

Defining Verifiability in e-Auction Protocols

Jannik Dreier
Université Grenoble 1, CNRS
jannik.dreier@imag.fr

Hugo Jonker
University of Luxembourg
hugo.jonker@uni.lu

Pascal Lafourcade
Université Grenoble 1, CNRS
pascal.lafourcade@imag.fr

ABSTRACT

An electronic auction protocol will only be used by those who trust that it operates correctly. Therefore, e-auction protocols must be *verifiable*: seller, buyer and losing bidders must all be able to determine that the result was correct. We pose that the importance of verifiability for e-auctions necessitates a formal analysis. Consequently, we identify notions of verifiability for each stakeholder. We formalize these and then use the developed framework to study the verifiability of two examples, the protocols due to Curtis et al. and Brandt, identifying several issues.

Categories and Subject Descriptors

C.2.2 [Network Protocols]: Protocol verification; K.4.4 [Electronic Commerce]: Security

Keywords

Verifiability; e-Auction; Formal Methods; Protocol Analysis

1. INTRODUCTION

Auctions provide sellers and buyers with a way to exchange goods for a mutually acceptable price. Unlike a marketplace, where the sellers compete with each other, auctions are a seller's market where *buyers* bid against each other over the goods for sale. There are many different types of auctions, varying how to determine winner and price. For example, in an English auction, buyers bid publicly against each other and the highest bid wins (e.g. [16]). A Vickrey auction is rather similar, except that the winning buyer pays the price of the second-highest bid (see e.g. [10, 14]). Conversely, in a Dutch auction, the price starts too high – the auctioneer keeps lowering the price until a buyer claims the good for that price (e.g. [17]). Sealed-bid auctions are auctions that run in one round: in the bidding phase, each buyer submits one sealed bid, which are then simultaneously opened in the opening phase. Sealed-bid auctions can be used to implement auctions where the bidder with the

highest bid wins (e.g. [4, 15, 1, 5]), but also Vickrey-style auctions, where the winner pays the second-highest price (e.g. [14]). In this paper, we focus on verifiability of sealed-bid auctions (though our results are applicable to English and, in a limited sense, Dutch auctions).

Auctions involve the following stakeholders:

- *Bidders*: prospective buyers, who want to pay as little as possible.
- *Seller*: a seller, striving to sell for as much as possible.
- *Auctioneer*: the party organizing the auction.

As is readily apparent, the interests of the various stakeholders are opposed. Buyers are in competition with each other for the goods on sale, sellers are in competition with buyers for the price of the goods, the auctioneer may profit directly from overvalued sale price (which provides an incentive to collude with the seller), but a reputation for undervalued sale prices will ensure many repeat customers (which provides incentive to collude with buyers). Consider e.g. the case where there are several bids for the same price. In such a case, an auctioneer might prefer the most “active” bidder instead of the normal tie-breaking rules, and so favor frequent customers over occasional ones.

There are thus no impartial parties to oversee the correctness of the process to determine selling price and winner. For this reason, an auction system must provide some form of *verifiability* for each involved party – irrespective of how the auction process is run and the winner is determined.

Auction verifiability is easy to achieve in isolation, as happens in English “shout-out” auctions. However, maintaining verifiability while ensuring other properties (non-repudiation, privacy, etc.) is far harder. Too often, newly proposed auction protocols proudly show how they achieve these other properties, while only acknowledging the requirement for verifiability in passing. Typically, verifiability is subsequently claimed without providing any formal proof, e.g., [3, 6]. To address this, we propose a generic formal framework applicable independent of the type of auction. The framework consists of formal tests of verifiability.

Contribution. The main contribution of this work is to identify a set of scheme-independent definitions, which, taken together, cover verifiability of auctions. To this end, we focus on the bidders (distinguishing verifiability for losing bidders from the winning bidder) and the seller. We present this framework as a set of formal verifiability tests. Moreover, we investigate the auction protocols by Curtis et al. [6] and by Brandt [3]. Both protocols claim verifiability, yet we identify issues with each.

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

ASIA CCS'13, May 8–10, 2013, Hangzhou, China.

Copyright 2013 ACM 978-1-4503-1767-2/13/05 ...\$15.00.

Related work. There are relatively few formal analyses of auction protocols. Dong et al. [7] study privacy properties of the protocol by Abe and Suzuki [1] in the applied π -calculus. More recently, Dreier et al. [8] used the applied π -calculus to formalize several properties (privacy, non-repudiation, non-cancellation and fairness, but not verifiability) for auction protocols, and studied (and found problems with) two auction protocols. Besides these, verifiability in auctions has (to the best of our knowledge) only been studied for particular schemes. However, in the field of voting several more generic definitions of verifiability have emerged, and we look there for inspiration.

In voting, the property of *individual verifiability* – a voter can verify that her vote counts correctly for the result – has been a well-established notion since the field’s inception [9, 2, 18, 11]. Sako and Killian [18] introduced the concept of *universal verifiability*: the property that any observer may verify (using only public information) the correctness of the result. Kremer et al. [12] introduced the notion of *eligibility verifiability*: the property whereby any observer may verify, using only public information, that the set of cast votes from which the result is determined originates only from eligible voters, and each eligible voter cast at most one vote. Finally, Küsters et al. [13] introduced the notion of *accountability*: when verifiability fails, it is possible to identify the person responsible for the failure.

While the intuition behind these notions carries (to some degree) over to auctions, we do note, that unlike voting, auctions involve *competing* parties – an illegal bid (e.g., by the seller) may increase the winning price, while not changing the winner. Hence, a lack of verifiability in auctions can compromise fairness. In addition, winner/price determination can be quite complex depending on the type of auction, e.g. for second-price or multi-good auctions. Thus, verification of voting systems does not translate directly to verification of auction systems.

Outline. In Section 2, we describe our modeling of auctions. In Section 3, we formalize the verifiability definitions, taking into account the point of view of the seller and the bidders. In Section 4, we then apply our model to two examples and exhibit problems in both cases.

2. MODELING AUCTION PROTOCOLS

We consider a set of bidders \mathcal{B} and a seller S . We do not model other parties as only bidders and the seller verify the execution of the protocol.

Bids are of type Bid (in the simplest case just a price). When being submitted the bids might be encrypted or anonymized to ensure privacy, hence we use the type $EBid$ for such bids. We assume that there is a public list \mathbb{L} of length n and type $List(EBid)$ of all submitted bids, for example a bulletin board. To define the soundness of the verification tests we need a mapping between both types, i.e. a function $getPrice: EBid \mapsto Bid$ that gives the bid for an encrypted bid. This function does not need to be computable for any party, as it is only used in the soundness definition.

Bidders have to register at some point, or are otherwise authenticated when bidding, in order to be able to obtain their goods once the auction has ended. This could for example be implemented using signatures, authentication tokens, MACs etc. Therefore we require a function $isReg: EBid \mapsto \{true, false\}$ that returns *true* if a bid was submitted by a

registered bidder, and not modified – this integrity protection is necessary to prevent manipulation of bids.

In addition, we require a public function that - given a list of bids - computes the index of the winning bid within the list of all bids: $win: List(Bid) \mapsto Index$. This might simply be the index of the maximal bid among all bids, but there may be more complex operations to determine this index depending on the type of auction or to deal with ties (i.e. several maximal bids).

Lastly, we assume that the variable $winBid$ of type $Index$ refers to the index of the announced winning bid at the end of the auction, and that each bidder has a variable $myBid$ of type $Index$ that refers to the index of his bid in \mathbb{L} .

Note that for a list l we write $l[i]$ to denote the i -th element of the list starting with 1, and $Indices(l)$ to denote the set of indices of l , i.e. $\{1, \dots, n\}$ if l contains n elements.

DEFINITION 1. *An auction protocol is a tuple $(\mathcal{B}, S, \mathbb{L}, getPrice, isReg, win, winBid)$ where*

- \mathcal{B} is the set of bidders,
- S is the seller,
- \mathbb{L} is a list of all submitted bids,
- $getPrice: EBid \mapsto Bid$ is a function mapping submitted bids to bids,
- $isReg: EBid \mapsto \{true, false\}$ is a function that returns *true* if a bid was submitted by a registered bidder,
- $win: List(Bid) \mapsto Index$ is a function that returns the index of the winning bid,
- $winBid$ is a variable referring to the index of the winning bid at the end of the auction.

3. DEFINING VERIFIABILITY

In this section, we formally define verifiability for auction protocols. In the first part we consider only first-price auctions. Thereafter we generalize the definitions to account for second-price, multi-price, and other types of auctions.

3.1 First-Price Auctions

To understand which verifications are needed, we start by discussing three different stakeholder’s perspectives:

- A *losing bidder* wants to be convinced that he actually lost, i.e. that:
 - the winning bid was actually superior to his bid (as defined by the win function), and
 - that the winning bid was submitted by another bidder (preventing both seller and auctioneer from maliciously adding or manipulating bids to influence the final price).
- A *winning bidder* wants to check that:
 - he actually submitted the winning bid,
 - the final price is correctly computed,
 - all other bids originated from bidders, and
 - no bid was modified.

Together, these verification checks ensure that the winning bidder is indeed the correct winner, for the correct price. Moreover, the last two checks ensure that the auction process was only influenced by legitimate bidders – neither seller nor auctioneer influenced the process.

- The *seller* wants to verify that:
 - the announced winner is correct, and
 - the winning price is correct,

in particular if the outcome of the auction was not determined publicly (e.g. privately by the auctioneer, or using distributed computations among the bidders).

To execute these verifications, we introduce the notion of *Verification Tests*.

We define a *Verification Test* as an efficient terminating algorithm that takes as input the data visible to a participant of an auction protocol and returns a Boolean value.

We deliberately do not specify more details at this point as they will depend on the underlying protocol model. Such a test could be a logical formula (whose size is polynomial in the input) in a symbolic model or a polynomial-time Turing-machine in a computational model. Obviously there can be different tests for different participants (e.g. for bidders and the seller), since they may have different views of the protocol execution.

We define verifiability as follows.

DEFINITION 2 (VERIFIABILITY - 1ST-PRICE AUCTIONS).
An auction protocol $(\mathcal{B}, S, \mathbb{L}, \text{getPrice}, \text{isReg}, \text{win}, \text{winBid})$ ensures Verifiability if we have Verification Tests $rv_s, rv_w, ov_l, ov_w, ov_s$ respecting the following soundness conditions:

1. *Registration and Integrity Verifiability (RV):*
 - *Anyone can verify that all bids on the list were submitted by registered bidders:*
 $rv_s = \text{true} \implies \forall b \in \mathbb{L}: \text{isReg}(b) = \text{true}$
 - *Anyone can verify that the winning bid is one of the submitted bids:*
 $rv_w = \text{true} \implies \text{winBid} \in \text{Indices}(\mathbb{L})$
2. *Outcome Verifiability (OV):*
 - *A losing bidder can verify that his bid was not the winning bid:*
 $ov_l = \text{true} \implies \text{myBid} \neq \text{win}(\text{getPrice}(\mathbb{L}))$
 - *A winning bidder can verify that his bid was the winning bid:*
 $ov_w = \text{true} \implies \text{myBid} = \text{win}(\text{getPrice}(\mathbb{L}))$
 - *The seller can verify that the winning bid is actually the highest submitted bid:*
 $ov_s = \text{true} \implies \text{winBid} = \text{win}(\text{getPrice}(\mathbb{L}))$

as well as the following completeness condition:

- *If all participants follow the protocol correctly, the above tests succeed (i.e., the implications hold in the opposite direction, \Leftarrow , as well).*

where – with abuse of notation – we write $\text{getPrice}(\mathbb{L})$ for $\text{getPrice}(\mathbb{L}[1]), \dots, \text{getPrice}(\mathbb{L}[n])$.

Consider the perspective of a losing bidder: He can verify that his bid was not the winning bid (ov_l), and that the winning bid was among the ones submitted by registered bidders, which were also not modified (rv_s and rv_w). Similarly a winning bidder can check that his bid was actually the winning bid (ov_w), and that the other bids were submitted by other bidders and not modified (rv_s). Lastly, the seller can also check that the bids used for computing the winner were submitted only by registered bidders (rv_s and rv_w), and that the outcome was correct (ov_s). Hence these tests cover all the verifications discussed above.

In the case of soundness, we require the conditions to hold even in the presence of malicious participants (since the tests should check if they did their work correctly), whereas in the case of completeness we only consider honest participants. This is necessary as otherwise e.g. a dishonest auctioneer could announce the correct result, but publish incorrect evidence. Hence the verification tests fail although the outcome is correct, but this is acceptable since the auctioneer did not “work correctly” in the sense that he deviated from the protocol specification.

Definition 2 can be applied to sealed-bid auctions, where all bids are submitted in a private way, as well as English auctions, where the price increases with each publicly announced bid. These latter are verified by applying the verification tests after each price increase.

Example.

Consider a simple auction system where all bidders publish their (not encrypted and not signed) bids on a bulletin board, and at the end of the bidding phase the auctioneer announces the winner. In this case there is a simple test for rv_w : anyone can simply test if the winning bid is one of the published ones. However there is no test for rv_s since bids are not authenticated. If we require bidders to sign their bids before publishing them, we also have a simple test for rv_s : verifying the signatures.

It is clear that we have simple tests for ov_l, ov_w and ov_s since everybody can compute the winner on the public list of unencrypted bids. This however means that the protocol ensures no privacy, and no fairness since a bidder can choose his price depending on the previously submitted bids. If we add encryption for the bids to address this shortcoming, the situation becomes more complex and the auctioneer has to prove that he actually computed the winner correctly, for example using zero-knowledge proofs.

3.2 Other Types of Auctions

Our definition can be extended to other auctions, including second-price auctions, more general $(M+1)$ st-price auctions, and even bulk-good auctions that have multiple winners at different prices. The price in these types of auctions may also depend on the other submitted bids – not only on the winning bid. To deal with this, we enrich our model of an auction protocol with a type *Price*. The function win now returns lists of winners and prices $\text{win}: \text{List}(\text{Bid}) \mapsto \text{List}(\text{Index}) \times \text{List}(\text{Price})$. We also assume that there are two variables winPrice and myPrice instantiated as the announced list of winning prices and the price announced to a winning bidder respectively. Similarly winBid is now instantiated as a list of indices of bids.

For such auctions, registration verifiability does not change,

but winner(s) and seller also want to verify the price they pay to prevent a malicious party from increasing price(s).

DEFINITION 3 (GENERALIZED VERIFIABILITY). *An auction protocol $(\mathcal{B}, S, \mathbb{L}, \text{getPrice}, \text{isReg}, \text{win}, \text{winBid}, \text{winPrice})$ ensures Verifiability if we have Verification Tests $rv_s, rv_w, ov_l, ov_w, ov_s$ respecting the following conditions:*

1. *Soundness:*

(a) *Registration and Integrity Verifiability (RV):*

- *Anyone can verify that all bids on the list were submitted by registered bidders:*
 $rv_s = true \implies \forall b \in \mathbb{L}: \text{isReg}(b) = true$
- *Anyone can verify that the winning bids are among the submitted bids:*
 $rv_w = true \implies$
 $\forall \mathbf{b} \in \text{winBid}: b \in \text{Indices}(\mathbb{L})$

(b) *Outcome Verifiability (OV):*

Let $(\text{indexes}, \text{prices}) = \text{win}(\text{getPrice}(\mathbb{L}))$.

- *A losing bidder can verify that his bid was not the winning bids:*
 $ov_l = true \implies \mathbf{myBid} \notin \text{indexes}$
- *A winning bidder can verify that his bid was among the winning bids, and that his price is correct:*
 $ov_w = true \implies \exists i: (\mathbf{myBid} = \text{indexes}[i] \wedge \mathbf{myPrice} = \text{prices}[i])$
- *The seller can verify that the list of winners and the winning prices are correctly determined:*
 $ov_s = true \implies$
 $(\text{winBid} = \text{indexes} \wedge \text{winPrice} = \text{prices})$

2. *Completeness: If all participants follow the protocol correctly, the above tests succeed (i.e., the implications hold in the opposite direction, \Leftarrow , as well).*

where – with abuse of notation – we write $\text{getPrice}(\mathbb{L})$ for $\text{getPrice}(\mathbb{L}[1]), \dots, \text{getPrice}(\mathbb{L}[n])$. Differences to Def. 2 are marked in **bold**.

Note that e.g. in the case of a second-price auction verifying the price, for example in test ov_w , may implicitly include some more registration verification, namely checking that the second-highest bid was actually submitted by a bidder. Otherwise a malicious seller could add a higher second-highest bid or manipulate the existing one to achieve a higher selling price. This is however included in our model as the function win only works on the list \mathbb{L} , hence adding another bid later on to manipulate the bidding price violates the test, and adding or manipulating a bid in \mathbb{L} violates rv_s .

4. CASE STUDIES

In this section, we discuss two case studies: The protocols by Curtis et al. [6] and Brandt [3].

4.1 Protocol by Curtis et al. [6]

The protocol by Curtis et al. [6] was designed to support any type of sealed-bid auction while guaranteeing fairness, privacy, verifiability and non-repudiation.

4.1.1 Informal Description

The main idea of the protocol is the following: using a public-key infrastructure (PKI), the bidders register with a trusted Registration Authority (RA), who issues pseudonyms that will then be used for submitting bids to the Seller (S). The seller eventually receives all bids in clear and can hence apply any possible win function. However, he can only link bids to pseudonyms, not to bidders. The protocol is split into three phases: Registration, Bidding, and Winner determination.

- *Registration:* Each bidder sends his identity, a hash of his bidding price b_i and a signature of $h(b_i)$ to the RA. The RA checks the identity and the signature using the PKI, and replies with an encrypted and signed message containing a newly generated pseudonym p and the hashed bid $h(b_i)$.
- *Bidding:* The RA generates a new symmetric key k . Each bidder will send $c = \text{Enc}_{pk_S}(b_i)$, his bid b_i encrypted with the seller's public key, and a signature of c , together with his pseudonym to the RA. The RA will reply with a signature on c , and encrypts the bidders message, together with the hashed bid $h(b_i)$ from phase one, using the symmetric key k . This encrypted message is then send to the seller.
- *Winner determination:* After all bids have been submitted, the RA will reveal the symmetric key k to the seller. The seller can then decrypt the bids, verify the correctness of the hash and determine the winner. To identify the winner using the pseudonym he can ask the RA to reveal the true identity.

4.1.2 Formal Model

We have the set of bidders \mathcal{B} and a seller S . We do not need to specify the type of bids Bid since the protocol supports any type of bids. The bids are published when the auctioneer reveals the symmetric key, i.e. \mathbb{L} contains bids of the following type: $(Pseudo \times PEnc(Bid) \times Hash)$, where $Pseudo$ is the type of pseudonyms, $PEnc$ is a public-key encryption and $Hash$ are hash values. The function getPrice will simply decrypt the encrypted bid (the second entry of the tuple). The function isReg will return $true$ if and only if the hash value is correct, the pseudonym was actually attributed by the RA and the bid was submitted correctly signed by the bidder with this pseudonym. The protocol is independent of the used auction mechanism and hence does not define win . The seller will simply decrypt all bids and can then apply any function win . He will publish the winning price and the winning bidders pseudonym, and winBid will denote the index of the bid containing this pseudonym.

4.1.3 Analysis

Since the seller does the winner determination on his own, there is a simple test for ov_s : He can check his own computations. As the computation of the winner is not specified in order to support any type of auction, we cannot give tests for ov_l and ov_w – they would have to be designed as a function of the used auction algorithm. Yet there is also a test for rv_w : Checking if the pseudonym appears in the list of bids.

However, the messages from the RA to the seller are not authenticated, hence there can be no suitable tests for rv_s .

once the (encrypted) bids are revealed. Even if they were authenticated, this still requires trusting the RA, since there is no way to verify if a pseudonym actually corresponds to a bidder. This also shows a simple attack: the RA can create a new pseudonym and submit a bid under this pseudonym, which may allow him to manipulate the auction outcome.

Curtis et al. explicitly state that the RA needs to be trusted. However, they themselves propose a property “robustness” which states that an auction protocol should be able to handle corrupt behavior. Clearly, their protocol is not robust to dishonest behavior by the RA, and verifiability of their protocol requires trust in the RA. We argue that basing verifiability tests on such a trust assumption at least partly contradicts the main point of verifiability, which is to eliminate such trust assumptions by providing evidence that a trusted party actually behaved honestly.

4.2 Protocol by Brandt [3]

The protocol by Brandt [3] realizes a first-price sealed-bid auction and was designed to ensure full privacy in a completely distributed way. It exploits the homomorphic properties of a distributed El-Gamal Encryption scheme for a secure multi-party computation of the winner.

4.2.1 Informal Description

The participating bidders and the seller communicate using a bulletin board, i.e. an append-only memory accessible for everybody. The bids are encoded as bit-vectors where each entry corresponds to a price. The protocol then uses linear algebra operations on the bid vectors to compute a function f_i , which returns a vector containing one entry “1” if bidder i submitted the highest bid, and different numbers ($\neq 1$) otherwise. To be able to compute this function in a completely distributed way, and to guarantee that no coalition of malicious bidders can break privacy, these computations are performed on the encrypted bids using homomorphic properties of a distributed El-Gamal encryption.

In a nutshell, the protocol realizes the following steps:

1. Firstly, the distributed key is generated: each bidder chooses his part of the secret key and publishes the corresponding part of the public key on the bulletin board.
2. Each bidder then computes the joint public key, encrypts his bid using this key and publishes it on the bulletin board.
3. Then function f_i is calculated for every bidder i using the homomorphic property of the encryption scheme.
4. The outcome of this computation (n encrypted values) are published on the bulletin board, and each bidder partly decrypts each value using his secret key.
5. These shares are sent to the seller, who can combine all to obtain the result (i.e. all f_i). He publishes part of the shares such that each bidder j can only compute his f_j to see if he won or lost (using his knowledge and the published shares), but not the other f_i .

4.2.2 Formal Model

We have a set of bidders \mathcal{B} and a seller S . The list of all submitted bids \mathbb{L} is published on the bulletin board. The

function $getPrice(C)$ decrypts the bid using the joint private key. The function win returns the index of the highest bid submitted, in case of ties the one submitted by the bidder with the smallest index. The protocol has two particularities: Firstly there is no registration (and hence no meaningful function $isReg$), and secondly the winner is not publicly announced – only the winning bidder and the seller know at the end who won. We can still assume that $winBid$ gives the index of the winning bid, although only the seller and the winning bidder have access to it. We assume that there is a “magical” function $isReg$ that can check if a bid was submitted by a registered bidder, however the absence of registration and authentication means that we cannot implement it.

4.2.3 Analysis

The protocol includes no authentication or registration, hence there is no suitable test for rv_s . An attacker may hence submit bids on behalf of a bidder, which cannot be detected using a verification test. Yet using the values published on the bulletin board everybody can check if the values used for the computation were the previously submitted bids, and as the winning index will be among them, we have a test for rv_w .

The author claims that the protocol is verifiable as the parties have to provide zero-knowledge proofs for their computations, however there are two problems.

Firstly a winning bidder cannot verify if he actually won. To achieve privacy, the protocol hides all outputs of f_i except for the entry containing “1”. This is done by exponentiation of all entries x_i of the return vector x with random values, i.e. by calculating $x_i^{\sum_j r_j}$. If x_i is one, this will still return one, but a random value for any other value of x . Yet these random values r_j may add up to zero (mod q), hence the returned value will be $x_i^0 = 1$ and the bidder will conclude that he won ($x_i = 1$), although he actually lost ($x_i \neq 1$). Hence simply verifying the proofs is not sufficient – such a test ov_w would not be sound. For the same reason the seller might observe two or more “1”-values even though all proofs are correct, and will be unable to decide which bidder actually won. He could even exploit such a situation to his advantage: He can simply tell both bidders that they won and take money from both, although there is only one good to sell. If the bidders do not exchange additional data there is no way for them to discover that something went wrong, since the seller is the only party having access to all values. The probability of the random values adding up to zero is low, yet this means that there are cases where the verifiability tests are not sound.

Secondly the author does not exactly specify the proofs that have to be provided in the joint decryption phase. If the bidders only prove that they use the same private key on all decryptions (and not also that it is the one they used to generate their public key), they may use a wrong one. This will lead to a wrong decryption where with very high probability no value is “1”, as they will be random. Hence all bidders will think that they lost, thus allowing a malicious bidder to block the whole auction, as no winner is determined. Hence, if we assume that ov_l consists in verifying the proofs, a bidder trying to verify that he lost using the proofs might perform the verification successfully, although the result is incorrect and he actually won – since he would have observed a “1” if the vector had been correctly

decrypted. This problem can be addressed by requiring the bidders to also prove that they used the same private key as in the key generation phase.

5. CONCLUSION

In this work, we identified the types of verifiability necessary for the stakeholders in auctions. We then formalized these requirements in a protocol-independent way, resulting in tests $rv_s, rv_w, ov_l, ov_w, ov_s$, which together constitute a general verifiability framework for auction protocols.

We illustrated the use of the proposed tests by two case studies, that analyzed the auction protocols by Curtis et al. [6] and by Brandt [3], respectively. The protocol by Curtis et al. is correct only for a trusted Registration authority – which runs contrary to the point of verification: that the authorities no longer need to be trusted. Brandt’s protocol does not have sound verifiability tests: it is technically possible for a losing bidder to conclude he won. Moreover, it may also be possible for a bidder to prevent anyone from winning by using a wrong decryption key. To prevent this, bidders must prove that the private key matches the previously announced public key. Additionally, the protocol does not provide sufficient registration and authentication mechanisms to allow registration verifiability.

Future work. We are currently working on a full application of these definitions to various auction protocols in both the symbolic model and the computational model.

Looking further ahead, we are interested in the full relationship between fairness and verifiability in auctions. As illustrated, there exist verifiability requirements without which violations of fairness may occur. The exact relationship between fairness and verifiability however is an open question.

6. ACKNOWLEDGMENTS

This work was partly supported by the ANR project ProSe (decision ANR-2010-VERS-004-01).

7. REFERENCES

- [1] M. Abe and K. Suzuki. Receipt-free sealed-bid auction. In *Proc. 5th Conference on Information Security*, volume 2433 of *LNCS*, pages 191–199. Springer, 2002.
- [2] J. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections (extended abstract). In *Proc. 26th Symposium on Theory of Computing, STOC '94*, pages 544–553, New York, NY, USA, 1994. ACM.
- [3] F. Brandt. How to obtain full privacy in auctions. *International Journal of Information Security*, 5:201–216, 2006.
- [4] C. Cachin. Efficient private bidding and auctions with an oblivious third party. In *Proc. 6th Conference on Computer and Communications Security (CCS'99)*, pages 120–127. ACM Press, 1999.
- [5] X. Chen, B. Lee, and K. Kim. Receipt-free electronic auction schemes using homomorphic encryption. In *Proc. 6th Conference on Information Security and Cryptology*, volume 2971 of *LNCS*, pages 259–273. Springer, 2003.
- [6] B. Curtis, J. Pieprzyk, and J. Seruga. An efficient eAuction protocol. In *Proc. 7th Conference on Availability, Reliability and Security (ARES'07)*, pages 417–421. IEEE Computer Society, 2007.
- [7] N. Dong, H. L. Jonker, and J. Pang. Analysis of a receipt-free auction protocol in the applied pi calculus. In *Proc. 7th Workshop on Formal Aspects in Security and Trust (FAST'10)*, volume 6561 of *LNCS*, pages 223–238. Springer-Verlag, 2011.
- [8] J. Dreier, P. Lafourcade, and Y. Lakhnech. Formal verification of e-auction protocols. In *Principles of Security and Trust (POST)*, 2013. To appear.
- [9] A. Fujioka, T. Okamoto, and K. Ohta. A practical secret voting scheme for large scale elections. In J. Seberry and Y. Zheng, editors, *Advances in Cryptology – AUSCRYPT '92*, volume 718 of *LNCS*, pages 244–251. Springer Berlin / Heidelberg, 1992.
- [10] M. Harkavy, J. D. Tygar, and H. Kikuchi. Electronic auctions with private bids. In *Proc. 3rd USENIX Workshop on Electronic Commerce*. Usenix, 1998.
- [11] M. Hirt and K. Sako. Efficient receipt-free voting based on homomorphic encryption. In *Proc. 19th Annual Conference on the Theory and Applications of Cryptographic Techniques: Advances in Cryptology (EUROCRYPT)*, volume 1807 of *LNCS*, pages 539–556. Springer, 2000.
- [12] S. Kremer, M. D. Ryan, and B. Smyth. Election verifiability in electronic voting protocols. In *Proc. 15th European Symposium on Research in Computer Security (ESORICS 2010)*, volume 6345 of *LNCS*, pages 389–404. Springer, 2010.
- [13] R. Küsters, T. Truderung, and A. Vogt. Accountability: definition and relationship to verifiability. In *Proc. 17th Conference on Computer and Communications Security (CCS'10)*, CCS '10, pages 526–535. ACM, 2010.
- [14] H. Lipmaa, N. Asokan, and V. Niemi. Secure vickrey auctions without threshold trust. In *Proc. 6th Conference on Financial Cryptography*, volume 2357 of *LNCS*, pages 87–101. Springer, 2003.
- [15] M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *Proc. 1st Conference on Electronic Commerce*, pages 129–139. ACM Press, 1999.
- [16] K. Omote and A. Miyaji. A practical english auction with one-time registration. In *Proc. 6th Australasian Conference on Information Security and Privacy*, volume 2119 of *LNCS*, pages 221–234. Springer, 2001.
- [17] T. E. Rockoff and M. Groves. Design of an internet-based system for remote dutch auctions. *Internet Research*, 5:10–16, 1995.
- [18] K. Sako and J. Kilian. Receipt-free mix-type voting scheme - A practical solution to the implementation of a voting booth. In *Proc. 14th Conference on the Theory and Applications of Cryptographic Techniques: Advances in Cryptology (EUROCRYPT'95)*, volume 921 of *LNCS*, pages 393–403. Springer, 1995.