

A Two-Server, Sealed-Bid Auction Protocol

Ari Juels and Michael Szydło

RSA Laboratories
Bedford, MA 01730, USA
E-mail: {ajuels,mszydło}@rsasecurity.com

Abstract. Naor, Pinkas, and Sumner introduced and implemented a sealed-bid, two-server auction system that is perhaps the most efficient and practical to date. Based on a cryptographic primitive known as oblivious transfer, their system aims to ensure privacy and correctness provided that at least one auction server behaves honestly. As observed in [19], however, the NPS system suffers from a security flaw in which one of the two servers can cheat so as to modify bids almost arbitrarily and without detection. We propose a means of repairing this flaw while preserving the attractive practical elements of the NPS protocol, including minimal round complexity for servers and minimal computation by players providing private inputs. Our proposal requires a slightly greater amount of computation and communication on the part of the two auction servers, but actually involves much *less* computation on the part of bidders. This latter feature makes our proposal particularly attractive for use with low-power devices. While the original proposal of NPS involved several dozen exponentiations for a typical auction, ours by contrast involves only several dozen modular multiplications.

The key idea in our proposal is a form of oblivious transfer that we refer to as *verifiable proxy oblivious transfer* (VPOT).

Key words: auction, sealed-bid auction, oblivious transfer, secure multi-party computation, secure function evaluation

1 Introduction

Cryptography offers a broad range of tools for distributing trust among computing entities in flexible and often unexpected ways. In an electronic auction setting, for example, a foundational cryptographic procedure known as *secure function evaluation* enables the submission and processing of sealed bids without the presence of a single, trusted auctioneer. As secure function evaluation is rather impractical in its general form, a large body of research, e.g., [1, 5, 11, 17, 19, 24], has focused on tailoring cryptographic protocols specifically to achieve efficient sealed-bid auction systems.

A recent architecture proposed and implemented by Naor, Pinkas, and Sumner [20] represents substantial progress toward the goal of practical sealed-bid

auctioning with distributed trust. In their system, bidders submit encrypted bids to a front-end server known as an *auctioneer*. With the involvement of a second, back-end server known as an *auction issuer*, any type of sealed-bid auction may be conducted, e.g., highest-bid auctions, Vickrey auctions, and so forth. The architecture distributes trust between the two servers in the following simple model: If at least one server is honest, the bids of all participants will remain private, and any auction outcome is assured to be correct. There is no robustness, however, in the sense that either server can cause the protocol to terminate. NPS report good scalability, claiming that their system can accommodate thousands of bidders with reasonable overhead. The computational requirement for bidders in their system is approximately one modular exponentiation per bit in a bid representation. See [20] for further details.

As identified in a footnote in work by Jakobsson and Juels [19], however, the NPS system has a serious flaw that permits tampering by one of the servers. Although not explicated in [19], it is easy to see that the auction issuer can modify any bit in any bid without detection. The underlying problem is a variant introduced by NPS on a cryptographic primitive known as *1-out-of-2 oblivious transfer* (1-2 OT), as we now explain.

Basic 1-2 OT is a procedure involving two players, a *Chooser* and a *Sender*. The Sender possesses a pair of values (t_0, t_1) . We refer to these values throughout our paper as *tags*. The Chooser elects to receive from the Sender one of these two tags t_b for $b \in \{0, 1\}$. The 1-2 OT procedure is oblivious in the sense that the Sender learns negligible information about b . An additional privacy property of 1-2 OT that the Chooser learns negligible information about t_{1-b} , i.e., the value that she did not select. NPS consider a variant on 1-2 OT, called *proxy oblivious transfer*. This variant involves an intervening third party known as a *Proxy*, who receives the value t_b on behalf of the Chooser, but herself learns negligible information about b and t_{1-b} . We provide details on the protocol below. Proxy oblivious transfer accomplishes the goal for which it was designed, namely privacy protection, but does not include any mechanism for *verifiability*. In particular, the proxy oblivious transfer protocol does not ensure that the Sender transmitted t_b as desired. In the NPS auction setting, the Sender (auction issuer) can substitute the tag t_{1-b} . Thus the auction issuer can tamper with bids!

In this paper, we introduce a protocol called *verifiable proxy oblivious transfer* (VPOT) that addresses the vulnerability in the NPS protocol. In principle, introducing verifiability into proxy oblivious transfer is not difficult using basic – and potentially expensive – cryptographic techniques such as zero-knowledge proofs. Our contribution in the design of VPOT is a collection of techniques that render the verification process computationally inexpensive and yet, at the same time, provably secure. When VPOT is introduced into the NPS auction protocol, it increases the computational burden on auction servers somewhat, but actually results in much *less* computation for bidders. This is particularly desirable given the fact that bidders in many settings may wish to employ low-

power, handheld devices. Thus, VPOT not only remedies the security flaw in the NPS architecture, but renders the system even more practical.

1.1 Background: Two-party computation and the NPS protocol

Secure function evaluation (also known as secure multi-party computation) began with the work of Yao [26], and Goldreich, Micali, and Wigderson [15], and has spawned an ever growing body of research papers. See [13] for a good overview of early work. The general goal of secure multi-party computation is to enable m players to apply a function F to respective private inputs X_1, X_2, \dots, X_m such that some subset of players learns $F(X_1, X_2, \dots, X_m)$, but no player learns additional, non-negligible information. Privacy and robustness against active (indeed, adaptive) adversaries are possible to achieve provided that the adversary controls at most a fraction of the players. Assuming the existence of a broadcast channel, this fraction is $1/2$; otherwise, it is $1/3$. For some recent work aiming to achieve practical multi-party protocols, see, e.g., [8, 18].

The two-player case of secure function evaluation is distinguished by its relative simplicity and practicality. The original secure function evaluation protocol of Yao [26] treated this case, and remains an important tool even now. In contrast to more general techniques, in which field operations such as addition and multiplication are the atomic unit of computation, the Yao protocol involves direct computation on boolean gates. While this is a limitation in the general case, many real-world protocols such as auctions involve intensive bitwise manipulation such that boolean circuits are in fact a natural form of representation for the required functions. The Yao protocol is appealing for other reasons as well. Provided that only one player is to learn the output, it is in fact possible to execute the Yao protocol with only one-round of interaction, an observation first set forth implicitly in [20] and explored in detail in [7]. While constant-round secure function evaluation is possible for multiple players, it requires both high overhead and the availability of a broadcast channel [3]. A model in which both players in the Yao protocol learn the output of such computation in a fair manner (given a passive trusted entity) is also possible, and is explored in [6].

Suppose that Alice and Bob wish to engage in the Yao protocol on respective private inputs X_A and X_B such that Bob learns the output $y = F(X_A, X_B)$. Alice constructs a “garbled” circuit representing F . She sends this circuit to Bob, along with a “garbled” representation of X_A . In order to evaluate the “garbled” circuit, Bob additionally needs a “garbled” version of his input X_B . He obtains this from Alice using some form of 1-2 *oblivious transfer* (OT) [21] protocol. This is the component of the Yao protocol that we focus on in detail in this paper. In the case where Alice may cheat, another important component in two-player secure function evaluation protocols are proofs of correct construction of the Yao circuits. A cut-and-choose protocol for this is proposed in [20], while [7] explores use of general non-interactive proof techniques. (If Alice wishes Bob to send y to her such that she can verify its correctness, she need merely embed a verification tag in her “garbled” version of F in the appropriate manner.)

Numerous variants of oblivious transfer have been considered in the literature [13]. Notions of combining bit commitment with oblivious transfer in a theoretical setting to achieve a committed or verifiable oblivious transfer are explored for example in [10] and [9]. These works explore theoretical approaches that treat oblivious transfer and bit commitment as black boxes, and are thus necessarily expensive. An alternative proposal makes use of a trusted initializer [22]. The key observation made by NPS in their auction system design is that by involving a Proxy in the oblivious transfer procedure, it is possible to expand application of basic Yao style two-server function evaluation so as to accept inputs from an arbitrarily large number of players, i.e., bidders, while the evaluation process is restricted to two auction servers.

Briefly stated, the NPS construction is as follows. The auction issuer (again, the back-end server) constructs a “garbled” representation of a function F that describes the auction protocol. The auctioneer (again, the front-end server) evaluates the circuit for F using “garbled” inputs representing the bids. In order to obtain the “garbled” input for a bit b in a given bid, it is necessary to invoke the proxy oblivious transfer protocol. In this protocol, the bidder plays the role of the Chooser, the auctioneer plays the role of the Proxy, and the auction issuer plays the role of the Sender. The Sender transmits “garbled” inputs (t_0, t_1) for the circuit corresponding to a ‘0’ bit and a ‘1’ bit in the bid. The Chooser selects t_b , which the Proxy receives through the transfer protocol. Having done this for all bits in all bids, the Proxy is able to evaluate F on the input bids and determine the outcome of the auction. The privacy properties of the proxy oblivious transfer protocol ensure that the Proxy does not learn b or t_{1-b} for any bit. The Proxy therefore does not learn any bidding information and can only evaluate F on correct bid amounts. Likewise, the Sender does not learn the bid amounts. Only if the Proxy and Sender collude is this privacy guarantee breached.

NPS include other security enhancements in the protocol. For example, for the auctioneer to ensure that the auction issuer has constructed F correctly, the two engage in a cut-in-choose protocol. Thus, the auctioneer in fact evaluates multiple, independent circuits representing F . We provide more details below.

1.2 Our contribution: Verifiable proxy oblivious transfer (VPOT)

The failing in the NPS protocol is that the auction issuer can transmit t_{1-b} instead of t_b without detection. To address this problem, we propose a protocol known as verifiable proxy oblivious transfer (VPOT). VPOT enables the Proxy (auctioneer) to ensure that the Sender (auction issuer) sent t_b , as required. VPOT retains all of the privacy characteristics of proxy oblivious transfer.

Here is a simplified overview of VPOT. The Sender provides commitments \mathcal{C}_0 and \mathcal{C}_1 to tags t_0 and t_1 (respectively representing a ‘0’ bit and ‘1’ bit in a bid). These commitments take the form of a randomly ordered pair $(\mathcal{C}_a, \mathcal{C}_{1-a})$, i.e., a is a randomly selected bit. The Sender also provides a commitment $E[a]$

to ordering a . Observe that the triple $(\mathcal{C}_0, \mathcal{C}_1, E[a])$ binds the Sender to values for t_0 and t_1 .

As usual in a 1-2 OT protocol, the Chooser selects a value t_b to be decommitted by the Sender. The Chooser in fact splits this bit b into two shares b_P and b_S such that $b = b_P \oplus b_S$. The Chooser sends the share b_S to the Sender. This is transmitted (via the Proxy) as a ciphertext $E[b_S]$. She sends the share b_P to the Proxy, also in a specially committed form that we do not describe here. It is the splitting of b into two shares that ensures privacy with respect to the two auction servers (provided that there is no collusion).

Finally, the Chooser transmits to the Proxy a secret value x that enables the Proxy to receive the selected tag t_b . The Sender decommits t_b for the Proxy, who then checks the correctness of the decommitment.

Here is a list of the more interesting cryptographic building blocks used in the construction of VPOT. While none is individually novel *per se*, our new constructions combine them in a novel way, providing a new fundamental building block useful for securely extending traditional two-party techniques to settings with multiple contributors.

- *Double commitment*: The Sender’s commitment $\mathcal{C}_k(t)$ on tag t in fact consists of a pair of values (Y_1, Y_2) . The first value, Y_1 , is the commitment on a key or witness k . In particular here, $Y_1 = H(k^3)$, where the cubing operation takes place over an RSA modulus provided by the Sender (as discussed in more detail below). H here is a hash function (modelled as a random oracle for security proofs on the system). Observe that as the hash of a cube, Y_1 is really a commitment within a commitment. It is for this reason that we refer to $\mathcal{C}_k(t)$ as a *double* commitment. The second value of the commitment pair, Y_2 , represents an encryption of t under k . In particular, $Y_2 = H(k) \oplus t$, where \oplus denotes the bitwise XOR operator. Knowledge of the witness k is sufficient both to open the commitment and obtain t and also to verify that the commitment has been correctly opened. This double commitment scheme may be seen to be both computationally binding and computationally hiding under the RSA assumption, with the random oracle model invoked for H .
- *RSA-based oblivious transfer*: Most oblivious transfer protocols designed for practical use in two-party secure function evaluation, e.g., in [20, 2], employ El Gamal-based encryption of tags [14]. The result is that the Chooser must perform at least one exponentiation per 1-2 OT invocation. In contrast, we introduce an RSA-based 1-2 OT scheme as the foundation for VPOT. The result is that the Chooser need only perform one RSA cubing, i.e., two modular multiplications, per 1-2 OT invocation. When employed in VPOT, this idea reduces the work of the Chooser by over an order of magnitude with respect to the proxy oblivious transfer protocol of NPS.
- *Goldwasser-Micali encryption*: The encryption function E in our brief description above is the Goldwasser-Micali cryptosystem [16]. Encryption in this system takes place with respect to an RSA modulus n . A ‘0’ bit is encrypted as a random quadratic non-residue over Z_n , while a ‘1’ bit is en-

encrypted as a random quadratic residue. The key property of this system is its additive homomorphism. In particular, given encryptions $E[b_0]$ and $E[b_1]$ of bits b_0 and b_1 respectively, the Proxy can non-interactively compute $E[b_0 \oplus b_1]$ as $E[b_0]E[b_1]$. Composition of commitments in this manner enables the Proxy to obtain an efficiently checkable proof of correct decommitment from the Sender, as we shall see. We sometimes refer to a Goldwasser-Micali ciphertext as a *quadratic-residue commitment*, abbreviated QR-commitment. We use the very loose notation $E[b]$ to denote a Goldwasser-Micali encryption of (QR-commitment to) a bit b .

1.3 Other work on auctions

Because of the difficulties involved in deploying standard general secure function evaluation techniques, a number of other secure protocols have been proposed that are specially tailored for auctions. One of the earliest of these is the scheme of Franklin and Reiter [12]. This scheme is not fully private, in the sense that it only ensures the confidentiality of bids until the end of the protocol (although the authors mention a fully private variant). Some more recent schemes include those of Harkavy, Tygar, and Kikuchi [17], Cachin [5], Stubblebine and Syverson [25], Sako [24], Di Crescenzo [11], and Jakobsson and Juels [19]. The Harkavy *et al.* scheme is fully privacy preserving, but involves intensive bidder involvement [17], and is not easily adaptable to different auction types or to related protocols. The scheme of Cachin involves two servers, and requires some communication among bidders. At the end of the protocol, a list of bidders is obtained, but not the bid amounts. The scheme of Di Crescenzo [11] requires no communication between bidders, and has low round complexity, but involves the participation of only a single server. The scheme of Sako [24] works on a different principle from these others, involving opening of bids in what is effectively a privacy-preserving Dutch-style auction. Efficient for small auctions, it involves costs linear in the range of possible bids, and does not accommodate second-price and other auction types. The Jakobsson and Juels [19] protocol aims at streamlining general secure multi-party computation for functions that involve intensive bitwise manipulation, of which auction protocols, as mentioned above, are a good example. A very recent protocol is that of Baudron and Stern [1]. This protocol is expensive, and involves only a single server, with privacy ensured under the condition that there is no collusion between the auction server and any bidder.

Organization

Section 2 reviews some cryptographic building blocks required for our construction. In section 3, we consider some new methods for combining bit commitment with oblivious transfer, and introduce our VPOT protocol. We show how to apply VPOT to the problem of secure function evaluation in section 4. In section 5, we discuss the motivating example: private auction computations. Due to

space constraints in this version of the paper, we do not provide formal security modelling.

2 Building Blocks and Background

We review several standard building blocks for our protocols. Further details regarding these primitives may be found in the literature. We let \in_U denote uniform, random selection from a set. A useful summary of details of the Yao construction may be found in [20].

Private channels: We assume the use of private, authenticated channels between all three possible pairings of the Chooser, Proxy, and Sender. The private channel between the Chooser and Sender involves the Proxy as an intermediary, for the sake of protocol simplification. We assume that messages are authenticated in a non-repudiable fashion. We do not devote attention to the cryptographic elements underlying these channels. In practice, private channels may be realized by way of, e.g., the Secure Sockets Layer protocol (SSL) with supplementary use of digital signatures.

RSA-based 1-2 OT: Recall from above that the aim of 1-2 OT is for the Chooser to obtain a tag t_b for $b \in \{0, 1\}$ from the Sender, who possesses the pair of tags (t_0, t_1) . The Chooser should not learn t_{1-b} , and the Sender should not learn b . Most of the proposed practical 1-2 OT protocols in the literature rely on use of El Gamal encryption or some close variant. As an example, we describe the proxy oblivious-transfer protocol of NPS in detail at the beginning of section 3.

In this paper, we introduce a special, RSA-based 1-2 OT protocol. We do not make direct use of the RSA cryptosystem as such in the construction of this primitive. We do, however, employ the familiar RSA setup [23], which we briefly review here. An RSA public key consists of an RSA modulus $n = pq$, where p and q are primes, and a public exponent e such that $\gcd(e, \phi(n)) = 1$. The corresponding private key d is such that $ed = 1 \pmod{\phi(n)}$. Our protocols involve exclusive knowledge and use of a private RSA key d by the Sender.

As a first step in the 1-2 OT protocol, the Sender must provide the Chooser with double commitments $\mathcal{C}_0 = \mathcal{C}_{k_0}(t_0)$ and $\mathcal{C}_1 = \mathcal{C}_{k_1}(t_1)$ on tags t_0 and t_1 respectively. The Sender additionally selects an integer $C \in_U Z_n^*$, which he sends to the Chooser. The Chooser, wishing to receive tag t_b , chooses an element $x \in_U Z_n^*$. If $b = 0$, the Chooser transmits $(x_0, x_1) = (x^3, Cx^3)$ to the Sender; otherwise, she transmits $(x_0, x_1) = (x^3/C, x^3)$. The Sender checks that $x_1/x_0 = C$. If so, he uses his private key to construct $(z_0, z_1) = (x_0^{1/3}k_0, x_1^{1/3}k_1)$, which he sends to the Chooser. The Chooser then makes use of x to extract k_b in the obvious fashion. Given k_b , the chooser can extract t_b from \mathcal{C}_b as desired.

Lacking knowledge of the cube root C , the RSA assumption implies that the Chooser cannot obtain k_{1-b} . In the random oracle model, then, it may be seen that t_{1-b} is hidden from the Chooser in a semantically secure manner. As

the Sender does not know for which element in the pair (x_0, x_1) the Chooser possesses the corresponding cube root, it may be seen that b is hidden in an information-theoretic sense from the Sender. Our VPOT protocol is essentially a variant on this basic 1-2 OT scheme.

As noted above, our choice of RSA for our protocols stems from a desire to render computation by the Chooser (corresponding to the bidder in an auction protocol) as efficient as possible. We employ $e = 3$, a common choice, in order to render these computations as rapid as possible, although none of our results depends on this fact.

Yao Circuit Evaluation: As discussed above, Yao circuit evaluation [26, 20] serves as the cornerstone of our VPOT protocol, as it does for NPS. Informally, the Yao construction encrypts an entire boolean function, using ciphertexts to represent the 0 and 1's in a table composing a "boolean gate". It is easy to see how any function with finite domain and range can be compiled into a *circuit*, namely a finite set of interdependent boolean gates. Construction of Yao circuits is conceptually straightforward for auction functions, which incorporate a collection of '>' comparisons.

Goldwasser-Micali encryption: The concept of probabilistic encryption was introduced [16] and elaborated on in [4] to set forth the notion of semantic security in an encryption scheme. The basic scheme employs a Blum integer $n = pq$; this is the product of two primes, where each prime is congruent to $3 \pmod{4}$. (To facilitate our security proofs, we assume that the Blum integer employed here is not the same as the RSA modulus employed for 1-2 OT. In practice, and to simplify our protocol descriptions, we believe that use of the same modulus in both cases is acceptable.) The two primes constitute the private decryption key. Encryption is bitwise: a '0' bit is encoded as a random square modulo n , and a '1' bit as a random non-square modulo n with Jacobi symbol 1. In other words, the quadratic residuosity (QR) of a ciphertext indicates the value of the plaintext bit. Knowledge of p and q enables efficient determination of the quadratic residuosity of an element in Z_n .

The Goldwasser-Micali encryption scheme can be employed straightforwardly as a commitment scheme for a player that does not know the factorization of n . To decommit a commitment C_b as a '0' bit, a player provides a square root of C_b modulo n ; to decommit as a '1' bit, she provides a square root of $-C_b$ modulo n . It is easy to see that the scheme is unconditionally binding. Privacy is reducible to the *quadratic residuosity assumption*. Recall from above that a key element of this encryption scheme, and indeed, our reason for employing it, is its useful additive homomorphism: $E[b_0]E[b_1] = E[b_0 \oplus b_1]$. We use this to prove the value of the XOR of two committed bits without revealing any additional information about the individual values of the bits themselves.

3 Verifiable Proxy Oblivious Transfer

As a basis for comparison, we begin by presenting details of the NPS proxy oblivious transfer protocol, whose intuition we sketched above. In an initialization process, the Chooser and Sender agree on a cyclic group G of order w over which computation of discrete logarithms is hard and an associated generator g , as well as a random value $C \in_U G$ whose discrete log is unknown to any player. As usual, we let b denote the choice of the bidder, the pair (t_0, t_1) , the tags held by the Sender. The protocol is as follows [20]:

1. The Chooser selects a private key $x \in Z_w$, and computes a pair $(PK_b, PK_{1-b}) = (g^x, C/g^x)$, and sends PK_0 to the Sender via the Proxy. She sends x to the Proxy.
2. The Sender computes $PK_1 = C/PK_0$. The Sender computes the pair $(z_0, z_1) = (E_{PK_0}[\rho(t_0)], E_{PK_1}[\rho(t_1)])$, where E_{PK_i} denotes El Gamal encryption under public key PK_i and ρ denotes a suitable error-detection function. The Sender transmits the pair (z_0, z_1) to the Proxy in a random order.
3. The Proxy attempts to decrypt both values in the pair using x . The Proxy knows he has obtained t_b when the error-detection function ρ shows that the decryption is successful.

It may be observed that provided there is no collusion, neither the Sender nor the Proxy can learn b . It may be shown that under the Decisional Diffie-Hellman assumption, even if the Proxy and Chooser collude, they cannot learn both t_0 and t_1 . The weakness in this protocol, hinted at above, is the fact that the Sender can choose to send $t_{b'}$ for b' of its choice simply by transmitting $(z_0, z_1) = (E_{PK_0}[\rho(t_{b'})], E_{PK_1}[\rho(t_{b'})])$. Even in the NPS auction protocol, neither the Chooser (i.e., a bidder) nor the Proxy (i.e., the auctioneer) can detect this tampering, which permits arbitrary alteration of bids.

We are now ready to remedy this problem by introducing our VPOT protocol. Due to space constraints, we provide only informal descriptions of the security properties of the protocol.

3.1 Verifiable Proxy Oblivious Transfer (VPOT)

VPOT detects the sort of cheating possible in NPS through use of a zero-knowledge proof based on QR-commitment. Here is a summary: the Sender provides a pair $(\mathcal{C}_0, \mathcal{C}_1)$ of commitments to the tags t_0 and t_1 , in a random order. The Sender also commits to an ordering a of these commitments. In particular, $a = 0$ if \mathcal{C}_0 represents a commitment to t_0 (and \mathcal{C}_1 represents a commitment of t_1); otherwise, $a = 1$. The Sender provides this bit a for the Proxy as a QR-commitment of the form $E[a]$. The Sender obtains a share b_S of b , the ciphertext $E[b_S]$ being observed by the Proxy. The Proxy therefore can compute a commitment to the bit $c = a \oplus b_S$; in particular, $E[c] = E[a]E[b_S]$. If the Sender decommits correctly,

the value c will specify whether the Proxy should be able to open \mathcal{C}_0 or \mathcal{C}_1 . In particular, if $c = 0$, the Proxy should be able to open \mathcal{C}_0 ; otherwise the Proxy should be able to open \mathcal{C}_1 . To prove correct behavior, the Sender decommits c for the Proxy by proving the quadratic residuity of $E[c]$. Observe that the bit c , since it is “masked” by the secret bit a , does not reveal information about b_S to the Proxy. Hence the Proxy does not learn the bit b specifying the tag requested by the Chooser. The following is the full VPOT protocol.

1. The Sender chooses his desired tags $t_0, t_1 \in_U \{0, 1\}^l$ and also an integer $C \in_U Z_n^*$.
2. The Sender computes commitments $\mathcal{C}_0 = \mathcal{C}_{k_a}(t_a)$ and $\mathcal{C}_1 = \mathcal{C}_{k_{1-a}}(t_{1-a})$. Let $u = E[a]$, i.e., a QR-commitment of a . The Sender also computes a commitment $CO = H[u]$ to ordering of $(\mathcal{C}_0, \mathcal{C}_1)$. The Sender transmits the pair $(\mathcal{C}_0, \mathcal{C}_1)$ to the Proxy, along with CO .
3. The Chooser receives C from the Proxy and splits b uniformly at random into bits b_P and b_S such that $b = b_P \oplus b_S$. She also selects $x \in_U Z_n^*$. If $b_P = 0$ she computes $x_0 = x^3$; otherwise, she computes $x_0 = x^3/C$. She also computes $v = E[b_S]$.
4. The Chooser sends (x_0, v, x) to the Proxy. She also sends (x_0, v) to the Sender (via the Proxy, if desired).
5. The Sender receives x_0 and computes $x_1 = Cx_0$. He then computes $y_0 = x_0^{1/3}$ and $y_1 = x_1^{1/3}$. He decrypts b_S . If $b_S = 0$, the Sender transmits the pair $(z_0, z_1) = (y_0k_0, y_1k_1)$ to the Proxy; otherwise he transmits the pair (y_0k_1, y_1k_0) .
6. The Sender transmits u to the Proxy (undoing the outer commitment in CO). The Sender then reveals $c = a \oplus b_S$ by decommitting $uv = E[a]E[b_S] = E[c]$. The decommitment of uv is provided as a value ρ such that $\rho^2 = uv$ if $c = 0$ and $\rho^2 = -uv$ if $c = 1$. The Proxy checks the correctness of these decommitments.
7. The Proxy first computes the cube of both z_0 and z_1 and checks that $H(z_0^3/x_0)$ and $H(z_1^3/x_1)$ are equal to the first element of \mathcal{C}_0 and the first element of \mathcal{C}_1 , in either order. As a final step, the Proxy checks that he can use x to open \mathcal{C}_0 if $c = 0$ and \mathcal{C}_1 if $c = 1$. This check ensures that the Sender decommitted in the correct order.

Security Features:

- The Sender learns no information about b_P . Under the quadratic residuity assumption governing the security of E , the Proxy does not learn b_S . Thus the Sender or Proxy cannot individually determine the Chooser’s choice b .
- The Proxy cannot feasibly compute the cube root of C under the RSA assumption, and therefore cannot learn the cube roots of both x_0 and x_1 . The unselected tag is in fact hidden in a semantically secure sense from the Proxy. This is true even if the Proxy cheats or colludes with the Chooser.
- The Proxy can verify that the Sender has correctly transmitted both t_0 and t_1 , even though he can extract only one of them.

- The Proxy can verify that the Sender has correctly sent him t_b for the bit b selected by the Chooser. Assuming that the Sender and Proxy do not collude, therefore, the Chooser can be assured that the Proxy has received the correct tag t_b .

4 Two-Server Secure-Function Evaluation

In this section we describe an architecture for secure computation based on Yao circuits and VPOT. Due to lack of space, we cannot provide security modeling and proofs for our auction protocol in this paper. As above, we do assume the availability of private, authenticated channels among participating players.

4.1 Putting together VPOT and Yao circuits

We now combine the VPOT protocol with Yao circuit to construct a secure function evaluation protocol involving two servers (evaluators) and multiple contributors of input values. For consistency with our protocol descriptions above, we refer to the two servers as the Proxy and the Sender. Our secure-computation protocol is designed to evaluate functions on inputs contributed by an arbitrarily large number m of players. We refer to these players as Choosers.

Our aim is to evaluate a function F on the m inputs of the Choosers. The Proxy and Sender together evaluate and publish the result of the function computation, and also provide each player with a receipt to guarantee correctness. The role of the Sender here is to construct Yao circuits and that of the Proxy, to evaluate these circuits. To compute input tags for the Yao circuits, these servers must process separate, parallel invocations of VPOT for every individual bit.

The complete function evaluation protocol is as follows:

Offline Steps

1. The Sender generates an RSA modulus n and publishes this for the VPOT invocations in the current function evaluation session. (**Note:** It is in fact critical that a new RSA modulus be published for each session so as to ensure the privacy properties of VPOT across sessions.)
2. The Sender constructs N copies of Yao circuits to evaluate the function F . He sends these circuits to the Proxy, with double commitments to the garbled input tags, as in VPOT. He also publishes a lookup hash table enabling Yao-output-to-plaintext translation.
3. The Proxy selects half of the circuits at random, and asks the Sender to “open” them.
4. The Sender “opens” the selected circuits and sends the keys to all of their committed garbled input tags. This enables verification of their correct construction. This constitutes half of a cut-and-choose proof, the remaining half involving verification of consistent output on the unopened circuits.

5. The Proxy verifies that the “opened” circuits and committed input tags do indeed calculate the correct function.

VPOT steps

1. The Choosers submit their inputs bitwise to the Proxy according to the VPOT protocol.
2. The Proxy forwards these choices to the Sender according to the VPOT protocol.
3. The Sender sends the garbled input tags according to VPOT for each input bit, and also each of the $N/2$ unopened circuits.
4. If either the Proxy or Sender detects the presence of an ill-formed input by a Chooser, this is proven to the other server. Together the two servers can annul the input of any Chooser, provided that F is suitably constructed. Details are straightforward, and omitted here.

Circuit Evaluation

1. The Proxy checks the garbled tags against the commitments, and evaluates the unopened $N/2$ Yao circuits.
2. The Proxy looks up the Yao circuit outputs in the lookup tables, and verifies that the results of the $N/2$ trials are identical.
3. The Proxy publishes the output entries of the Yao tables, along with the function output. If the entries and output are correct, then the Sender certifies the output.

We remark that the Proxy should not publish the garbled output strings if the outputs are inconsistent. Such a situation only occurs if the Sender cheats, and revealing the outputs might leak information about input values. Once the correct output values are published, the result is verifiable by any outside party.

5 Application to Auctions

The two-server secure-function evaluation scheme presented in the previous section can be applied straightforwardly, of course, to create a sealed-bid auction system. As auctions are our key motivating application for the work in this paper, it is worth a brief, concluding discussion of the particular benefits of our approach to auctions.

As explained above, our scheme in this paper addresses the flaw in the NPS auction protocol [20]. The NPS protocol is privacy preserving, but effectively operates (unbeknownst to the authors) under the assumption that both the servers are honest. The flaw in this paper is simple: the sender may flip or set constant the two tags which he sends in the oblivious transfer for a given bit, e.g., he can send only '0' tags for a given bit submitted by a bidder. This allows

the Sender to change the bid of any bidder to any value that the Sender chooses. Nonetheless, we believe that NPS offer a key insight in suggesting a two-server model to exploit the high computational efficiency and low round complexity of the Yao construction. This represents an important step toward the realization of practical, privacy-preserving auction protocols.

The secure-function evaluation procedure that we propose in section 4 not only fixes the flaw in the NPS protocol but, as already noted, has the additional benefit of substantially reducing the computation required by bidders. In summary, then, our proposed architecture offers the following features:

1. *Non-interactivity*: Bidders submit bids in a non-interactive fashion. That is, they present their bids to the servers, but need not participate subsequently in the auction protocol except to learn the outcome.
2. *Auction adaptability*: Our auction protocol is readily adaptable with little overhead to a range of auction types, such as highest-price auctions and Vickrey auctions.
3. *Full privacy*: We characterize privacy in terms of a static, active adversary that controls at most one server and an arbitrary number of bidders. The only information revealed to such an adversary at the conclusion of the protocol about the bids of any honest bidders is the outcome of the auction. In a highest-price auction, for example, such an adversary learns only the winning bid and the identity of the winning bidder.
4. *Correctness*: Any player can be assured of the correctness of the auction execution assuming that the two auction servers do not collude.
5. *Robustness*: While we do not achieve robustness against failure by either of the auction servers, the servers can eliminate any ill-formed bids and process the remaining ones correctly.
6. *Low round-complexity*: The protocol involves only five communication passes; this includes the offline cut-and-choose proof of correct Yao circuit construction.
7. *High computational efficiency*: Our protocol is highly efficient in terms of computational requirements. For bidders, it is more so than any other cryptographically based privacy-preserving auction scheme in the literature. The requirement for a bidder in a typical auction would be several tens of modular multiplications (as opposed to a comparable number of modular *exponentiations* in NPS). The cost for the servers is about twice that in NPS. (While in general it is desirable to shed server load in favor of computation on the part of clients, the NPS protocol is so computationally intensive for clients as to pose a likely bottleneck even for reasonably powerful handheld devices.)

The principal drawback of our scheme is that, like the NPS protocol, it does not extend to a trust model involving more than two servers. Whether or not the NPS scheme can incorporate multiple servers is an open research question.

We emphasize that given the successful implementation experiments of NPS, our proposed architecture is likely to be amenable to practical software deployment. With this in mind, we provide a brief efficiency analysis in the appendix.

References

1. O. Baudron and J. Stern. Non-interactive private auctions. In S. Haber, editor, *Financial Cryptography '01*, pages 303–313, 2001.
2. D. Beaver. Minimal-latency secure function evaluation. In B. Preneel, editor, *Advances in Cryptology - Eurocrypt '00*, pages 335–350. Springer-Verlag, 2000. LNCS no. 1807.
3. M. Bellare, S. Micali, and P. Rogaway. The round complexity of secure protocols. In *ACM CCS '90*, pages 503–513. ACM Press, 1990.
4. M. Blum and S. Goldwasser. An efficient probabilistic public-key encryption scheme which hides all partial information. In G.R. Blakely and D. Chaum, editors, *Advances in Cryptology - Crypto '84*, pages 289–299. Springer-Verlag, 1985. LNCS No. 196.
5. C. Cachin. Efficient private bidding and auctions with an oblivious third party. In G. Tsudik, editor, *ACM CCS '99*, pages 120–127. ACM Press, 1999.
6. C. Cachin and J. Camenisch. Optimistic fair secure computation. In M. Bellare, editor, *Advances in Cryptology - Crypto '00*, pages 94–112. Springer-Verlag, 2000. LNCS no. 1880.
7. C. Cachin, J. Camenisch, J. Kilian, and J. Muller. One-round secure computation and secure autonomous mobile agents, 2000.
8. R. Cramer, I. Damgård, and J.B. Nielsen. Multiparty computation from threshold homomorphic encryption. In B. Pfitzmann, editor, *Advances in Cryptology - Eurocrypt '01*, pages 280–300. Springer-Verlag, 2001. LNCS no. 2045.
9. Claude Crépeau. Verifiable disclosure of secrets and applications. In J.J. Quisquater and J. Vandewalle, editors, *Advances in Cryptology - Eurocrypt '89*, pages 181–191. Springer-Verlag, 1990. LNCS no. 434.
10. Claude Crépeau, van de Graaf, Jeroen, and Alain Tapp. Committed oblivious transfer and private multi-party computation. In D. Coppersmith, editor, *Advances in Cryptology - Crypto '95*, pages 110–123. Springer-Verlag, 1995. LNCS No. 963.
11. G. Di Crescenzo. Private selective payment protocols. In P. Syverson, editor, *Financial Cryptography '00*, 2000.
12. M. Franklin and M. Reiter. The design and implementation of a secure auction server. *IEEE Transactions on Information Theory*, 22(5):302–312, 1996.
13. M. Franklin and M. Yung. Varieties of secure distributed computing. In *Proc. Sequences II, Methods in Communications, Security and Computer Science*, pages 392–417, 1991.
14. T. El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31:469–472, 1985.
15. O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *STOC '87*, pages 218–229. ACM Press, 1987.
16. S. Goldwasser and S. Micali. Probabilistic encryption. *J. Comp. Sys. Sci.*, 28(1):270–299, 1984.
17. M. Harkavy, J.D. Tygar, and H. Kikuchi. Electronic auctions with private bids. In *3rd USENIX Workshop on Electronic Commerce*, pages 61–73, 1999.
18. M. Hirt, U. Maurer, and B. Przydatek. Efficient secure multi-party computation. In T. Okamoto, editor, *Advances in Cryptology - Asiacrypt '00*, pages 143–161. Springer-Verlag, 2000. LNCS No. 1976.
19. M. Jakobsson and A. Juels. Mix and match: Secure function evaluation via ciphertexts. In T. Okamoto, editor, *Advances in Cryptology - Asiacrypt '00*, pages 162–177. Springer-Verlag, 2000. LNCS No. 1976.

20. M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *1st ACM Conf. on Electronic Commerce*, pages 129–139. ACM Press, 1999.
21. M. Rabin. How to exchange secrets by oblivious transfer, 1991. Tech. Memo TR-81 Aiken Computation Laboratory, Harvard University.
22. R. L. Rivest. Unconditionally secure commitment and oblivious transfer schemes using private channels and a trusted initializer, 1999.
23. R. L. Rivest, A. Shamir, and L. M. Adelman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21:120–126, 1977.
24. K. Sako. An auction protocol which hides bids of losers. In H. Imai and Y. Zheng, editors, *PKC '00*, pages 422–432. Springer-Verlag, 2000. LNCS no. 1751.
25. Stuart G. Stubblebine and Paul F. Syverson. Fair on-line auctions without special trusted parties. In *Financial Cryptography*, pages 230–240, 1999.
26. A.C. Yao. Protocols for secure computations (extended abstract). In *FOCS '82*, pages 160–164. IEEE Computer Society, 1982.

A Efficiency Considerations

The protocol VPOT involves both offline and online calculations; the latter may be considered of more practical relevance. A typical implementation might use a 1024-bit RSA modulus with exponent 3. Disregarding the cost of hash functions computations, which is relatively small, we observe that the Sender must compute seven modular multiplications offline. Online, the Sender must calculate three modular exponentiations. The Proxy has much less computational expense: only five modular multiplications and two modular divisions. Best of all, the Chooser need only calculate five modular multiplications per bit selection. Note that these are the costs for only one invocation of VPOT. A full auction protocol will involve many, of course, as we now consider.

A.1 A typical auction

To provide a flavor of the resource requirements for our proposed architecture, we summarize the computational requirements in a typical auction setting. We omit the relatively small cost of private channel establishment (via, e.g., SSL) and negligible cost of hash calculations; we count circuit evaluations as a unit.

In our example there are 10 bidders in the auction, the bids are 10 bits long, and 10 circuits out of 20 remain after the cut-and-choose step. The Sender must create the 20 circuits offline, and he can also calculate 10,000 of his modular multiplications off-line. During the protocol, he must calculate 2000 modular multiplications and 2000 modular exponentiations. The Proxy must evaluate 20 circuits accepting 100 inputs each, calculate 10,000 modular multiplications, and 2000 modular divisions. About half of this effort can be done off-line before bidding commences. Finally, the Choosers (bidders) need only perform at most 50 modular multiplications each in total to construct their bids.