here called "false-name bidding." Shill bidding and false-name bidding are both difficult to detect in on-line auctions.

Limitations of Existing Methods Against False-Name Bidding

Many on-line auction houses are aware of the practice of shill bidding and are developing techniques against it. However, they are not yet aware of the damage that false-name bidding can do, especially in sealed-bid combinatorial auctions. The lack of discussions on auction forums or media coverage of false-name bidding does not mean that the fraud does not exist. As will be shown below mathematically and illustrated with examples, false-name bidding is a type of fraud that can shift the seller's profit to dishonest buyers who collude among false identities.

Once on-line auction houses recognize the existence and extent of false-name bidding, they will probably try to combat it by adapting the techniques currently used against shill bidding, since the two frauds have many similarities. But techniques that rely on computing technology of some kind, such as IP address monitoring or detection through statistical analysis, are far from effective and sufficient, because dishonest agents who are technically savvy can circumvent them. For example, the distributed nature of the Internet allows a dishonest agent to compromise a legitimate user's computer and use the stolen IP address. In fact, IP address monitoring is not a viable approach at all. Many companies and ISPs use proxy servers and dynamic IP addresses, which means that there can be many users for a single IP address. Statistical analysis is more useful. However, the lack of authentication on the Internet allows one user to create an unlimited number of on-line identities and bidding records. As a result, finding correlations among these identities would be very difficult, even with powerful data-mining tools. In addition to the limitations of computing

and statistical methods, on-line auction houses like eBay only keep the most recent bidding records, restricting the scope of fraud investigation.

Although there are state-of-the-art cryptographic protocols to enhance on-line authentication, they may verify false identities and cannot prevent them from damaging the desired system outcome. For instance, standard on-line Web transactions involve Secure Socket Layer (SSL). However, SSL mostly authenticates the identity of the Web business to the Web user, not the other way around. If the full functionalities of SSL are used—that is, if every Web user needs to subscribe to a digital certificate to authenticate himself cryptographically to a Web business—the burden of verifying the binding between on-line identity and real-world identity will shift from the Web business to a Certificate Authority like VeriSign. Unfortunately, the certification business is by no means as rigorous as the entities issuing Social Security numbers or driver's licenses in the physical economy. At present, the user-side authentication of SSL is optional, and the majority of Web users choose not to use it. Thus, although technical solutions to enhance on-line authentication are available, the way they are being used is far from what the design engineers had in mind.

Nevertheless, most of the on-line authentication processes applied in current practice, such as using the .htpasswd and .htaccess facilities of Unix/Linux Web servers, only check whether a login name/password combination maps with a pair already registered on the Web server or the user database of the Web business. These authentication methods, however, cannot tell which login names are false identities. To control false identities, a Web business can make the registration of on-line login names stricter by asking for social security numbers, driver's licenses, or other unique user information, but this may scare users away because it will raise doubts about whether the Web business can keep their critical information secure and private. Once a hacker gains access to a user's sensitive information, such as a Social Security number, to perpetrate fraud, the victim's life will be miserable. To avoid storing and protecting users' sensitive information and at the same time not annoy users or discourage registration, most Web registrations only require certain noncritical user information, such as name, mailing address, and e-mail addresses. Credit card information is requested only when the user is making an on-line transaction, but not for registration. Non-critical user information can be stolen, manipulated, or even created on the fly, which makes it difficult to pinpoint a registered login name with a real-world identity. Besides, even when a Web business knows that a user has registered a second login name (database technology requires that each registered login name be unique), it will not prohibit the user from doing so, because it has no way of knowing the purpose of the second-time registration. The user may have fraudulent intentions, but on the other hand may simply be re-registering because of having lost the original password. A Web business may offer users the option of retrieving existing passwords, but customers often do not remember that they had an account with that business before, or will not bother to recover their passwords unless the Web business provides discounts for long-term customer loyalty, and this is not a very wide practice. For whatever reason, many users have multiple on-line identities, even for a single Web business. Those that are used for fraudulent purposes are here called false identities.

Fraudulent on-line identities are not just a problem in auctions but can cause difficulties in other on-line transactions as well. For example, many Web businesses give discounts or even free merchandise to first-time customers. This creates incentives for customers to register multiple times under false identities in order to get multiple gifts. This phenomenon has urged Web businesses to rethink how to practice this traditional marketing tactic.

Applying Mechanism Design for E-commerce Security

Opportunities

On-line auction fraud is only one aspect of the e-commerce security problem. E-commerce security is a broad issue and has become a serious concern of firms and individuals who rely on distributed digital processes in their daily operations. In May 2000, concerns of this kind were highlighted when a simple virus called "love bug" bombarded corporate and individual windows-based e-mail systems, costing up to $7.7 billion in lost productivity. More recently, in June 2001, some hackers, through the Internet, attacked Cal-ISO, the California electric grid parent company, and attempted to compile code to penetrate firewalls to access the actual grid-control computer. The intrusion went undetected for more than two weeks until the hackers had brought too much attention to themselves.

These are only two examples out of numerous on-line security breaches. Nowadays, computerized operations are widespread in daily activities and even in mission-critical tasks. But sadly, security breaches have become the by-products of e-commerce, and firms and individuals have to accept and endure the economic losses resulting from occasional digital disasters. Is there a way to eliminate or control these digital disasters that damage e-commerce systems?

To minimize and control digital disasters, rather than simply clean up the damage after they have occurred, one needs to understand the economic reasons behind them. The pervasive spread of the "love bug," for instance, was made possible by the negative effect of the network externality associated with the design of Microsoft Outlook software. Unfortunately, while there are hackers who are fond of exploiting such vulnerable externalities, software providers are surprisingly not motivated to eliminate or reduce the software faults that may be exploited by hackers. Metaphorically, Microsoft has left its door wide open and allowed polluters to dump garbage into the Microsoft river that runs through public water-supply systems worldwide.

Once one understands the economic reasons behind a digital disaster, it is possible to deploy schemes that weaken the reasons or reduce their effects. In the case of the "love bug," there are two options: either increase software providers' responsibilities or decrease the scope of negative externality by disabling the automatic accessibility to Outlook address books. To repeat the metaphor, Microsoft should better control the interface between its river and the public water system.

Fortunately, unlike natural disasters, which are the acts of God, digital disasters are man-made and consequently can be controlled by devising human

motivations through economic designs. The solution to the virus problem may reside in better quality control in software production, changes in the relationship between software providers and users, and refinement of contracts and payment schemes to increase software providers' liabilities and hence their incentives for delivering safer software. Moreover, systems can be devised to discourage people to take the role of hackers or frauds.

Mechanism design in economics has been widely deployed to promote and guard traditional commerce against man-made disasters. It uses the tools of economics and game theory to design "rules of interaction" for economic transactions that will, in principle, yield desired outcomes defined in objective functions [2]. It is a controlled game, for its designer defines and mediates the structure of the game played by participating agents, taking into account their potentially bad behaviors.

Let us return to the issue of false-name bidding. Because of the limitations of existing methods mentioned above, false name registration exists and is hard to detect. A better way to solve the problem is to discourage such dishonest behavior, and to do this it is necessary to reduce the incentives that lead users to create false identities for the purpose of fraud. Incentive design is a proactive way of dealing with security problems. It is much more cost-effective than the detective and corrective methods applied through computing enhancement and statistic analysis.

As mentioned earlier, the revelation principle in mechanism design guarantees that for each mechanism there is a counterpart with the same outcome where the agent's best strategy is truth-telling. With truthful mechanisms, agents can only lose by not revealing their true intentions. Truthful mechanisms are especially important in e-commerce because current practices cannot tell whether or not participants reveal their true intentions, or even their true identities.

Following the revelation principle in mechanism design, trading processes can be redesigned to discourage participating agents from using the anonymity of the Internet maliciously. It is an approach to modifying their incentives to voluntarily reveal their true identities. This compliance is more important than the typical concept of authentication—verify that one is who one claims to be (i.e., who one registers oneself as). Hence, it will be more effective than the ordinary practice of authentication.

In addition to helping to solve security problems, mechanism design is a method well suited to the e-commerce environment in other ways. Mechanism design involves distributed agents interacting in economic transactions. This feature fits naturally with e-commerce, where agents are, in fact, represented through distributed digital processes implemented with software. Most agents can even be automated, implemented in distributed computing agent technology like Java, where the human agents' intentions and decision structures are embedded within the computer agent software.

Since mechanism design works best when the participating agents are rational, the outcome of mechanism design for an e-commerce system will be much more predictable if the system only deals with automated computing agents. This is because automated computing agents can behave rationally if designed to do so. The issue of complex decisions is excluded, since in such cases it is difficult or even impossible to program rational behaviors. Never-

theless, by contrast, human participants are not consistently rational. A person in an emotional state, for instance, may not act rationally.

Another benefit from using automated processes to conduct interactions in mechanisms is the reduction of the strategy space, as compared with human interactions. At least for now, digital processes cannot "observe" or "demonstrate" threats or eagerness to "opponent" digital processes. Thus, a less complicated mechanism design may be acceptable for e-commerce where software agents take certain roles of human players.

Challenges

Despite the fit between mechanism design and e-commerce, challenges do exist. Mechanisms that function well in traditional commerce may need to be updated to accommodate new perspectives raised in e-commerce.

As an example, traditional mechanism designs assume a fixed set of known agents (e.g., sellers and buyers in a transaction) who normally look out for their own interests with precise objective functions. However, in e-commerce, this set has to be extended to include unknown or unrevealed agents. These fall into two categories: (1) additional agents created by the participating agents for some strategic purpose, such as the false identities created by auction sellers or buyers for fraudulent purposes, (2) external malicious agents (e.g., hackers) who wish to disturb the mechanism, cause problems for other agents, and manipulate the process outcome.

The need for mechanism updates is inevitable because environmental changes require modification of trading rules [13]. Traditional mechanisms, if applied for on-line processes, have to be rethought with respect to distributed computing issues, including authentication problems.

Another inevitability is the tradeoff between optimization results and added constraints. The traditional mechanisms discussed in economics are often first-best choices. They are based on mathematical optimization, abstract from execution environment and technology. Considerations related to special environments and technologies may necessitate the sacrifice of all or part of some objectives, such as a reduction in the social surplus. This comes about if the added constraints worsen the optimization results. As discussed below, the protocol suggested here enhances security but sometimes reduces the social surplus.

To achieve certain desirable properties, such as resilience to on-line fraud, tradeoffs have to be made to the best traditional mechanisms. The goal here, therefore, is to search for second-best mechanisms that are still nearly optimal in the e-commerce setting. The discussion will now proceed to a mathematical explanation of why false-name bidding in on-line auctions is harmful and will introduce a modified mechanism against this type of fraud.

False-Name Bidding in On-line Auctions

Most on-line auction houses have adopted traditional auction mechanisms. Although these function well in traditional commerce, they have several short-

comings when applied in e-commerce. As a result, on-line auction houses have been struggling to prevent fraud because their trading rules are not designed to handle on-line fraud. In addition, the attraction offered by on-line auction houses—operational efficiency—is a double-edged sword. Compared to traditional physical auctions, like the ones at Sotheby's, they charge very low listing fees and commission rates. This encourages such frauds as shill bidding. Sellers only incur a low risk in bidding on their own items, since a seller whose bid wins only loses the low listing and commission fees, whereas a seller who is outbid by someone else attains a substantial increase in utility [12].

To fight against on-line auction fraud, existing auction mechanisms need to be examined and even modified with respect to the special conditions attendant to Internet auctions. A new fee schedule has been suggested as a way to control the seller's incentives for shill bidding [12]. Similarly, auctions can be redesigned to discourage false-name bidding.

As a start, it is necessary to understand why false-name bidding is profitable to dishonest agents. The following examples show how a dishonest bidder can increase utility in multi-unit auctions by submitting false-name bids.

Effects of False-Name Bidding on Ascending-Bid Auctions

In ascending-bid auctions, bidders can create false identities and bids to indirectly secure a winning position and block other bidders. This is also called bid-shielding. For example, suppose there are two printers for auction and two buyers to bid. Bidder A only wants to buy one printer for $80, whereas B wants to buy two printers for $100 each. Under normal circumstances B would buy both printers. However, A places an $80 bid first, immediately after that creates a false identity F, and using this identity submits a bid for two printers for $120 each. By the time B tries to bid, the false bid is already in, and B's only alternative is to outbid F. If unwilling to pay that much, B must leave the auction. After the bid closes and the auctioneer fails to collect the payment from F, the auctioneer sells the printer to the next-highest bidder, who is A.

Effects of False-Name Bidding on Sealed-Bid Auctions

Consider the Generalized Vickrey Auction (GVA) protocol [8]. GVA allocates a set of goods to a set of agents. This allocation and the agents' corresponding payments are all based on the agents' declared valuations of each subset of goods. The GVA protocol allocates the goods to maximize the social surplus, which is the sum of all the valuations. For each agent it is also calculated what social surplus could be achieved had the agent not been present. The difference between the optimal surplus and the surplus that can be achieved in the agent's absence is subtracted from the agent's bid to calculate the agent's payment. In effect, all the bidders get credit for their contributions to the social surplus.

Let M be the set of goods to be auctioned, $N = \{1,\ldots n\}$ be the set of agents, $G: M \mapsto \mathrm{N} \cup \{0\}$ be an allocation of goods to agents, $G^{-1}(\{x\})$ be the set of goods allocated to agent x, and $v_x(G)$ be agent x's declared valuation for the set of

goods allocated. The allocation under the GVA is the optimal allocation G^*, that maximizes $\sum_{x=1}^{n} \upsilon_x (G)$. The payment of agent x is:

$$\begin{aligned} p_x &= \upsilon_x - \left[\sum_{y \in N} \upsilon_y (G^*) - \sum_{y \in N \setminus \{x\}} \upsilon_y (G^*_{\sim x}) \right] \\ &= \sum_{y \in N \setminus \{x\}} \upsilon_y (G^*_{\sim x}) - \sum_{y \in N \setminus \{x\}} \upsilon_y (G^*) \end{aligned} \tag{1}$$

where $G^*_{\sim x}$ is the optimal allocation when agent x does not participate in the auction.

Clearly, when there is only one item for auction, the GVA mechanism assigns it to the highest bidder, who pays the second-highest bid, which is identical to the classic Vickrey auction protocol.

In the GVA protocol, truth-telling is an optimal strategy, which means that one can never gain any advantage by not declaring one's true valuation. As is clearly shown in Equation (1), for a given allocation each agent's payment is independent from the agent's own valuation.

The GVA mechanism assumes that agents try to maximize their own valuations without colluding with other agents. However, the assumption no longer holds when an auction is conducted over the Internet, where agents can easily collude among their fictitious identities. In such a setting GVA is no longer incentive-compatible because an agent can obtain a reduced payment by submitting bids under false names. In other words, even though agent x's valuation does not appear in Equation (1), x can still manipulate a payment through valuations under false identities. The following example from Yokoo et al. illustrates this problem [15]:

> *Example 1. The GVA is not incentive-compatible in the presence of false-name bidding.*

Suppose that there are two concert tickets to be auctioned, and two bidders, Alice and Bob. Alice would like to buy two tickets for $10 each, but she does not want to go to the concert alone, so a single ticket is worthless for her. Bob would like to buy two tickets for $15 each and does not mind going alone, so he would be willing to pay $15 for a single ticket. From (1) it follows that Bob will get both tickets and has to pay $20. But Bob can get the tickets cheaper by cheating. He can create a false identity, Charlie, who will bid $15 for a single ticket, and Bob himself will also bid $15 for one ticket. Now the optimal allocation gives the tickets to Bob and Charlie, which gives $30 social surplus. But if one of them had not bid, the optimal allocation would have given both tickets to Alice for $20 social surplus. This means that Bob and Charlie each contribute $10 to the social surplus, thus both get that much discount from their bids, so they pay only $5 a ticket. In this way, Bob can get two tickets for half the price he would normally have to pay, by "cashing" the social surplus credit twice.

The following example shows even more adverse effects of false-name bidding:

Example 2. The GVA is not efficient in the presence of false-name bidding.

In a GVA auction, there are two tickets for sale. Suppose that all the bidders except Bob bid truthfully. Let v_1 be the highest valuation of a single ticket among the truthful bidders, and let v_2 be the highest valuation for the bundle of two tickets. It is assumed that $v_2 > 2(v_1 + \varepsilon)$ for some $\varepsilon > 0$. This is usually true, because most people do not like to go to concerts alone. If Bob submits two bids under two different bidder identities at $v_2 - v_1 - \varepsilon > v_2/2$ each, he would win both tickets. If either of Bob's two bids were removed, the truthful bidder with valuation v_2 would win, because Bob's remaining bid combined with the highest truthful single bid, v_1, only adds up to $v_2 - \varepsilon > v_2$. This means that both of Bob's bids get $v_2 - 2v_1 - 2\varepsilon$ discount, so Bob's total payment is $2(v_1 + \varepsilon)$. Therefore, if Bob's actual valuation for two tickets is more than $2v_1$, he has an incentive to submit two bids for a single ticket each at more than $v_2/2$, even if his valuation is less than $v_2/2$, which means that the allocation does not maximize the social surplus. Moreover, the higher Bob's two identities bid, the less they pay, as long as the bids are less than $v_2 - v_1$.

Let us take a numerical example. Suppose that Alice bids \$12 for one ticket and \$30 for two. Bob also wants two tickets, but does not want to pay more than \$27. Bob has a lower valuation than Alice, but could still win by submitting two bids for \$17 for a single ticket. This way Bob would win both tickets and pay only \$26. In more detail, the auctioneer would calculate that the social surplus corresponding to this allocation is \$34, based on Bob's false bids. If one of Bob's bid were removed, the optimal allocation would give both tickets to Alice for \$30 social surplus, because it is better than giving one ticket to Bob's other identity and one ticket to Alice, which would only achieve \$17 + \$12 = \$29 social surplus. This means that each of Bob's identities can receive \$4 discount for their contribution to the social surplus. Bob ends up paying only \$26 for the two tickets. The two tickets are not given to Alice, the bidder with the highest valuation. If Bob had bid truthfully, the auctioneer would have gotten \$27 payment, or more in case there were some other bidders with valuations higher than Bob's. With Bob's false-name bidding, both the efficiency and the seller's income are reduced.

This shows that truthful bidding is no longer an equilibrium strategy in GVA in the presence of false-name bidding, and the GVA with false-name bidding does not always maximize the social surplus. Moreover, in Example 2 Bob has an incentive to overbid his true value to reduce his payment.

Robust Auction Protocols Against False-Name Bidding

The Leveled Division Set Auction Protocol

The Leveled Division Set (LDS) auction protocol introduced by Yokoo et al. is robust against false-name bids because an agent cannot profit from creating false identities and submitting bids under fictitious names [15]. It is the intricacy of handling leveled division sets that makes the protocol robust. For a complete proof, see the work by Yokoo, Sakurai, and Matsubara [15]. They also

prove that LDS remains incentive-compatible (i.e., since one's payment is not affected by one's own valuation, the best strategy is to bid one's true valuation) and individually rational (i.e., one's payment never exceeds one's valuation).

The LDS protocol applies to combinatorial sealed-bid auctions. It predefines a set of divisions of goods and a reserve price for each auctioned item. A division D contains disjoint sets of goods, that is, $D \subset P(M)$, where $P(M)$ is the power set of M, and $\forall S \forall S'(S \in D \wedge S' \in D \Rightarrow S \cap S' = \emptyset)$. An allocation G is said to be allowed under D if each set of goods allocated to agents is in D, that is, $\forall x(x) \in N \Rightarrow G^{-1}(\{x\}) \in D \cup \{\emptyset\})$.

The divisions of goods are arranged into leveled division sets. The first level has a single division $\{M\}$ where the only allowed allocation is to allocate all the goods to one agent. Each subsequent level is a "refinement" of the previous one. The LDS protocol introduces a recursive algorithm that picks a level of divisions and applies a modified GVA protocol that chooses from a restricted set of allocations allowed under one of the divisions of that level. Unallocated goods are allocated to a dummy agent whose valuation for each good is equal to its reserve price.

If there is no false-name bidding, the GVA protocol maximizes social surplus. Otherwise, as seen in Example 2, the GVA is not efficient. In addition, the winner-determination problem in GVA is NP-hard, which means that it is practically impossible to calculate the optimal allocation of goods, even for as few as 30 bidders, because of the computational and communication complexity of the GVA protocol [6].

In the LDS protocol, on the other hand, the achieved social surplus and the computational complexity depend on the choice of leveled division sets and the reserve prices. Since the LDS protocol only has to consider allocations allowed under some division from the leveled division sets, it is computationally less complex than the GVA protocol, which has to consider all possible allocations.

But Yokoo et al. do not explain how to choose the best leveled division sets and reserve prices that either minimize computation complexity or maximize social surplus, or both. It is not known how the choice of parameters affects the achieved social surplus and computational complexity of the LDS protocol, and how they compare to the optimal allocation achieved by the GVA protocol under no false-name bidding.

To avoid the complexity of GVA and LDS, a mechanism is introduced below that only deals with identical goods. Specialization of this type greatly simplifies the algorithm.

The Leveled Partition Set Protocol

Overview

Whether or not an economic mechanism is applicable in the real world often depends on its complexity and its trade-offs with the expected outcome. Striving for the simplicity and efficiency of a mechanism is important, especially when the mechanism is applied in e-commerce, which is full of uncertainties. It is important not only because it is easier to implement and more practical to

use, but also because it is easier to verify the correctness of the implementation. Verification is critical to the integrity of e-commerce [10].

Although better than GVA with respect to security, LDS is so complicated that it may be impractical to implement and verify through feasible techniques. Therefore, a specialized auction protocol is introduced that is not only robust against false-name bidding but also less complex.

Definition

The Leveled Partition Set (LPS) protocol is a modified version of the LDS protocol. LPS applies to sealed-bid auctions selling multi-unit identical goods. Let k be the number of items for sale. An allocation defines the number of goods obtained by each agent; that is, it is a function $A : N \mapsto \{0, \ldots k\}$ such that $\sum_{x=0}^{n} A(x) = k$. Agent x gets $A(x)$ goods, and $A(0)$ goods remain unallocated. A partition P is a multiset of positive integers whose sum is less than or equal to k.[1] An allocation A is allowed under a partition P if the number of goods received by each agent is either zero or an element of P. In addition, if the number of agents who get exactly i goods allocated is j, then i must appear in P at least j times. In other words, the multiset $\{A(x) : x \in \{1, \ldots n\} \wedge A(x) > 0\}$ is a subset of P.

A leveled partition set is defined as follows:

1. For each level $i = 1, \ldots max\_level$, a partition set $SP_i = \{P_{i1}, P_{i2}, \ldots\}$ is defined.
2. $SP_1 = \{\{k\}\}$
3. For each partition $P_{ij} \in -SP_i$, the sum of any two or more elements of P_{ij} is an element of a partition of an earlier level, that is, $\forall P$ (P is a multiset $\wedge (P \subset P_{ij}) \wedge (\mid P \mid \geq 2) \Rightarrow \exists l \exists m (l < i \wedge \Sigma P \in P_{lm})$).

For an allocation A, each agent x participating in the auction declares a valuation $v_x(A)$ that only depends on the number of goods the agent would like to receive. In addition, the seller sets a reserve price r for a single item, the minimum price at which the agent is willing to sell. The seller's valuation is defined as $v_0(A) = rA(0)$. The utility of agent x under an allocation is the declared valuation minus the payment: $u_x(A) = v_x(A) - p_x(A)$. The utility of the seller is the total collected payment minus the reserve price of the sold goods: $u_0(A) = \sum_{x=1}^{n} p_x(A) - (n - A(0))r$. The social surplus is

$$
\begin{aligned}
U(A) &= \sum_{x=1}^{n} u_x(A) + u_0(A) \\
&= \sum_{x=1}^{n} \upsilon_x(A) - (n - A(0))r \\
&= \sum_{x=0}^{n} \upsilon_x(A) - nr.
\end{aligned}
$$

The LPS protocol is a recursive algorithm, as follows:

Procedure LPS(i):

1. If there is no agent whose valuation is at least the reserve price under an allocation A allowed under a partition from SP_i, apply LPS(i+1), or terminate if i=*max_level*.
2. If there is exactly one agent whose valuation is at least the reserve price under an allocation A allowed under a partition from SP_i, only this agent will obtain some goods. The agent gets the amount of goods that maximizes utility, considering all allowed allocations under SP(i) , where the agent pays the reserve price, and the allocation under LPS(i+1), where the agent pays the determined price.
3. Otherwise chose an allocation A^* allowed under SP_i that maximizes the social surplus $U(A)$. The payment of agent x is

$$p_x = U(A^*_{\sim x}) - U(A^*) + v_x(A^*) \tag{2}$$

$$= \sum_{y \in (N \cup \{0\}) \setminus \{x\}} (v_y(A^*_{\sim x}) - v_y(A^*)) \tag{3}$$

where $A^*_{\sim x}$ is the allocation allowed under SP_i that does not allocate anything to agent x and maximizes $U(A)$.

Proof

Lemma 1. One never pays more than the declared valuation for the goods one obtains.

Proof. This follows immediately from $U(A^*_{\sim x}) \leq$ U (A*), which is true since A^* maximizes $U(A)$.

Lemma 2. Under LPS each agent pays at least the reserve price for the obtained goods.

Proof. Let A' be the allocation that is the same as A^* except that the goods allocated to agent x in A^* are unallocated under A', therefore $A'(0) = A^*(0) + A^*(x)$. Due to the maximal choice of $A^*_{\sim x}$, $U(A^*_{\sim x}) \geq (A')$ and (2), the following holds:

$$\begin{aligned} p_x &\geq U(A') - U(A^*) + v_x(A^*) \\ &\geq v_0(A') - v_0(A^*) \\ &\geq r(A'(0) - A^*(0)) = rA^*(x). \end{aligned}$$

Lemma 3. Under LPS an agent cannot increase utility by submitting false-name bids.

Proof. Assume that agent x can improve utility by submitting bids under two identities x and x', and the final allocation is A'. Assume that $A'(x) > 0$

and $A'(x') > 0$. This is only possible if the allocation is a result of step 3 of LPS(i) for some i. From lemma 2, $p_x + p_{x'}' \geq r\,(A'(x) + A'(x'))$. From part 3 of the definition of the LPS, there exists $j < i$ such that $A'(x) + A'(x') \in P_{jm}$ for some $P_{jm} \in SP_j$. In this case agent x would have been better off submitting a single bid at the reserve price for $A'(x) + A'(x')$ goods, resulting from an allocation based on step 2 of LPS(j).

Lemma 3 means that one can only lose by splitting up one's bid.

Examples

In Example 1, discussed earlier, if both Bob and Charlie submit a $15 bid for a single ticket, neither of them will win. Only bids for two tickets are considered in level 1, so if neither Bob nor Charlie bids for two tickets, Alice would get the tickets at the reserve price. On the other hand, if Bob bids $30 for two tickets, he will win and will have to pay $20, which is Alice's bid.

The situation in Example 2 is similar when using the LPS protocol. Since Alice has a valid bid for two units, her bid will be considered at level 1. Bob can only win if he submits a bid for two units more than Alice's bid. But in this case Bob would have to pay Alice's bid, which is more than his valuation. Bob cannot win by splitting his bids because single-unit bids can only win if there is no valid bundle bid.

Computational Complexity

As mentioned earlier, an auctioneer who wishes to calculate the allocation of goods in a GVA auction has to solve an NP-hard problem [6]. There have been several attempts to come up with computationally manageable combinatorial auction mechanisms that are still close to optimal, some of which are surveyed by de Vries and Vohra [9]. Computational experiments show that the GVA protocol can be computationally infeasible depending on the bids received (e.g., for 50 items with 30 bidders) [5]. Despite this complexity, GVA may still be inefficient in the presence of false-name bidding. In general, the LDS protocol is less computationally intensive because it restricts the set of possible allocations. LPS is even simple, for the sale items are identical. Therefore, one need only consider the number of units allocated to each winner, and the set of possible allocations is further reduced by the choice of partition sets. There is a tradeoff between computational complexity and the achievable social surplus. Partition sets that allow more allocations improve the social surplus, but also require more computing resources. The exact computational complexity of LPS also depends on the bids received and the predefined partitions.

Conclusion

A series of research projects was conducted aimed at introducing a new way of looking at e-commerce security that integrates economic reasoning (e.g.,

mechanism design) with computing advances (e.g., program verification). Approaching security issues in this integrated fashion is the best way to understand and mitigate a very real and potentially damaging threat to the growth of e-commerce. The goal is to provide a full package of the tools and methods necessary for designing and verifying economically feasible and technically correct e-commerce systems.

The authors have researched the question of how to enhance e-commerce security from the computer science perspective using modern formal verification techniques (model-checking) to assure correct technical design and implementation of e-processes, given a set of system specifications and procedures [10].

While model-checking can eliminate or reduce the impact of agents' bad behavior, mechanism design can proactively discourage agents from making undesirable choices. This is why designers and developers, when defining system specifications and procedures, should take account of the special issues pertaining to the Internet and ensure that mischievous user behaviors are not allowed or at least discouraged to prevent them from compromising the desired process outcome. As demonstrated in the paper, very often such assurance can only be obtained by redesigning the trading rules.

This paper analyzes a special security issue: false-name bidding in on-line auctions. Bidders can increase their utility by splitting their bids for a large bundle to several smaller bundles under false identities. As illustrated by examples, the Generalized Vickrey Auction protocol is neither incentive-compatible nor efficient in the presence of false-name bidding. To counteract this problem this paper introduces a Leveled Partition Set (LPS) protocol for multi-unit sealed-bid auctions. LPS is incentive-compatible, robust against false-name bidding, and less computationally complex than the GVA and LDS protocols. In follow-up research, the authors have also developed an auction mechanism called the Binary Vickrey Auction (BVA) protocol that achieves much the same goals as LPS. BVA is dramatically different from LPS but is also robust against false-name bidding. It provides more efficient allocation, but in doing so partially sacrifices incentive-compatibility [11].

Although mechanism design and formal verification are totally different methods, they are connected. Mechanism design is a supplementary to formal verification. It enables designers to better anticipate the behavior of agents, analyze their types and choice space, and, more important, modify the incentives affecting them.

Continuing research by the authors will study how mechanism design and formal verification can be integrated and how current management practices can be changed to fully benefit from these advanced techniques. As a demonstration of tools and methods, the authors plan to implement the mechanisms developed in an on-line ticket-sale example with dynamic pricing. The prototype involves a set of distributed digital processes interacting over the Internet in real time on behalf of the participating agents: sellers, buyers, and intermediaries. These agents provide input/output to the auction system and have the potential to significantly deteriorate auction performance. Formal verification methods will be used to mathematically verify the correctness of the new mechanisms and their implementations, and hence

demonstrate the applicability of the principles outlined in this paper in real e-commerce applications.

NOTE

1. A multiset is a set where the same element can appear more than once.

REFERENCES

1. Internet Fraud Watch. 2000 internet fraud statistics. 2001 (www.fraud.org/internet/lt00totstats.htm).
2. Mas-Colell, A., and Whinston, M.D. *Microeconomic Theory*. New York: Oxford University Press, 1995.
3. Milgrom, P.R. Auction theory. In *Advances in Economic Theory: Fifth World Congress*. Cambridge: Cambridge University Press, 1987, pp. 1–32.
4. National Consumers League. On-line auction survey summary. 2001 (www.nclnet.org/on-lineauctions/auctionsurvey2001.htm).
5. Parkers, D.C., and Ungar, L.H. Iterative combinatorial auctions: Theory and practice. In *Proceedings of the 17th National Conference on Artificial Intelligence*, August 2000, pp. 74–81.
6. Rothkopf, M.H.; Pekec, A.; and Harstad, R.M. Computationally manageable combinatorial auctions. *Management Science, 44*, 8 (1998), 1131–1147.
7. Schwartz, J., and Dobrzynski, J.H. 3 men are charged with fraud in 1,100 art auctions on ebay. *New York Times*, March 9, 2001, A1.
8. Varian, H.R. Economic mechanism design for computerized agents. In *Proceedings of the First USENIX Workshop on Electronic Commerce*. New York, July 1995, pp. 13–21.
9. de Vries, S., and Vohra, R.V. Combinatorial auctions: A survey. Working paper, June 2001 (www.m9.ma.tum.de/~devries/comb_auction_supplement/comauction.pdf).
10. Wang, W.; Hidvégi, Z.; Bailey, A.D., Jr.; and Whinston, A.B. E-process control and assurance using model checking. *IEEE Computer, 33*, 10 (October 2000), 48–53.
11. Wang, W.; Hidvégi, Z.; and Whinston, A.B. BVA—a protocol against false-name bidding in multi-unit auctions. Working paper, June 2001.
12. Wang, W.; Hidvégi, Z.; and Whinston, A.B. Shill bidding in English auctions. Working paper, April 2001.
13. Wilson, R. Game-theoretic analyses of trading processes. In Truman Bewley (ed.), *Advances in Economic Theory: Fifth World Congress*. Cambridge: Cambridge University Press, 1987, pp. 33–70.
14. Year2000.com. Year2000.com auction and reward information. 2000 (www.year2000.com/auction/).
15. Yokoo, M.; Sakurai, Y.; and Matsubara, S. Robust combinatorial auction protocol against false-name bids. In *Proceedings of the 17th National Conference on Artificial Intelligence*, August 2000, pp. 110–115.

WENLI WANG (Wenli_Wang@bus.emory.edu; wenli@ede02.bus.utexas.edu) is a visiting assistant professor in the Department of Decision and Information Analysis at the Goizueta Business School, Emory University, and a visiting scholar at the Center for Research in Electronic Commerce. She received her B.Sc. in computer networking and telecommunications at the Beijing University of Posts and Telecommunications in 1994 and her Ph.D. in management information systems from the University of Texas at Austin in 2000. Dr. Wang's research concentrates on electronic commerce security, control, and assurance services. Her current studies focus on Internet auctions and the use of economic principles and computing advances to enhance reliability and security in e-businesses and e-markets. She also has a general interest in the impact of emerging technologies on the digital economy and vice versa.

ZOLTÁN HIDVÉGI (hzoli@ede02.bus.utexas.edu) is a doctoral student in an interdisciplinary study between computer science and management information systems at the University of Texas at Austin. He is also a software engineer at IBM Corporation in Austin, where he writes cycle simulation software for the functional verification of chip designs. He received his M.Sc. in mathematics from Eötvös University in Budapest. His current research interests focus on formal verification and economics.

ANDREW B. WHINSTON (abw@uts.cc.utexas.edu; crec.bus.utexas.edu) holds the Hugh Cullen chair of information systems, computer science, and economics at the University of Texas at Austin. He obtained his Ph.D. in management from Carnegie-Mellon University in 1962 and subsequently held academic posts at the University of Texas at Austin, Purdue University, the University of Virginia, Yale University, and the University of California, Los Angeles. Over the years, his research has explored artificial intelligence, information systems, decision support systems, and electronic commerce. Currently it spans a range of issues in electronic commerce, including resource allocation, bundle markets, trust and assurance, marketing and market design. He is editor-in-chief of *Decision Support Systems*, a member of the editorial board of *Annals of Mathematics and Artificial Intelligence*, and editor-in-chief of the *Journal of Organizational Computing and Electronic Commerce*. He is also the director of the Center for Research in Electronic Commerce.