Secure Vickrey Auctions without Threshold Trust

Helger Lipmaa[1], N. Asokan[2], and Valtteri Niemi[2]

[1] Laboratory for Theoretical Computer Science
Department of Computer Science and Engineering
Helsinki University of Technology
P.O.Box 5400, FIN-02015 HUT, Espoo, Finland
{helger}@tcs.hut.fi
[2] Nokia Research Center
P.O.Box 407, FIN-00045 NOKIA GROUP, Finland
{n.asokan,valtteri.niemi}@nokia.com

Abstract. We argue that threshold trust is not an option in most of the real-life electronic auctions.We then propose two new cryptographic Vickrey auction schemes that involve, apart from the bidders and the seller S, an auction authority A so that unless S and A collude the outcome of auctions will be correct, and moreover, S will not get any information about the bids, while A will learn bid statistics. Further extensions make it possible to decrease damage that colluding S and A can do, and to construct $(m+1)$st price auction schemes. The communication complexity between the S and A in medium-size auctions is at least one order of magnitude less than in the Naor-Pinkas-Sumner scheme.

Keywords: cryptographic auction schemes, homomorphic encryption, range proofs, Vickrey auctions.

1 Introduction

Vickrey auctions [Vic61] are sealed-bid auctions where the highest bidder is awarded the item, but is required to pay only the second-highest bid. Despite attractive theoretical properties, Vickrey auctions are relatively rarely used in practice since a cheating seller could either (1) change the outcome of auctions or (2) reveal bidders' private information. As argued in [RTK90,RH95], in the first case, a honest bid taker will not choose a Vickrey auction, while in the second case, a cheating bid taker eventually destroys the trust on which the use of Vickrey auctions depends. Therefore, Vickrey auctions are certainly more widely applicable when secured cryptographically, so that the seller is forced to follow the auction mechanism and no extra information is revealed to him. Attractive properties of Vickrey auctions together with these observations have motivated a huge body of research on cryptographic Vickrey auction schemes, starting with [NS93].

Now, most of the cryptographic auction schemes work in one of the following two trust models: (1) The *threshold (symmetric) trust model* where the tasks of the seller are executed by $N > 1$ servers, with less than $N/3$ (or $N/2$, depending on the precise cryptographic model) servers assumed to be dishonest; and (2) The *two-party (asymmetric) model* with a seller S and an auction authority A, where at least one of S and A is assumed to be honest. In this model, S and A are usually assigned complementary duties and are supposed to verify the actions of each other.

We argue that the threshold trust model is not suitable for many real-life electronic auctions. This trust model requires that each replicated server should be run by an independent auction authority, of which a majority is more trustworthy than the seller. The number of such trusted authorities is likely to be rather small compared to the number of distinct sellers. Therefore, every auction authority will participate in many auctions conducted by many different sellers. But in each auction, this authority has to do the same amount of work as the sellers. That may quickly lead to the auction authorities becoming bottlenecks either in the sense of security or efficiency. Simplistic use of the threshold trust approach is therefore not scalable in the case of applications like electronic auctions. (See also [NPS99] for additional motivation.)

In this paper we propose two different schemes that work in the two-party model. The first scheme is intended to illustrate the basic properties of this model. In this scheme, S blindly shuffles encrypted bids before forwarding them to A. After that, A computes the second highest bid X_2 and sends it together with a pointer to the winner's encrypted bid to S. The seller S then identifies the winner. At the end, any bidder can complain if he believes the result to be incorrect. In particular, if all bidders confirm the linear order between their bid b and X_2 (i.e., whether $b < X_2$, $b = X_2$ or $b > X_2$), A becomes accountable for his actions. However, this simple scheme has several vulnerabilities that we outline later.

The main contribution of this paper is *the homomorphic auction scheme*. In this scheme, a bid b is encoded as B^b, B being the (maximum allowed) number of bidders. The ith bidder encrypts his bid b_i with A's public key in a suitable homomorphic encryption scheme, and sends it to S. S multiplies all the received encrypted bids, and sends the resulting encryption $B^{\sum_i b_i}$ to A. After decrypting this result, A finds out the bid statistics (that is, how many bidders bid b for any possible bid b) but is not able to connect any bidders with their bids. Then, A sends the second highest bid to S. Every action in this scheme is accompanied with an *efficient* (statistical) zero-knowledge correctness proof. By using recently proposed cryptographic range proofs, we achieve that both the bidder-seller and the seller-authority communication complexity are of order $\Theta(V \cdot \log_2 B)$ bits, where V is the maximum possible number of different bids. In medium-size auctions this results in amount of interaction, at least one order of magnitude less than in the Naor-Pinkas-Sumner scheme [NPS99] that was the only previously known secure Vickrey auction scheme without threshold trust.

Our schemes use a few cryptographic observations that might be of independent interest. First, the homomorphic scheme uses a property of some known homomorphic public-key cryptosystems that we call *coin-extractability*; the same property might be also important in other applications. (In the context of auction schemes, coin-extractability of the Paillier cryptosystem was already used in [BS01].) We propose a range proof in exponents that corrects a few mistakes in the proof system from [DJ01, Section 5]; it seems to be the most efficient currently known scheme with perfect zero-knowledge that works with arbitrary intervals; a more efficient statistical zero-knowledge proof system that works only with a prime base was recently proposed in [Lip01]. Based on either of these proof systems and a recent range proof from [Lip01] we describe an efficient noninteractive statistical zero-knowledge proof system for proving that an encrypted value is the second highest value in some set. This proof system is used by A to prove that he computed a correct second-highest bid X_2 and can be easily extended to prove that an encrypted value is the $(m+1)$st highest value for a small $m > 1$. This results, in particular, in efficient $(m+1)$st price auctions.

Road-map. We start with a short overview of the existing auction schemes in Section 2. We give the necessary (cryptographic) preliminaries for the rest of the paper in Section 3. In Section 4, we describe several auxiliary protocols for homomorphic public-key cryptosystems. Our new auction schemes are described in Section 5, some extensions to them are given in Section 6, followed by some discussion in Section 7. We compare our schemes with the Naor-Pinkas-Sumner scheme in 8, and conclude the paper in Section 9.

2 State of the Art

We will briefly survey the known cryptographic Vickrey auction schemes that do not rely on the threshold trust. A few auction schemes [Cac99,BS01] are based on the Yao's millionaire's problem. Such schemes avoid threshold trust by using an oblivious third party for bid comparison. Without a collusion between the seller and the third party, the seller will get to know some partial order among the bids but not the bid values themselves. While [BS01] also discusses how to extend their auction scheme to the Vickrey auctions, at least the extension proposed in their paper would also reveal the identity of the second highest bidder. This together with the partial leak of information to the untrusted seller poses a serious problem, since it demotivates potential high bidders to participate in that type of an auction.

The auction scheme of Naor, Pinkas and Sumner [NPS99] uses a third party A (that we call *an auction authority*) and no unnecessary information is leaked unless the seller S and the third party A collude. The Naor-Pinkas-Sumner scheme bases on the two-party secure computation model of Yao, where A constructs a garbled circuit, transports it (off-line) to S and then helps S (on-line) to execute it. The circuit can be designed to completely satisfy all possible security requirements. However, A may "misdesign" the circuit to perform whatever computations he likes. A serious drawback of this scheme is that a corrupt third party can only be detected by "cut-and-choose" techniques [NPS99, Section 2.4] that would introduce a severe overhead to the protocol. The authors suggest that A (called an auction issuer in their scheme) should be an established service provider with high reputation, while S is a (usually considerably less trusted) seller. They argue that even cheating once would ruin the reputation of A.

On the other hand, even if the cut-and-choose technique is not used, circuit transfer would imply a huge communication complexity between A and S. Even if done off-line, the amount of information transferred is clearly infeasible in many real-life scenarios. Another drawback is that the circuit depends on the maximum number of bidders and hence the seller has to estimate this number before the auction relatively precisely.

Currently, [NPS99] seems to be the only published secure Vickrey auction scheme that neither reveals any unnecessary information nor relies on the threshold trust. Moreover, we are aware of only one other secure Vickrey auction scheme, recently proposed by Kikuchi [Kik01]. Kikuchi's scheme has smaller communication complexity than the Naor-Pinkas-Sumner scheme but relies on threshold trust. Moreover, the number of bidders in Kikuchi's scheme is bounded above by the number of auction servers, which makes it unusable in many practical situations. (However, it is still applicable, for example, in radio frequency spectrum or wireless spectrum license auctions, where the number of competitors is relatively small.)

The Sakurai-Miyazaki auction scheme [SM00] is secure without an *explicit* threshold-trust assumption. However, this scheme uses a bulletin board, a secure implementation of which introduces implicit threshold trust. It also bases on some relatively ad hoc security primitives. Finally, there are also schemes where threshold trust is w.r.t. the bidders. However, in these schemes, the threshold trust assumption seems to have even less ground, since in many practical cases, there is no guarantee that even a single bidder will be honest.

3 Preliminaries

Notation. Let B be the (maximum) number of bidders, let V be the (maximum) number of different bids. After an auction, let $(X_1, \dots, X_B)$ be the vector of bids in a nonincreasing order, and let Y_j be the bidder who bid X_j. Let t denote the security parameter. For a probabilistic public-key cryptosystem (G, E, D), let $c = E_K(m; r)$ denote the encryption of m by using a random coin r under the key K. In general, we denote the message space by $\mathcal{M}$, the key space by $\mathcal{K}$, the nonce space by $\mathcal{R}$ and the ciphertext space by $\mathcal{C}$.

Homomorphic encryption. Let G be the key generation algorithm, E the encryption algorithm and D the decryption algorithm. We say that a public-key cryptosystem $\Pi = (G, E, D)$ is *doubly homomorphic* if the sets $\mathcal{M}$ and $\mathcal{R}$ are (additive) Abelian groups, and $E_K(m_1; r_1) \cdot E_K(m_2; r_2) = E_K(m_1 + m_2; r_1 + r_2)$ for every $(K, m_1, m_2, r_1, r_2) \in \mathcal{K} \times \mathcal{M}^2 \times \mathcal{R}^2$. If Π is doubly homomorphic then $E_K(em; er) = E_K(m; r)^e$ for all e, and $E_K(m; r) = E_K(0; r) \cdot E_K(m; 0)$. In most of the known (doubly) homomorphic public-key cryptosystems, all spaces $\mathcal{M}$, $\mathcal{R}$ and $\mathcal{C}$ are key-dependent: In such cases we assume that the corresponding key K is understood from the context. With this in mind, we denote $M := \lceil \log_2 |\mathcal{M}| \rceil$, $R := \lceil \log_2 |\mathcal{R}| \rceil$, $C := \lceil \log_2 |\mathcal{C}| \rceil$.

Damgård-Jurik cryptosystem [DJ01]. Damgård-Jurik cryptosystem is an extension of the Paillier cryptosystem with the main difference that the size of message space can be increased polynomially without increasing $|\mathcal{R}|$ at the same time. Here, $K = n$ is an RSA modulus and s is a public parameter. The message space $\mathcal{M} = \mathbb{Z}_{n^s}$, coin space $\mathcal{R} = \mathbb{Z}^*_{n^{s+1}}$ and ciphertext space $\mathcal{C} = \mathbb{Z}^*_{n^{s+1}}$ vary together with the key n. In one variant of this cryptosystem, a message m is encrypted by generating a random number r and letting $E_K(m; r) := (1 + n)^m \cdot r^{n^s} \mod n^{s+1}$.

Coin-extractability. We say that (G, E, D, R) is a *coin-extractable public-key cryptosystem* if (G, E, D) is a public-key cryptosystem and R is an efficient algorithm, such that $R_K(E_K(m; r)) = r$ for all m and r. The Damgård-Jurik cryptosystem is coin-extractable since after decrypting $c = E_K(m; r)$, receiver obtains $r^{n^s} \mod n^{s+1}$. Since he knows factorization of n, he can then easily find r. Coin-extrability of the Paillier cryptosystem was also used in [BS01]. Note that we let $\mathcal{R}$ to be an additive group even if in the Damgård-Jurik cryptosystem, $\mathcal{R} = \mathbb{Z}^*_{n^{s+1}}$ is a multiplicative group.

Proofs of knowledge. For some (unknown) bit-string α and predicate $P(\cdot)$, $\mathrm{PK}(y = P(\alpha))$ is a (usually, honest-verifier zero-knowledge) proof-of-knowledge between two parties that given the publicly known value y, the first party knows a value of α, such that the predicate $P(\alpha)$ is true. The convention is that Greek letters denote the knowledge proved, whereas all other parameters are known to the verifier. We assume that the "Greek variables" are scoped within one proof-of-knowledge. For example, $\mathrm{PK}(c = E_K(m; \rho))$ is a proof that given as common input a ciphertext c, plaintext m and a public key K, the prover knows a nonce ρ, such that $c = E_K(m; \rho)$. For many predicates, the corresponding proofs are already known; for a few predicates we will devise new proofs in Section 4.

In our subsequent protocols, we will need two proofs-of-knowledge from [DJ01]. The first proof system is for $\mathrm{PK}(c = E_K(m_1; \rho) \vee c = E_K(m_2; \rho))$; we call this a 1-*out-of*-2 *proof system*. A noninteractive version of this proof is $3t + 2R$ bits long. The second proof system is for $\mathrm{PK}(c_1 = E_K(\mu_1; \rho_1) \wedge c_2 = E_K(\mu_2; \rho_2) \wedge c_3 = E_K(\mu_3; \rho_3) \wedge \mu_1\mu_2 = \mu_3)$; we call this a *proof system for multiplicative relationship*. A noninteractive version of this proof is $2t + M + 2R$ bits long.

Range proof. In rest of this paper we will assume that $\mathcal{M}$ is a linearly ordered ring so that we could apply the recent range proof by Lipmaa [Lip01] for $\mathrm{PK}(c = E_K(\mu; \rho) \wedge \mu \in [\mathsf{L}, \mathsf{H}])$. We will briefly outline this proof system. Prover and Verifier use both a doubly homomorphic public-key cryptosystem and an integer commitment scheme [DF01]. Now, $\mu \in [\mathsf{L}, \mathsf{H}]$ can be proven by first showing that $\mu - \mathsf{L} \geq 0$ and then showing that $\mathsf{H} - \mu \geq 0$. Thus, it suffices to describe a proof system for $\mathrm{PK}(c = E_K(\mu; \rho) \wedge (\mu \geq 0))$ that proceeds as follows: (1) Prover commits to μ and proves in statistical zero-knowledge that the committed number is equal to $\mu \pmod{|\mathcal{M}|}$. (2) Prover finds a representation $\mu = \mu_1^2 + \mu_2^2 + \mu_3^2 + \mu_4^2$ of μ. (Such representation exists iff $\mu \geq 0$ as shown by Lagrange. An efficient algorithm for finding μ_i was proposed by Rabin and Shallit [RS86].) Prover commits to $(\mu_1, \mu_2, \mu_3, \mu_4)$ and then proves in statistical zero-knowledge that $\sum_{i=1}^{4} \mu_i^2 = \mu$. With suitable security parameters, a noninteractive version of this proof is $\approx 3366 + \frac{5}{8} \lceil \log_2 \mathsf{H} \rceil$ bytes long. For more details see [Lip01].

4 Auxiliary Proofs

Range Proof in Exponents. In the following we will need a proof system for $\mathrm{PK}(c = E_K(B^\mu; \rho) \wedge (\mu \in [0, \mathsf{H}]))$; we call such a proof system a *range proof in exponents*. Our proof system is based on the observation that $\mu \in [0, \mathsf{H}]$ iff $\mu = \sum_{j=0}^{\lfloor \log_2 \mathsf{H} \rfloor} \mathsf{H}_j \cdot \mu_j$

for some $\mu_j \in \{0,1\}$ and $\mathsf{H}_j := \lfloor (\mathsf{H}+2^j)/2^{j+1} \rfloor$. For example, $\mu \in [0,10]$ iff $\mu = 5\mu_0 + 3\mu_1 + \mu_2 + \mu_3$, whereas $\mu \in [0,9]$ iff $\mu = 5\mu_0 + 2\mu_1 + \mu_2 + \mu_3$. Equivalently, $\mu \in [0,\mathsf{H}]$ iff $B^\mu = \prod_{j=0}^{\lfloor \log_2 \mathsf{H} \rfloor} (B^{\mathsf{H}_j})^{\mu_j}$ for some $\mu_j \in \{0,1\}$. Based on this, we can prove that

Theorem 1. *Let (G,E,D) be the Damgård-Jurik cryptosystem; we follow our notational convention that $\mathcal{R}$ is an additive group. For $j \in [0, \lfloor \log_2 \mathsf{H} \rfloor]$, let $\mathsf{H}_j := \lfloor (\mathsf{H}+2^{j-1})/2^j \rfloor$. Then the next protocol is a complete, HVZK and specially sound proof system for $\mathrm{PK}(c = E_K(B^\mu;\rho) \wedge \mu \in [0,\mathsf{H}])$. Let $c_{2,-1} \leftarrow E_K(1;0)$. For all $j \in [0, \lfloor \log_2 \mathsf{H} \rfloor]$ do:*

- *Both Verifier and Prover precompute H_j. Prover generates an r_j s.t. $\sum_{j=0}^{\lfloor \log_2 \mathsf{H} \rfloor} r_j = \rho$, and a $c_{1j} \leftarrow E_K((B^{\mathsf{H}_j})^{\mu_j}; r_j)$. Prover sends c_{1j} to Verifier. Both parties compute $c_{2j} = \prod_{k=0}^{j} c_{1k}$. Prover proves to Verifier that c_{1j} is an encryption of either 1 or B^{H_j} by using a 1-out-of-2 proof system from [DJ01] and that $\mathrm{PK}(c_{2,j-1} = E_K(\mu_1;\rho_1) \wedge c_{1j} = E_K(\mu_2;\rho_2) \wedge c_{2j} = E_K(\mu_3;\rho_3) \wedge \mu_1\mu_2 = \mu_3)$, by using a proof system for multiplicative relationship from [DJ01].*

Moreover, $\mathrm{PK}(c = E_K(B^\mu;\rho) \wedge \mu \in [\mathsf{L},\mathsf{H}])$ can be proven similarily by taking $\mathsf{H}_j = \lfloor (\mathsf{H}-\mathsf{L}+2^j)/2^{j+1} \rfloor$ and adding an extra addend $\mathsf{H}_{-1} = \mathsf{L}$. A proof for $\mathrm{PK}(c = E_K(B^\mu;\rho) \wedge \mu \geq \mathsf{L})$ can now be derived by letting $\mathsf{H} = |\mathcal{M}| - 1$.

Note that if $c = E_K(B^\mu;\rho)$ for $\mu \in [0,\mathsf{H}]$ then $(B^{\mathsf{H}_j})^{\mu_j} \in \{1, B^{\mathsf{H}_j}\}$. Completeness of this proof system follows since $c_{2j} \leftarrow E_K(\sum_{k=0}^{j} (B^{\mathsf{H}_k})^{\mu_k}; \sum_{k=1}^{j} r_k)$ (hence $c_{2,\lfloor \log_2 \mathsf{H} \rfloor} = c$), $D_K(c_{2j})/D_K(c_{2,j-1}) = D_K((B^{\mathsf{H}_j})^{\mu_j}) = D_K(c_{1j})$, and both protocols from [DJ01] are complete.

Let $\ell_1 = 3t + 2R$ be the length of the 1-out-of-2 proof system from [DJ01] and let $\ell_2 = 2t + M + 2R$ be the length of the proof system for multiplicative relationship from Lemma [DJ01]. Clearly, a noninteractive version of the protocol from Theorem 1 is then $(\lfloor \log_2 \mathsf{H} \rfloor + 1) \cdot (C + \ell_1 + \ell_2) = (\lfloor \log_2 \mathsf{H} \rfloor + 1) \cdot (C + 5t + 4R + M) \leq \log_2 V \cdot (C + 5t + 4R + M)$ bits long.

Damgård and Jurik presented a very similar range proof in exponents in [DJ01]. However, their proof system had a subtle flaw of working only when H is $2^j - 1$ for some j. The sole difference between our proof system and the one in [DJ01, Section 5] is in the choice of the values H_j: Namely, Damgård and Jurik chose H_j to be the j-th bit in the binary expansion of H, while we have chosen different H_j, so that values $\mu \in [0,\mathsf{H}]$ have at least one (but possibly several different) representations $\sum \mathsf{H}_j \mu_j$, but values from outside of this interval do not have such representations. Our protocol has identical complexity to the protocol from [DJ01] since the values H_j can be precomputed by both parties separately. We were recently acknowledged [Dam01] that the authors of [DJ01] were aware of the flaw in [DJ01] and have a different solution to it. However, their new protocol, described in the upcoming journal version [DJN], requires in particular approximately $2\log_2 \mathsf{H}$ additional proof systems for multiplicative relationship. This means that compared to their solution we save a constant factor in the size of interactive protocol.

Protocol 1 Proof that $(X_2, \mathsf{tiebreak})$ is correctly computed.

1. P finds $x \leftarrow D_P(c)$ and $r \leftarrow R_P(c)$. He decodes x in base B as $\sum_j x_j B^j$. Based on this, P finds $X_1 \leftarrow \max\{j : x_j > 0\}$ and X_2.
2. If $\mathsf{tiebreak} = 0$ (there is no tie-break), then do:
   (a) P proceeds as follows. Let $r_1 \leftarrow_R \mathcal{R}$, $c_1 \leftarrow E_P(B^{X_1}; r_1)$ and $c_2 \leftarrow E_P(x - B^{X_2} - B^{X_1}; r - r_1)$. Send (c_1, c_2) to V.
   (b) V verifies that $c_1 \cdot E_P(B^{X_2}; 0) \cdot c_2 = c$.
   (c) After that, P proves to V, that
      i. c_1 encrypts a $(> X_2)$th power of B: $\mathrm{PK}(c_1 = E_P(B^\mu; \rho) \wedge \mu \in [X_2 + 1, V])$.
      ii. $\mathrm{PK}(c_2 = E_P(\mu; \rho) \wedge \mu \in \left[0, (B-2) \cdot B^{X_2 - 2} - 1\right])$ by using the range proof.
3. Otherwise (if there is a tie-break), do:
   (a) P sends $c_2 \leftarrow E_P(x - 2(B^{X_2}); r)$ to V.
   (b) V verifies that $(E_P(B^{X_2}; 0))^2 \cdot c_2 = c$.
   (c) P proves to V that c_2 encrypts a value less than $(B-2) \cdot B^{X_2}$: $\mathrm{PK}(c_2 = E_P(\mu; \rho) \wedge \mu \in \left[0, (B-2) \cdot B^{X_2} - 1\right])$.

Proof That X_2 is Second Largest in Set. Let (G, E, D, R) be a coin-extractable doubly homomorphic public-key cryptosystem like the Damgård-Jurik cryptosystem. In Protocol 1, Prover P and Verifier V have a common input $(X_2, \mathsf{tiebreak}, c)$, where $D_P(c) = \sum_j x_j B^{b_j}$ for some $x_j \in [0, B-1]$. Prover has to prove to Verifier that (1) If $\mathsf{tiebreak} = 0$, then there is exactly one j_0, such that $b_{j_0} > X_2$, and exactly one j_1, such that $b_{j_1} = X_2$, and (2) If $\mathsf{tiebreak} = 1$, then there are no such j-s, for which $b_j > X_2$, but there exist $j_0 \neq j_1$, such that $b_{j_0} = b_{j_1} = X_2$. Let ℓ_1 be the length of the used range-proof-in-exponents, and ℓ_2 be the length of the used range proof. Then a noninteractive version of Protocol 1 is $\leq 2C + \ell_1 + \ell_2$ bits long.

5 New Auction Schemes

We now present our auction schemes. In the following schemes, all parties are assumed to have a public encryption key and a signature key that are in public knowledge. The signature scheme should be secure against the chosen-message attack. In the scheme of Section 5.2, A also has a public integer commitment key. We assume that the key-distribution mechanism is secure.

5.1 Simple Scheme

Protocol 2 depicts a simple auction scheme that puts more trust on A, compared to the later scheme from Section 5.2, but avoids elaborated cryptographic protocols and does not put as severe limites on the values B and V as the latter. If $\mathsf{tiebreak} = 0$ (no tie-break), a successful protest constitutes bidder i_w proving (in zero-knowledge) that he did not bid more than X_2, or some other bidder proving that he bid also more than X_2. If $\mathsf{tiebreak} = 1$, a successful protest means i_w proving that he bid less than X_2, or some other bidder proving that he bid more than X_2. All such proofs can be based on the previously described range proofs that originate from [Lip01].

Protocol 2 The simple auction scheme.

BIDDING PHASE

1. Bidder i encrypts b_i by using A's public key and sends the resulting ciphertext c_i together with $\mathrm{sig}_i(c_i)$ to S via a confidential channel.
2. S verifies the signatures (complains, if necessary). He computes $z \leftarrow \mathrm{sig}_S(\{c_i\})$, where $\{c_i\}$ is represented in some fixed order that does not depend on i-s. (For example, in lexicographic order with respect to the c_i-s.). He broadcasts $(\{c_i, \mathrm{sig}_S(c_i)\}, z)$ to all bidders.
3. Every bidder i obtains $(\{c_{i'}, \mathrm{sig}_S(c_{i'})\}, z)$, and complains if c_i is missing. He also verifies the signature z.

BID OPENING PHASE

1. A does the following. Obtain $\{(c_i, \mathrm{sig}_S(c_i))\}$ and z, and verify the signatures $\mathrm{sig}_S(c_i)$, $\forall i$, and z. For all i: Decrypt c_i and obtain b_i. Compute the second highest bid X_2. Set $\mathsf{tiebreak} \leftarrow 1$ if there is a tie-break and $\mathsf{tiebreak} \leftarrow 0$, otherwise. Send $(X_2, \mathsf{tiebreak})$, together with signature $z' \leftarrow \mathrm{sig}_A(X_2, \mathsf{tiebreak}, \{c_i\})$ to S.
2. S verifies the signature z'. He then broadcasts $(X_2, \mathsf{tiebreak}, z, z')$ to all bidders.
3. After obtaining $(X_2, \mathsf{tiebreak}, z, z')$, all bidders verify the signature z', in particular that it is given over the same set $\{c_i\}$ as z.
4. A points to S a c_{i_w}, such that $D_A(c_{i_w}) = X_1$. S identifies i_w and declares him the winner.
5. Bidders can now protest against the choice of i_w.

In this auction scheme, A will get to know the winner and the bid statistics, but cannot bind bids with concrete bidders. A malicious A can change X_2 to X_2', $X_1 > X_2' \geq X_2$. The seller S will get to know only the minimal amount of information: That is, X_2, and the winner (or all winners if there is a tie-break). If S and A do not collude, the seller cannot deviate from the protocol without being detected.

5.2 Homomorphic Scheme

Protocol 3 depicts *the homomorphic scheme*, where every bid b_i is encoded as B^{b_i}. This encoding will allow everybody to compute, given encryptions of B^{b_i}, an encryption of $\sum_i B^{b_i}$ without knowing the corresponding decryption key. Note that for this scheme to work correctly it is necessary that $B^V < |\mathcal{M}|$. We assume implicitly that communication goes over a confidential channel.

The auction authority will get to know the bid statistics but cannot bind them with the bidders. The seller will get to know only the minimal amount of information, X_2 and the winner (or all winners if there is a tie-break). If S and A do not collude, neither S nor A can deviate from the protocol without being detected.

In many situations, knowing bid statistics might not be very valuable for A: First, even if the authority does sell the statistics to the seller of a subsequent auction, the new seller will most probably not have exactly the same set of bidders. Second, if S would use designated verifier signatures (with verifier A), A would be unable to convince the new seller that he is actually selling correct data, unless carefully building up reputation of a "honest cheater". However, such a reputation is unlikely to stay hidden from bidders

Protocol 3 The homomorphic auction scheme.

BIDDING PHASE

1. Bidder i encodes and encrypts his bid b_i by using A's public key, $c_i = E_A(B^{b_i}; r_i)$, signs it, and sends $(c_i, \mathrm{sig}_i(c_i))$ to S. Bidder i proves to S that the bid is correctly computed by performing a proof for $\mathrm{PK}(c_i = E_A(B^\mu; \rho) \wedge (\mu < V + 1))$.
2. S does the following: Verify the signatures and complain if necessary. Let $c \leftarrow \prod_i c_i = E_A(\sum_i B^{b_i}; \sum_i r_i) = E_A(\sum_j x_j B^j; \sum_i r_i)$, $\mathcal{C} \leftarrow \{c_{\pi(i)}\}$ for random permutation π, $h \leftarrow H(\mathcal{C})$ and $z = \mathrm{sig}_S(h, c)$. Send $\mathcal{C}$ to all bidders. Post (c, z).
3. For all i, bidder i verifies that $c_i \in \mathcal{C}$, $c = \prod_{c' \in \mathcal{C}} c'$ and $z = \mathrm{sig}_S(h, \mathcal{C})$.

BID OPENING PHASE

1. A obtains (c, z) and verifies the signature z.
2. After that, A decrypts c, obtains $D_K(c) = \sum_j x_j B^j$ and then computes the second highest bid X_2 and a bit $\mathsf{tiebreak}$, such that $\mathsf{tiebreak} = 1$ iff there is a tie-break. He sends $(X_2, \mathsf{tiebreak})$, together with his signature $z = \mathrm{sig}_A(X_2, \mathsf{tiebreak})$, to S.
3. S verifies z.
4. A proves to S that the pair $(X_2, \mathsf{tiebreak})$ is correctly computed. (See Section 4 for the corresponding proof.)
5. S publishes X_2 on an authenticated medium together with A's and his own signatures.
6. Bidders can now participate in the confirmation phase with S. If $\mathsf{tiebreak} = 0$, the bidder who confirms that he bid more than X_2 will be the winner. If $\mathsf{tiebreak} = 1$, a previously announced rule (for example, the equal probability rule) is used to determine the winner.

for extended periods of time. Third, even if it is impossible to verify for sure whether A abuses the bid statistics, but too obvious abuses will certainly be noticed and ruin his reputation.

The confirmation step is optional, since the proof of step 4 already shows that X_2 is correctly computed. The highest bidder has to participate in the confirmation phase to claim the item. However, if he does not, one can apply a mandatory protocol where every bidder has either to confirm or revoke that he is eligible to win.

6 Refinements to Our Auction Schemes

Using a prime B. If B is a prime, the range proofs in exponents can be made considerably shorter as shown in [Lip01]. Without going into more details, we note that a noninteractive version of this proof has length $2636 + \lceil \log_2 |\mathcal{M}| \rceil + \frac{5}{16} \log_2 \mathsf{H}$ bytes. Now, restricting B to be a prime is not a big obstacle in our auction scheme. Really, by the prime number theorem, the average gap $p_{i+1} - p_i$ between two consequent primes less than n is $\Theta(\log_2 n)$. In particular, the largest prime gap between primes less than 1200 is 22. Thus, the seller has to introduce approximately $\Theta(\log_2 B)$ dummy bidders that do not actually participate in the auction.

Extension to $(m + 1)$st price auctions. Vickrey auction mechanism can be generalized to the $(m + 1)$st price auction mechanism, where m copies of the same item are given to m highest bidders for the $(m+1)$st highest bid. The $(m+1)$st price auction scheme is a direct incentive-compatible mechanism [Vic61]. A trivial modification to our homomorphic auction scheme results in a $(m+1)$st price auction scheme with additional communication of about $(m-2) \cdot (C+\ell)$ bits, where $\ell \leq 2(C+5t+4R+M)\log_2 V$ is the length of the range proof in exponents from Section 4. Briefly, this modification consists of changing Protocol 1 so that instead of proving that $D_K(c) = B^{\chi_1} + B^{X_2} + \tau$ with $\chi_1 > X_2$ and $\tau < B^{X_2+1}$, P proves that $D_K(c) = \sum_{i=1}^{m} B^{\chi_i} + B^{X_2} + \tau$ with $\chi_i > X_2$ and $\tau < B^{X_2+1}$. On the other hand, if we assume that B is a prime then usually $\ell \leq 3$ KB. The only previous secure $(m+1)$st-price auction schemes that we are aware of by Kikuchi [Kik01] and by Naor, Pinkas and Sumner [NPS99]. In the latter scheme, circuit for the mth price auctions is about $m \log_2 V$ times bigger than circuit for the first price auctions [Pin01].

Thresholding. It is possible to distribute A and/or S using the threshold trust model. For example, when trust on A is distributed in a way where bid statistics will only be leaked if at least $1/3$rd of the A-servers are faulty then in the homomorphic scheme the thresholded A does not have to prove that X_2 was correctly computed. This introduces a new interesting *bipartite threshold trust model*, where some of the functionality is controlled by one set $\mathcal{S}$ of servers (operated by one or more parties), while some other functionality is controlled by another set $\mathcal{A}$ of servers (operated by one or more parties, independent from the parties who operate the set $\mathcal{S}$). Server sets $\mathcal{S}$ and $\mathcal{A}$ check that the other set behaves correctly. Our auction schemes stay secure unless significant fractions of both $\mathcal{S}$ and $\mathcal{A}$ cheat. We feel that this bipartite threshold trust model might also be interesting in many other applications, like e-voting.

Reducing the influence of collusions. Let H be a secure commitment scheme; in practice one may assume that H is a hash function. Damage caused by colluding A and S can be reduced in both of our auction schemes when the bidders first send a signed commitment to their bid to S, who then broadcasts all commitments together with his signature on the tuple of commitments. Only after that, actual encrypted bids are sent to S.

When this "meta-scheme" is employed, an auction will stay correct even when S and A collude. The only use from the colluding is that then A and S will obtain additional information: Namely, they will be able to connect every bidder with his bid. However, they will *not* be able to artificially raise X_2 or declare a false winner. The same simple but very useful method works in conjunction with almost every auction (and voting) scheme.

There are three scenarios how this meta-scheme itself could be abused. First, a bidder can cheat by sending a commitment but then refusing to send the bid. However, since the commitments are signed the offending bidder is identifiable in some sense. Second, colluding S and A can delete some bids that are not to their liking, arguing that the corresponding encrypted bid was not submitted. However, in this case corresponding bidders can prove by showing their bids that S (and A) were faulty. Third, S and A can arrange a shill to submit a very low fake bid. If then the results are not to their liking, they can claim that the shill failed to send the encryption. However, this shill is again identifiable. Hence, this concern might not be very serious especially in the local electronic auctions. (A more general solution would be to use a fair exchange instead of receipts during the bid commitment.)

Avoiding replay attacks. Because of the reliance on homomorphism, encrypted bids cannot contain any other information but B^{b_i}. This opens up cut-and-paste attacks that may compromise bid privacy. As a simple example, a crook seller that wants to find the highest bid in an auction could replay the winning bid, along with a bunch of zero bids and one artificially high bid $b = V - 1$ to A.

Replay attacks can be avoided by once again using the coin-extractability property of the cryptosystem, as far as the random coins r_i are not revealed to the seller. Namely, accompany each bid with $E_A(\mathsf{transaction\_id}; r_i)$, where r_i is the same coin that was used to encrypt the bid, and transaction_id is guaranteed to be unique (i.e, something that A can detect in case of a replay).

Preferably transaction_id should also contain a commitment of auction parameters. For example, transaction_id can be computed as $H(\mathsf{auction\_advertisement})$ where auction_advertisement includes all relevant details about the auction (e.g., seller name, sequence number added by seller, auction mechanism, deadlines, etc.). Since this is the only communication channel (for arbitrary data) from bidders to the auction authority, it should be used to send all security-critical information. For example, in the absence of such a communication between bidders and A, S could advertise a Vickrey auction to the bidders, but tell A that it is a first-price auction.

Avoiding replay attacks in e-voting schemes. Similar cut-and-paste replay attacks can be applied to the voting schemes that base on homomorphic cryptosystems. Here one can use exactly the same solution as in the previous paragraphs. Even if replaying is hard to mount to voting systems, our proposed defence mechanisms are so simple that one might consider using them.

7 Discussion

Local electronic auctions. In *local electronic auctions*, the bidders are physically present at an auction house, and participate via a local wireless network by using some mobile devices for computations. Local online auctions have some specific positive properties that simplify their organization and decrease the trust requirements. First, due to the locality assumption, the bidders can closely examine the goods before they decide to bid. Similarly, the winning bidder is physically present and payment can be enforced as in traditional auctions. Hence, the two most common source of complaints about Internet auctions are avoided. Second, we can assume high bandwidth capacity and sufficiently reliable communications between the seller and the bidders. In particular, the audience is captive: Bidders will stay available. Therefore, a multiple-round auction is not a problem.

Our auction schemes were designed with local electronic auctions in mind though not solely for them. Partially due to these remarks, we have assumed that law enforcement is out of the scope of the current paper: It is certainly easy to enforce correct behaviour in local auctions, but in remote auctions one must use additional protocols

$\lfloor \log_2 \lvert\mathcal{M}\rvert \rfloor, V$	100	200	300	400	500	1000
1024	1209	34	10	5	4	2
1536	$4.2 \cdot 10^4$	205	34	14	8	2
2048	$1.4 \cdot 10^6$	1209	110	34	17	4
3072	$1.7 \cdot 10^9$	$4.2 \cdot 10^4$	1209	205	70	8
4096	∞	$1.4 \cdot 10^6$	12884	1209	292	17

Table 1. Example values of maximum possible B and V for some common cardinalities of message spaces. In general, a greater $\lvert\mathcal{M}\rvert$ means that either higher security parameter has to be used, or the bid should consist of several encryptions. "∞" means that the number of possible bidders far exceeds the population of Earth, and is hence virtually unlimited.

that are not described in this paper, e.g., punish bidders who refuse to co-operate in the confirmation phase.

Moreover, in local electronic auctions the bidder and the seller are in the same room or at least building, while the authority might be kilometers away. The same authority might be involved in many other auctions in parallel. This observation motivated us to prioritize the seller-authority communication complexity over the seller-bidders communication complexity.

Limitation on the number of valuations. A disadvantage of the homomorphic scheme from Section 5.2 is that the number of different valuations is small. Namely, if the plaintext message space is $\mathcal{M}$ then the maximum number of valuations V and the maximum number of bidders B are bounded by $V \cdot \log_2 B < \log_2 \lvert\mathcal{M}\rvert$. While one can increase the size of message space, doing this will greatly increase the computational complexity of the homomorphic scheme, with $\log_2 \lvert\mathcal{M}\rvert = 3072$ being almost the limit with the current computational technology. Still, it means that for smaller V, the number of bidders is almost unlimited, as seen from Table 1.

We feel that choice $V \leq 500$ is sufficient in most of the auctions. For example, in an auction of a second-hand item, the bid $b \in [0, V]$ could correspond to the price $\frac{5}{V}Pb$, where P is the original price of the sold item. A price increase of 1% of P seems to be sufficiently precise. Moreover, the mapping between the bids and actual prices does not have to be linear, it only has to be strictly monotonic and publicly known. In particular, one might use higher precision with large bids than with small bids. The fact that in our scheme there might be more bidders than available bid options should not be a big concern. (Similar encoding was used in [DJ01] in the context of electronic voting where V—the number of candidates—is usually much less than 500.)

8 Comparison to Naor-Pinkas-Sumner Scheme

We will next compare the homomorphic scheme from Section 5.2 to the Naor-Pinkas-Sumner scheme [NPS99], the only previous cryptographic Vickrey auction scheme that does not rely on the threshold trust.

In our scheme, A receives more information than in [NPS99]. On the other hand, detecting misbehavior by the auction authority A is considerably more complicated in [NPS99]. Basically, to catch a cheating authority with probability $1 - 2^{-m}$, the off-line complexity in their scheme will increase m times, compared to the basic scheme. In the homomorphic scheme, the actions of A are verifiable; verifiability can be omitted but this would decrease the interaction only by a half.

First, the off-line communication complexity of Naor-Pinkas-Sumner scheme (*without* applying the cut-and-choose method) is $300B \cdot \log_2 V$ bytes, the on-line communication complexity between the servers is about $Bt \cdot \log_2 V$ and the communication complexity of each bidder is $\Theta(t \cdot \log_2 V)$. This makes the total communication overhead of the auctioneer to be $\Theta(Bt \cdot \log_2 V)$.

In the homomorphic scheme, without the confirmation phase, bidders' sole communication with S consists of one encryption and a proof of bid correctness that takes together $(C + 5t + 4R + M)O(\log_2 V)$ bits, and requires $\Theta(\log_2 V)$ encryptions. If B is a prime then a constant number of encryptions and communication of $\Theta(V \cdot \log_2 B)$ bits suffices. Confirmation phase has the same complexity. Therefore, the total communication between the auctioneer and the bidders is $\Theta(BV \cdot \log_2 B)$.

As motivated in Section 7, our primary concern as regards communication complexity is the communication between S and A that is dominated by the noninteractive proof that $(X_2, \mathsf{tiebreak})$ was correctly calculated. When B is a prime then the asymptotic seller-authority communication complexity will be $\Theta(V \cdot \log_2 B)$ that is close to optimal for large-scale auctions. We will next give a comparison of the seller-authority communication complexity for some concrete values of B and V. We will use the Damgård-Jurik cryptosystem, where $R \approx C \approx \frac{s+1}{s}M$. We will suppose that $M = 1024$ and choose s as $s \leftarrow \lceil V \cdot \log_2 B/M \rceil$. We will also assume that $t = 80$ and that $V \in \{300, 500\}$. For these special cases, efficiency comparison of the homomorphic scheme from Section 5.2 with the Naor-Pinkas-Sumner scheme [NPS99] is presented in Table 2. (For the homomorphic scheme, this table shows the combined total length of a range proof and a range proof in exponents, since both are used in the correctness proof. This table also presents two versions of the homomorphic scheme, one that works with a generic B and another one which works only with a prime B.)

As seen from Table 2, if $B \geq 50$ then the seller-authority interaction in the prime-B homomorphic scheme is less than $30V \cdot \log_2 B$ bytes. The prime-B homomorphic scheme and the Naor-Pinkas scheme have roughly equal communication complexity for small-scale auctions. On the other hand, the prime-B version of the homomorphic scheme is at least 10 (resp., at least 100) times more communication-efficient than the Naor-Pinkas-Sumner scheme in medium-scale auctions (resp., in large-scale auctions). The difference in interaction will be even greater when the cut-and-choose method is applied to the Naor-Pinkas-Sumner scheme.

Finally, note that in our scheme, the computations of the seller consist of a few exponentiations per every bidder and a few exponentiations when communicating with the authority. Every bidder has to do a few exponentiations. On the other hand, the computational complexity of the authority is somewhat higher due to the use of a range proof; however range proofs from [Lip01] seem to provide an adequate efficiency.

V			3	4	8	16	32	64	128	256	512	1024
300	Our	gener. B	24.6	24.6	24.6	38.1	38.1	38.1	51.6	51.6	51.6	51.6
		prime B	5.7	5.7	5.7	8.4	8.4	8.4	11.0	11.0	11.0	11.0
			$s = 1$			$s = 2$			$s = 3$			
	[NPS99]		7.2	9.6	19.3	38.6	77.1	154.3	308.6	617.2	1234.3	2468.6
500	Our	gener. B	26.5	26.5	41.0	41.0	55.6	55.6	70.1	70.1	84.6	84.6
		prime B	5.7	5.7	8.4	8.4	11.0	11.0	13.6	13.6	16.2	16.2
			$s = 1$		$s = 2$		$s = 3$		$s = 4$		$s = 5$	
	[NPS99]		7.9	10.5	21.0	42.0	84.1	168.1	336.2	672.4	1344.9	2679.7

Table 2. Communication efficiency of the homomorphic scheme (for both generic B and a prime B) and the Naor-Pinkas-Sumner scheme [NPS99] for $V \in \{300, 500\}$ and varying B. The proof lengths are given in kilobytes. In the case of our scheme we also mention the size of s.

9 Conclusions

We proposed two different auction schemes that work in a setting without threshold trust and are practical for a large number of bidders. In both schemes we have an seller S and an auction authority A that are assumed not to collude. In the second, homomorphic, scheme we embed many bids in one encryption. This allows the communication complexity to be reduced substantially.

The homomorphic scheme achieves, especially compared to [NPS99], (1) Similar level of security for other parties w.r.t. seller or bidders; (2) Verifiability of A during the protocol execution: A can change the outcome of auctions only when colluding with S; (4) Both bidder-seller and seller-authority communication complexities are reduced to $\Theta(V \cdot \log_2 B)$ bits with a moderate-size hidden constant in the Θ-expression. This makes it possible to use our auction scheme in large-scale auctions.

On the other hand, the main drawbacks of our scheme are (1) Somewhat lower level of confidentiality for other parties w.r.t. auction authority; and (2) Limited number of possible bids. However, as argued before, both drawbacks might not be that serious. In particular, we feel that scalability in the number of possible bidders is more important than in the number of possible bids. Finally, note that the homomorphic scheme from Section 5.2 can be used as a backbone for voting scheme, modulo the change that A, instead of sending back $(X_2, \mathsf{tiebreak})$ and proving its correctness, just sends back $x = D_A(c)$ together with a proof of correct decryption.

Acknowledgments and Further Work

The first author was partially supported by Nokia Research. We would like to thank Ivan Damgård for pointing out their new range proof in exponents [DJN] and for comments on Section 4, and Benny Pinkas for useful discussions and comments. We were recently reported by the author of [Lip01] that there exists a really efficient range proof in exponents also if one does not assume that B is a prime.

From the efficiency point of view, our auction scheme is an interesting complement to the Naor-Pinkas-Sumner auction scheme, since it has communication complexity of

$\Theta(V \cdot \log_2 B)$ as compared to $\Theta(B \cdot \log_2 V)$ (with a considerably smaller constant in the big-Θ expression). It would be interesting to find a secure Vickrey auction scheme without threshold trust where the communication complexity is polylogarithmic in both B and V. It would also be interesting to know how to avoid A getting to know the bid statistics without using threshold trust.

References

[BS01] Olivier Baudron and Jacques Stern. Non-interactive Private Auctions. In Paul Syverson, editor, *Financial Cryptography — Fifth International Conference*, Lecture Notes in Computer Science, Grand Cayman, BWI, 19–22 February 2001. Springer-Verlag. To appear.

[Cac99] Christian Cachin. Efficient Private Bidding and Auctions with an Oblivious Third Party. In *6th ACM Conference on Computer and Communications Security*, pages 120–127, Singapore, 1–4 November 1999. ACM Press.

[Dam01] Ivan Damgård. Personal communication, November 2001.

[DF01] Ivan Damgård and Eiichiro Fujisaki. An Integer Commitment Scheme Based on Groups with Hidden Order. Technical Report 064, IACR, 13 August 2001. Available at `http://eprint.iacr.org/2001/064/`.

[DJ01] Ivan Damgård and Mads Jurik. A Generalisation, a Simplification and Some Applications of Paillier's Probabilistic Public-Key System. In Kwangjo Kim, editor, *Public Key Cryptography '2001*, volume 1992 of *Lecture Notes in Computer Science*, pages 119–136, Cheju Island, Korea, 13–15 February 2001. Springer-Verlag.

[DJN] Ivan B. Damgård, Mads J. Jurik, and Jesper B. Nielsen. A Generalization, a Simplification and Some Applications of Paillier's Probabilistic Public-key System. *International Journal of Information Security*. To appear. Manuscript, available from authors.

[Kik01] Hiroaki Kikuchi. $(M+1)$st-Price Auction Protocol. In Paul Syverson, editor, *Financial Cryptography — Fifth International Conference*, Lecture Notes in Computer Science, Grand Cayman, BWI, 19–22 February 2001. Springer-Verlag. To appear.

[Lip01] Helger Lipmaa. Statistical zero-knowledge proofs from diophantine equations. Cryptology ePrint Archive, Report 2001/086, 20 November 2001. http://eprint.iacr.org/.

[NPS99] Moni Naor, Benny Pinkas, and Reuben Sumner. Privacy Preserving Auctions and Mechanism Design. In *The 1st ACM Conference on Electronic Commerce*, Denver, Colorado, November 1999.

[NS93] Hannu Nurmi and Arto Salomaa. Cryptographic Protocols for Vickrey Auctions. *Group Decision and Negotiation*, 2:363–373, 1993.

[Pin01] Benny Pinkas. Personal communication, October 2001.

[RH95] Michael H. Rothkopf and Ronald M. Harstad. Two Models of Bid-Taker Cheating in Vickrey Auctions. *Journal of Business*, 68(2):257–267, April 1995.

[RS86] Michael O. Rabin and Jeffrey O. Shallit. Randomized Algorithms in Number Theory. *Communications in Pure and Applied Mathematics*, 39:239–256, 1986.

[RTK90] Michael H. Rothkopf, Thomas J. Teisberg, and Edward P. Kahn. Why are Vickrey Auctions Rare? *The Journal of Political Economy*, 98(1):94–109, February 1990.

[SM00] Kouichi Sakurai and Shingo Miyazaki. An Anonymous Electronic Bidding Protocol Based on a New Convertible Group Signature Scheme. In Ed Dawson, Andrew Clark, and Colin Boyd, editors, *Fifth Australasian Conference on Information Security and Privacy*, volume 1841 of *Lecture Notes in Computer Science*, pages 385–399, Brisbane, Australia, 10–12 July 2000. Springer-Verlag. ISBN 3-540-67742-9.

[Vic61] William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, March 1961.