

# Covert Security with Public Verifiability: Faster, Leaner, and Simpler

Cheng Hong<sup>1</sup>, Jonathan Katz<sup>2\*</sup>, Vladimir Kolesnikov<sup>3\*\*</sup>,  
Wen-jie Lu<sup>4</sup>, and Xiao Wang<sup>5</sup>

<sup>1</sup> Alibaba, Hangzhou, China  
vince.hc@alibaba-inc.com

<sup>2</sup> University of Maryland, College Park, USA  
jkatz@cs.umd.edu

<sup>3</sup> Georgia Tech, Atlanta, USA  
kolesnikov@gatech.edu

<sup>4</sup> University of Tsukuba, Tsukuba, Japan  
riku@mdl.cs.tsukuba.ac.jp

<sup>5</sup> MIT/BU, Cambridge/Boston, USA  
wangxiao@northwestern.edu

**Abstract.** The notion of *covert security* for secure two-party computation serves as a compromise between the traditional semi-honest and malicious security definitions. Roughly, covert security ensures that cheating behavior is detected by the honest party with reasonable probability (say,  $1/2$ ). It provides more realistic guarantees than semi-honest security with significantly less overhead than is required by malicious security.

The rationale for covert security is that it dissuades cheating by parties that care about their reputation and do not want to risk being caught. But a much stronger disincentive is obtained if the honest party can generate a publicly verifiable *certificate* when cheating is detected. While the corresponding notion of publicly verifiable covert (PVC) security has been explored, existing PVC protocols are complex and less efficient than the best covert protocols, and have impractically large certificates.

We propose a novel PVC protocol that significantly improves on prior work. Our protocol uses only “off-the-shelf” primitives (in particular, it avoids signed oblivious transfer) and, for deterrence factor  $1/2$ , has only 20–40% overhead compared to state-of-the-art *semi-honest* protocols. Our protocol also has, for the first time, *constant-size* certificates of cheating (e.g., 354 bytes long at the 128-bit security level).

As our protocol offers strong security guarantees with low overhead, we suggest that it is the best choice for many practical applications of secure two-party computation.

---

\* Work supported in part by a grant from Alibaba.

\*\* Work supported in part by Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA-0003525.

## 1 Introduction

Secure two-party computation allows two mutually distrusting parties  $P_A$  and  $P_B$  to evaluate a function of their inputs without requiring either party to reveal their input to the other. Traditionally, two security notions have been considered [7]. Protocols with *semi-honest security* can be quite efficient, but only protect against passive attackers who do not deviate from the prescribed protocol. *Malicious security*, in contrast, categorically prevents an attacker from gaining any advantage by deviating from the protocol; unfortunately, despite many advances over the past few years, protocols achieving malicious security are still noticeably less efficient than protocols with semi-honest security.

The notion of *covert security* [3] was proposed as a compromise between semi-honest and malicious security. Roughly, covert security ensures that while a cheating attacker may be successful with some small probability, the attempted cheating will fail *and be detected by the other party* with the remaining probability. The rationale for covert security is that it dissuades cheating by parties that care about their reputation and do not want to risk being caught. Covert security thus provides stronger guarantees than semi-honest security; it can also be achieved with better efficiency than malicious security [3, 9, 6, 17].

Nevertheless, the guarantee of covert security is not fully satisfactory. Covert security only ensures that when cheating is unsuccessful, the honest party detects the fact that cheating took place—but it provides no mechanism for the honest party to *prove* this fact to anyone else (e.g., a judge or the public) and, indeed, existing covert protocols do not provide any such mechanism. Thus, a cheating attacker only risks harming its reputation with one other party; even if the honest party publicly announces that it caught the other party cheating, the cheating party can simply counter that it is being falsely accused.

Motivated by this limitation of covert security, Asharov and Orlandi [2] proposed the stronger notion of *publicly verifiable* covert (PVC) security. As in the covert model, any attempted cheating is detected with some probability; now, however, when cheating is detected the honest party can generate a *publicly verifiable certificate* of that fact. This small change would have a significant impact in practice, as a cheating attacker now risks having its reputation publicly and permanently damaged if it is caught. Alternatively (or additionally), the cheating party can be brought to court and fined for its misbehavior; the parties may even sign a contract in advance that describes the penalties to be paid if either party is caught. Going further, the parties could execute a “smart contract” in advance of the protocol execution that would automatically pay out if a valid certificate of cheating is posted on a blockchain. All these consequences are infeasible in the original covert model and, overall, the PVC model seems to come closer to the original goal of covert security.

Asharov and Orlandi [2] mainly focus on feasibility; although their protocol is implementable, it is not competitive with state-of-the-art semi-honest protocols since, in particular, it requires a stronger variant of oblivious transfer (OT) called *signed OT* and thus is not directly compatible with OT extension. Subsequent work by Kolesnikov and Malozemoff [13] shows various efficiency improvements

to the Asharov-Orlandi protocol, with the primary gain resulting from a new, dedicated protocol for signed-OT extension. (Importantly, signed-OT extension does not follow generically from standard OT extension, and so cannot take advantage of the most-efficient recent constructions of the latter.)

Unfortunately, existing PVC protocols [2, 13] seem not to have attracted much attention; for example, to the best of our knowledge, they have never been implemented. We suggest this is due to a number of considerations:

- **High overhead.** State-of-the-art PVC protocols still incur a significant overhead compared to known semi-honest protocols, and even existing covert protocols. (See Section 6.)
- **Large certificates.** Existing PVC protocols have certificates of size at least  $\kappa \cdot |\mathcal{C}|$  bits, where  $\kappa$  is the (computational) security parameter and  $|\mathcal{C}|$  is the circuit size.<sup>1</sup> Certificates this large are prohibitively expensive to propagate and are incompatible with some of the applications mentioned above (e.g., posting a certificate on a blockchain).
- **Complexity.** Existing PVC protocols rely on signed OT, a non-standard primitive that is less efficient than standard OT, is not available in existing secure-computation libraries, and is somewhat complicated to realize (especially for signed-OT extension).

## 1.1 Our Contributions

In this work we put forward a new PVC protocol in the random oracle model that addresses the issues mentioned above. Specifically:

- **Low overhead.** We improve on the efficiency of prior work by roughly a factor of  $3\times$  for deterrence factor  $1/2$ , and even more for larger deterrence. (An exact comparison depends on a number of factors; we refer to Section 6 for a detailed discussion.) Strikingly, our PVC protocol (with deterrence factor  $1/2$ ) incurs *only 20–40% overhead compared to state-of-the-art semi-honest protocols based on garbled circuits*.
- **Small certificates.** We achieve, for the first time, *constant-size* certificates (i.e., independent of the circuit size or the lengths of the parties’ inputs). Concretely, our certificates are small: at the 128-bit security level, they are only 354 bytes long.
- **Simplicity.** Our protocol avoids entirely the need for signed OT, and relies only on standard building blocks such as (standard) OT and circuit garbling. We also dispense with the XOR-tree technique for preventing selective-failure attacks; this allows us to avoid increasing the number of effective OT inputs. This reduction in complexity allowed us to produce a simple and efficient (and, to our knowledge, the first) implementation of a PVC protocol.

---

<sup>1</sup> We observe that the certificate size in [13] can be improved to  $O(\kappa \cdot n)$  bits (where  $n$  is the parties’ input lengths) by carefully applying ideas from the literature. In many cases, this is still unacceptably large.

**Overview of the paper.** In Section 2 we provide an overview of prior PVC protocols and explain the intuition behind the construction of our protocol. After some background in Section 3, we present the description of our protocol in Section 4 and prove security in Section 5. Section 6 gives an experimental evaluation of our protocol and a comparison to prior work.

## 2 Technical Overview

We begin by providing an overview of the approach taken in prior work designing PVC protocols. Then we discuss the intuition behind our improved protocol.

### 2.1 Overview of Prior Work

At a high level, both previous works constructing PVC protocols [2, 13] rely on the standard cut-and-choose paradigm [18] using a small number of garbled circuits, with some additional complications to achieve public verifiability. Both works rely crucially on a primitive called *signed OT*; this is a functionality similar to OT but where the receiver additionally learns the sender’s signatures on all the values it obtains. Roughly, prior protocols proceed as follows:

1. Let  $\lambda$  be a parameter that determines the deterrence factor (i.e., the probability of detecting misbehavior).  $P_A$  picks random seeds  $\{\text{seed}_j\}_{j=1}^\lambda$  and  $P_B$  chooses a random index  $\hat{j} \in \{1, \dots, \lambda\}$  that will serve as the “evaluation index” while the  $j \neq \hat{j}$  will be “check indices.” The parties run signed OT using these inputs, which allows  $P_B$  to learn  $\{\text{seed}_j\}_{j \neq \hat{j}}$  along with signatures of  $P_A$  on all those values.
2.  $P_A$  generates  $\lambda$  garbled circuits, and then sends signed commitments to those garbled circuits (along with the input-wire labels corresponding to  $P_A$ ’s input wires). Importantly,  $\text{seed}_j$  is used to derive the (pseudo)randomness for the  $j$ th garbling as well as the  $j$ th commitment.  
The parties also use signed OT so that  $P_B$  can obtain the input-wire labels for its inputs across all the circuits.
3. For all  $j \neq \hat{j}$ , party  $P_B$  checks that the commitment to the  $j$ th garbled circuit is computed correctly based on  $\text{seed}_j$  and that the input-wire labels it received are correct; if this is not the case, then  $P_B$  can generate a certificate of cheating that consists of the inconsistent values plus their signatures.
4. Assuming no cheating was detected,  $P_B$  reveals  $\hat{j}$  to  $P_A$ , who then sends the  $\hat{j}$ th garbled circuit and the input-wire labels corresponding to its own inputs for that circuit.  $P_B$  can then evaluate the garbled circuit as usual.

Informally, we refer to the  $j$ th garbled circuit and commitment as the  *$j$ th instance of the protocol*. If  $P_A$  cheats in the  $j$ th instance of the protocol, then it is caught with probability at least  $1 - \frac{1}{\lambda}$  (i.e., if  $j$  is a check index). Moreover, if  $P_A$  is caught, then  $P_B$  has a signed seed (which defines what  $P_A$  was supposed to do in the  $j$ th instance) and also a signed commitment to an incorrect garbled circuit or

incorrect input-wire labels. These values allow  $P_B$  to generate a publicly verifiable certificate that  $P_A$  cheated.

As described, the protocol still allows  $P_A$  to carry out a selective-failure attack when transferring garbled labels for  $P_B$ 's input wires. Specifically, it may happen that a malicious  $P_A$  corrupts a single input-wire label (used as input to the OT protocol) for the  $j$ th garbled circuit—say, the label corresponding to a ‘1’ input on some wire. If  $P_B$  aborts, then  $P_A$  learns that  $P_B$ 's input on that wire was equal to 1. Such selective-failure attacks can be prevented using the XOR-tree approach [18].<sup>2</sup> This approach introduces significant overhead because it increases the number of effective inputs, which in turn requires additional signed OTs. The analysis in prior work [3, 2, 13] shows that to achieve deterrence factor (i.e., probability of being caught cheating)  $1/2$ , a replication factor of  $\lambda = 3$  with  $\nu = 3$  is needed. More generally, the deterrence factor as a function of  $\lambda$  and the XOR-tree expansion factor  $\nu$  is  $(1 - \frac{1}{\lambda}) \cdot (1 - 2^{-\nu+1})$ .

**Practical performance.** Several aspects of the above protocol are relatively inefficient. First, the dependence of the deterrence factor on the replication factor  $\lambda$  is not optimal due to the XOR tree, e.g., to achieve deterrence factor  $1/2$  at least  $\lambda = 3$  garbled circuits are needed (unless  $\nu$  is impractically large); the issue becomes even more significant when a larger deterrence factor is desired. In addition, the XOR-tree approach used in prior work increases the effective input length by at least a factor of 3, which necessitates  $3\times$  more signed OTs; recall these are relatively expensive since signed-OT extension is. Finally, prior protocols have large certificates. This seems inherent in the more efficient protocol of [13] due to the way they do signed-OT extension. (Avoiding signed-OT extension would result in a much less efficient protocol overall.)

## 2.2 Our Solution

The reliance of prior protocols on signed OT and their approach to preventing selective-failure attacks affect both their efficiency as well as the size of their certificates. We address both these issues in the protocol we design.

As in prior work, we use the cut-and-choose approach and have  $P_B$  evaluate one garbled circuit while checking the rest, and we realize this by having  $P_A$  choose seeds for each of  $\lambda$  executions and then allowing  $P_B$  to obviously learn all-but-one of those seeds. One key difference in our protocol is that we utilize the seeds chosen by  $P_A$  not only to “derandomize” the garbled-circuit generation and commitments, but *also to derandomize the entire remainder of  $P_A$ 's execution*, and in particular its execution of the OT protocol used to transfer  $P_B$ 's input-wire labels to  $P_B$ . This means that after  $P_B$  obviously learns all-but-one of the seeds of  $P_A$ , the rest of  $P_A$ 's execution is entirely deterministic; thus,  $P_B$  can verify correct execution of  $P_A$  during the entire rest of the protocol for all-but-one of the seeds. Not only does this eliminate the need for signed OT for the input-wire

<sup>2</sup> For reasonable values of the parameters, the XOR-tree approach will be more efficient than a coding-theoretic approach [18].

labels, but it also defends against the selective-failure attack described earlier *without* the need to increase the effective input length at all.

As described, the above allows  $P_B$  to *detect* cheating by  $P_A$  but does not yet achieve public verifiability. For this, we additionally require  $P_A$  to sign its protocol messages; if  $P_A$  cheats,  $P_B$  can generate a certificate of cheating from the seed and the corresponding signed inconsistent transcript.

Thus far we have focused on the case where  $P_A$  is malicious. We must also consider the case of a malicious  $P_B$  attempting to frame an honest  $P_A$ . We address this by also having  $P_B$  commit in advance to *its* randomness<sup>3</sup> for each of the  $\lambda$  protocol instances. The resulting commitments will be included in  $P_A$ 's signature, and will ensure that a certificate will be rejected if it corresponds to an instance in which  $P_B$  deviated from the protocol.

Having  $P_B$  commit to its randomness also allows us to avoid the need for signed OT in the first step, when  $P_B$  learns all-but-one of  $P_A$ 's seeds. This is because those seeds can be reconstructed from  $P_B$ 's view of the protocol, i.e., from the transcript of the (standard) OT protocol used to transfer those seeds plus  $P_B$ 's randomness. Having  $P_A$  sign the transcripts of those OT executions serves as publicly verifiable evidence of the seeds used by  $P_A$ .

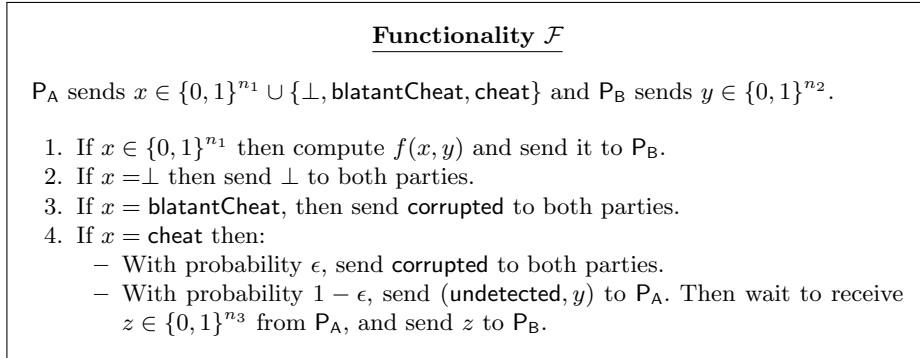
We refer to Section 4 for further intuition behind our protocol, as well as its formal specification.

### 3 Covert Security with Public Verifiability

Before defining the notion of PVC security, we review the (plain) covert model [3] it extends. We focus on the strongest formulation of covert security, namely the *strong explicit cheat* formulation. This notion is formalized via an ideal functionality that explicitly allows an adversary to specify an attempt at cheating; in that case, the ideal functionality allows the attacker to successfully cheat with probability  $1 - \epsilon$ , but the attacker is caught by the other party with probability  $\epsilon$ ; see Figure 1. (As in [2], we also allow an attacker to “blatantly cheat,” which guarantees that it will be caught.) For simplicity, we adapt the functionality such that only  $P_A$  has this option (since this is what is achieved by our protocol). For conciseness, we refer to a protocol realizing this functionality (against malicious adversaries) as having *covert security with deterrence*  $\epsilon$ .

The PVC model extends the above to consider a setting wherein, before execution of the protocol,  $P_A$  has generated keys  $(pk, sk)$  for a digital-signature scheme, with the public key  $pk$  known to  $P_B$ . We *do not* require that  $(pk, sk)$  is honestly generated, or that  $P_A$  gives any proof of knowledge of the secret key  $sk$  corresponding to the public key  $pk$ . In addition, the protocol is augmented with two additional algorithms, **Blame** and **Judge**. The **Blame** algorithm is run by  $P_B$  when it outputs **corrupted**. This algorithm takes as input  $P_B$ 's view of the protocol execution thus far, and outputs a certificate **cert** which is then sent

<sup>3</sup> As an optimization, we have  $P_B$  commit to seeds, just like  $P_A$ , and then use those seeds to generate the (pseudo)randomness to use in each instance. (This optimization is critical for realizing constant-size certificates.)



**Fig. 1.** Functionality  $\mathcal{F}$  for covert security with deterrence  $\epsilon$  for two-party computation of a function  $f$ .

to  $P_A$ . The `Judge` algorithm takes as input  $P_A$ 's public key  $pk$ , (a description of) the circuit  $\mathcal{C}$  being evaluated, and a certificate `cert`, and outputs 0 or 1.

A protocol  $\Pi$  along with algorithms `Blame`, `Judge` is said to be *publicly verifiable covert with deterrence  $\epsilon$*  for computing a circuit  $\mathcal{C}$  if the following hold:

**Covert security:** The protocol  $\Pi$  has covert security with deterrence  $\epsilon$ . (Since the protocol includes the step of possibly sending `cert` to  $P_A$  if  $P_B$  outputs **corrupted**, this ensures that `cert` itself does not violate privacy of  $P_B$ .)

**Public verifiability:** If the honest  $P_B$  outputs `cert` in an execution of the protocol, then we know  $\text{Judge}(pk, \mathcal{C}, \text{cert}) = 1$ , except with negligible probability.

**Defamation freeness:** If  $P_A$  is honest, then the probability that a malicious  $P_B$  generates a certificate `cert` for which  $\text{Judge}(pk, \mathcal{C}, \text{cert}) = 1$  is negligible.<sup>4</sup>

As in prior work on the PVC model, we assume the `Judge` algorithm learns the circuit  $\mathcal{C}$  through some “out-of-band” mechanism; in particular, we do not include  $\mathcal{C}$  as part of the certificate. In some applications (such as the smart-contract example), it may indeed be the case that the party running the `Judge` algorithm is aware of the circuit being computed in advance. When this is not the case, a description of  $\mathcal{C}$  must be included as part of the certificate. However, we stress that the description of a circuit may be much shorter than the full circuit; for example, specifying a circuit for computing the Hamming distance between two  $10^6$ -bit vectors requires only a few lines of high-level code in modern secure-computation platforms even though the circuit itself may have millions of gates. Alternately, there may be a small set of commonly used “reference circuits” that can be identified by ID number rather than by their complete wiring diagram.

<sup>4</sup> Note that defamation freeness implies that the protocol is also non-halting detection accurate [3].

## 4 Our PVC Protocol


### 4.1 Preliminaries

We let  $[n] = \{1, \dots, n\}$ . We use  $\kappa$  for the (computational) security parameter, but for compactness in the protocol description we let  $\kappa$  be an implicit input to our algorithms. For a boolean string  $y$ , we let  $y[i]$  denote the  $i$ th bit of  $y$ .

We let  $\text{Com}$  denote a commitment scheme. We assume for simplicity that it is non-interactive, but this restriction can easily be removed. The decommitment  $\text{decom}$  is simply the random coins used during commitment.  $H$  is a hash function with  $2\kappa$ -bit output length.

We say a party “uses randomness derived from  $\text{seed}$ ” to mean that the party uses a pseudorandom function (with  $\text{seed}$  as the key) in CTR mode to obtain sufficiently many pseudorandom values that it then uses as its random coins. If  $m_1, m_2, \dots$  is a transcript of an execution of a two-party protocol (where the parties alternate sending the messages), the *transcript hash* of the execution is defined to be  $\mathcal{H} = (H(m_1), H(m_2), \dots)$ .

We let  $\Pi_{\text{OT}}$  be an OT protocol realizing a parallel version of the OT functionality, as in Figure 2.



**Fig. 2.** Functionality  $\mathcal{F}_{\text{OT}}$  for parallel oblivious transfer.

**Garbling.** Our protocol relies on a (circuit) *garbling scheme*. For our purposes, a garbling scheme is defined by algorithms  $(\text{Gb}, \text{Eval})$  having the following syntax:

- $\text{Gb}$  takes as input the security parameter  $1^\kappa$  and a circuit  $\mathcal{C}$  with  $n = n_1 + n_2$  input wires and  $n_3$  output wires. It outputs *input-wire labels*  $\{X_{i,0}, X_{i,1}\}_{i=1}^n$ , a *garbled circuit*  $\text{GC}$ , and *output-wire labels*  $\{Z_{i,0}, Z_{i,1}\}_{i=1}^{n_3}$ .
- $\text{Eval}$  is a deterministic algorithm that takes as input a set of input-wire labels  $\{X_i\}_{i=1}^n$  and a garbled circuit  $\text{GC}$ . It outputs a set of output-wire labels  $\{Z_i\}_{i=1}^{n_3}$ .

Correctness is defined as follows: For any circuit  $\mathcal{C}$  as above and any input  $w \in \{0, 1\}^n$ , consider the experiment in which we first run  $(\{X_{i,0}, X_{i,1}\}_{i=1}^n, \text{GC}, \{Z_{i,0}, Z_{i,1}\}_{i=1}^{n_3}) \leftarrow \text{Gb}(1^\kappa, \mathcal{C})$  followed by  $\{Z_i\} := \text{Eval}(\{X_{i,w[i]}\}, \text{GC})$ . Then, except with negligible probability, it holds that  $Z_i = Z_{i,y[i]}$  and  $Z_i \neq Z_{i,1-y[i]}$  for all  $i$ , where  $y = \mathcal{C}(w)$ .



A garbling scheme can be used by (honest) parties  $P_A$  and  $P_B$  to compute  $\mathcal{C}$  in the following way: first,  $P_A$  computes  $(\{X_{i,0}, X_{i,1}\}_{i=1}^n, \text{GC}, \{Z_{i,0}, Z_{i,1}\}_{i=1}^{n_3}) \leftarrow \text{Gb}(1^\kappa, \mathcal{C})$  and sends  $\text{GC}, \{Z_{i,0}, Z_{i,1}\}_{i=1}^{n_3}$  to  $P_B$ . Next,  $P_B$  learns the input-wire labels  $\{X_{i,w[i]}\}$  corresponding to some input  $w$ . (In a secure-computation protocol,  $P_A$  would send  $P_B$  the input-wire labels corresponding to its own portion of the input, while the parties would use OT to enable  $P_B$  to learn the input-wire labels corresponding to  $P_B$ 's portion of the input.) Then  $P_B$  computes  $\{Z_i\} := \text{Eval}(\{X_{i,w[i]}\}, \text{GC})$ . Finally,  $P_B$  sets  $y[i]$ , for all  $i$ , to be the (unique) bit for which  $Z_i = Z_{i,y[i]}$ ; the output is  $y$ .

We assume the garbling scheme satisfies the standard security definition [10, 15]. That is, we assume there is a simulator  $\mathcal{S}_{\text{Gb}}$  such that for all  $\mathcal{C}, w$ , the distribution  $\{\mathcal{S}_{\text{Gb}}(1^\kappa, \mathcal{C}, \mathcal{C}(w))\}$  is computationally indistinguishable from

$$\left\{ (\{X_{i,0}, X_{i,1}\}_{i=1}^n, \text{GC}, \{Z_{i,0}, Z_{i,1}\}_{i=1}^{n_3}) \leftarrow \text{Gb}(1^\kappa, \mathcal{C}) : (\{X_{i,w[i]}\}, \text{GC}, \{Z_{i,0}, Z_{i,1}\}_{i=1}^{n_3}) \right\}.$$

As this is the “minimal” security notion for garbling, it is satisfied by garbling schemes including all state-of-the-art optimizations [14, 4, 20].

## 4.2 Our Scheme

We give a high-level description of our protocol below; a formal definition of the protocol is provided in Figure 3. The Blame algorithm is included as part of the protocol description (cf. Step 6) for simplicity. The Judge algorithm is specified in Figure 5.

We use a signature scheme  $(\text{Gen}, \text{Sign}, \text{Vrfy})$ . Before executing the protocol,  $P_A$  runs  $\text{Gen}$  to obtain public key  $pk$  and private key  $sk$ ; we assume that  $P_B$  knows  $pk$  before running the protocol. As noted earlier, if  $P_A$  is malicious then it may choose  $pk$  arbitrarily.

The main idea of the protocol is to run  $\lambda$  parallel instances of a “basic” garbled-circuit protocol that is secure against a semi-honest  $P_A$  and a malicious  $P_B$ . Of these instances,  $\lambda - 1$  will be checked by  $P_B$ , while a random one (the  $\hat{j}$ th) will be evaluated by  $P_B$  to learn its output. To give  $P_B$  the ability to verify honest behavior in the check instances, we make all the executions deterministic by having  $P_A$  use (pseudo)randomness derived from corresponding seeds  $\{\text{seed}_j^A\}_{j \in [\lambda]}$ . That is,  $P_A$  will uniformly sample each seed  $\text{seed}_j^A$  and use it to generate (pseudo)randomness for its  $j$ th instance. Then  $P_A$  and  $P_B$  run an OT protocol  $\Pi_{\text{OT}}$  (with malicious security) that allows  $P_B$  to learn  $\lambda - 1$  of those seeds. Since  $P_A$ 's behavior in those  $\lambda - 1$  instances is completely determined by  $P_B$ 's messages and those seeds, it is possible for  $P_B$  to check  $P_A$ 's behavior in those instances.

The above idea allows  $P_B$  to catch a cheating  $P_A$ , but not to generate a publicly verifiable certificate that  $P_A$  has cheated. To add this feature, we have  $P_A$  sign the transcripts of each instance, including the transcript of the execution of the OT protocol by which  $P_B$  learned the corresponding seed. If  $P_A$  cheats in, say, the  $j$ th instance ( $j \neq \hat{j}$ ) and is caught, then  $P_B$  can output a certificate

**Protocol  $\Pi_{\text{pvc}}$**

**Private inputs:**  $P_A$  has input  $x \in \{0, 1\}^{n_1}$  and keys  $(pk, sk)$  for a signature scheme.  $P_B$  has input  $y \in \{0, 1\}^{n_2}$  and knows  $pk$ .

**Public inputs:** Both parties also agree on a circuit  $\mathcal{C}$  and parameters  $\kappa, \lambda$ .

**Protocol:**

1.  $P_B$  chooses uniform  $\kappa$ -bit strings  $\{\text{seed}_j^B\}_{j \in [\lambda]}$ , sets  $h_j \leftarrow \text{Com}(\text{seed}_j^B)$  for all  $j$ , and sends  $\{h_j\}_{j \in [\lambda]}$  to  $P_A$ .
2.  $P_A$  chooses uniform  $\kappa$ -bit strings  $\{\text{seed}_j^A, \text{witness}_j\}_{j \in [\lambda]}$ , while  $P_B$  chooses uniform  $\hat{j} \in [\lambda]$  and sets  $b_{\hat{j}} := 1$  and  $b_j := 0$  for  $j \neq \hat{j}$ .  
 $P_A$  and  $P_B$  run  $\lambda$  executions of  $\Pi_{\text{OT}}$ , where in the  $j$ th execution  $P_A$  uses  $(\text{seed}_j^A, \text{witness}_j)$  as input, and  $P_B$  uses  $b_j$  as input and randomness derived from  $\text{seed}_j^B$ . Upon completion,  $P_B$  obtains  $\{\text{seed}_j^A\}_{j \neq \hat{j}}$  and  $\text{witness}_j$ . Let  $\text{trans}_j$  be the transcript of the  $j$ th execution of  $\Pi_{\text{OT}}$ .
3. For each  $j \in [\lambda]$ ,  $P_A$  garbles  $\mathcal{C}$  using randomness derived from  $\text{seed}_j^A$ . Denote the  $j$ th garbled circuit by  $\text{GC}_j$ , the input-wire labels of  $P_A$  by  $\{A_{j,i,b}\}_{i \in [n_1], b \in \{0,1\}}$ , the input-wire labels of  $P_B$  by  $\{B_{j,i,b}\}_{i \in [n_2], b \in \{0,1\}}$ , and the output-wire labels by  $\{Z_{j,i,b}\}_{i \in [n_3], b \in \{0,1\}}$ .  
 $P_A$  and  $P_B$  then run  $\lambda$  executions of  $\Pi_{\text{OT}}$ , where in the  $j$ th execution  $P_A$  uses  $\{(B_{j,i,0}, B_{j,i,1})\}_{i=1}^{n_2}$  as input, and  $P_B$  uses  $y$  as input if  $j = \hat{j}$  and  $0^{n_2}$  otherwise. The parties use  $\text{seed}_j^A$  and  $\text{seed}_j^B$ , respectively, to derive all their randomness in the  $j$ th execution. In this way,  $P_B$  obtains  $\{B_{j,i,y[i]}\}_{i \in [n_2]}$ . We let  $\mathcal{H}_j$  denote the transcript hash for the  $j$ th execution of  $\Pi_{\text{OT}}$ .
4.  $P_A$  computes commitments  $h_{j,i,b}^A \leftarrow \text{Com}(A_{j,i,b})$  for all  $j, i, b$ , and then computes the commitments  $c_j \leftarrow \text{Com}(\text{GC}_j, \{h_{j,i,b}^A\}_{i \in [n_1], b \in \{0,1\}}, \{Z_{j,i,b}\}_{i \in [n_3], b \in \{0,1\}})$  for all  $j$ , where each pair  $(h_{j,i,0}^A, h_{j,i,1}^A)$  is randomly permuted. All randomness in the  $j$ th instance is derived from  $\text{seed}_j^A$ . Finally,  $P_A$  sends  $\{c_j\}_{j \in [\lambda]}$  to  $P_B$ .
5. For each  $j \in [\lambda]$ ,  $P_A$  computes  $\sigma_j \leftarrow \text{Sign}_{sk}(\mathcal{C}, j, h_j, \text{trans}_j, \mathcal{H}_j, c_j)$  and sends  $\sigma_j$  to  $P_B$ . Then  $P_B$  checks that  $\sigma_j$  is a valid signature for all  $j$ , and aborts with output  $\perp$  if not.
6. For each  $j \neq \hat{j}$ ,  $P_B$  uses  $\text{seed}_j^A$  and the messages it sent to simulate  $P_A$ 's computation in steps 3 and 4, and in particular computes  $\hat{\mathcal{H}}_j, \hat{c}_j$ . It then checks that  $(\hat{\mathcal{H}}_j, \hat{c}_j) = (\mathcal{H}_j, c_j)$ . If the check fails for some  $j \neq \hat{j}$ , then  $P_B$  chooses a uniform such  $\hat{j}$ , outputs corrupted, sends  $\text{cert} := (j, \text{trans}_j, \mathcal{H}_j, c_j, \sigma_j, \text{seed}_j^B, \text{decom}_j^B)$  to  $P_A$ , and halts.

**Fig. 3.** Full description of our PVC protocol (part I).

that includes  $P_B$ 's view (including its randomness) in the execution of the  $j$ th OT protocol (from which  $\text{seed}_j^A$  can be recomputed) and the transcript of the  $j$ th instance, along with  $P_A$ 's signature on the transcripts. Note that, given the randomness of both  $P_A$  and  $P_B$ , the entire transcript of the instance can be recomputed and anyone can then check whether it is consistent with  $\text{seed}_j^A$ . We remark that nothing about  $P_B$ 's inputs is revealed by a certificate since  $P_B$  uses a dummy input in all the check instances.

**Protocol  $\Pi_{\text{pvc}}$**

7.  $P_B$  sends  $(\hat{j}, \{\text{seed}_j^A\}_{j \neq \hat{j}}, \text{witness}_j)$  to  $P_A$ , who checks that  $\{\text{seed}_j^A\}_{j \neq \hat{j}}, \text{witness}_j$  are all correct and aborts if not.
8.  $P_A$  sends  $\text{GC}_{\hat{j}}, \{A_{\hat{j},i,x[i]}\}_{i \in [n_1]}, \{h_{\hat{j},i,b}^A\}_{i \in [n_1], b \in \{0,1\}}$  (in the same permuted order as before), and  $\{Z_{\hat{j},i,b}\}_{i \in [n_3], b \in \{0,1\}}$  to  $P_B$ , along with decommitments  $\text{decom}_{\hat{j}}$  and  $\{\text{decom}_{\hat{j},i,x[i]}^A\}$ . If  $\text{Com}(\text{GC}_{\hat{j}}, \{h_{\hat{j},i,b}^A\}, \{Z_{\hat{j},i,b}\}; \text{decom}_{\hat{j}}) \neq c_{\hat{j}}$  or  $\text{Com}(A_{\hat{j},i,x[i]}; \text{decom}_{\hat{j},i,x[i]}^A) \notin \{h_{\hat{j},i,b}^A\}_{b \in \{0,1\}}$  for some  $i$ , then  $P_B$  aborts with output  $\perp$ .  
 Otherwise,  $P_B$  evaluates  $\text{GC}_{\hat{j}}$  using  $\{A_{\hat{j},i,x[i]}\}_{i \in [n_1]}$  and  $\{B_{\hat{j},i,y[i]}\}_{i \in [n_2]}$  to obtain output-wire labels  $\{Z_i\}_{i \in [n_3]}$ . For each  $i \in [n_3]$ , if  $Z_i = Z_{\hat{j},i,0}$ , set  $z[i] := 0$ ; if  $Z_i = Z_{\hat{j},i,1}$ , set  $z[i] := 1$ . (If  $Z_i \notin \{Z_{\hat{j},i,0}, Z_{\hat{j},i,1}\}$  for some  $i$ , then abort with output  $\perp$ .) Output  $z$ .

**Fig. 4.** Full description of our PVC protocol (part II).

There still remains the potential issue of *defamation*. Indeed, an honest  $P_A$ 's messages might be deemed inconsistent if  $P_B$  includes in the certificate fake messages different from those sent by  $P_B$  in the real execution. We prevent this by having  $P_B$  commit to its randomness for each instance at the beginning of the protocol, and having  $P_A$  sign those commitments. Consistency of  $P_B$ 's randomness and the given transcript can then be checked as part of verification of the certificate.

As described, the above would result in a certificate that is linear in the length of  $P_B$ 's inputs, since there are that many OT executions (in each instance) for which  $P_B$  must generate randomness. We compress this to a constant-size certificate by having  $P_B$  also generate its (pseudo)randomness from a short seed.

The above description conveys the main ideas of the protocol, though various other modifications are needed for the proof of security. We refer the reader to Figures 3 and 5 for the details.

### 4.3 Optimizations

Our main protocol is already quite efficient, but we briefly discuss some additional optimizations that can be applied.

**Commitments in the random-oracle model.** When standard garbling schemes are used, all the values committed during the course of the protocol have high entropy; thus, commitment to a string  $r$  can be done by simply computing  $H(r)$  (if  $H$  is modeled as a random oracle) and decommitment requires only sending  $r$ .

**Using correlated oblivious transfer.** One optimization introduced by Asharov et al. [1] is using correlated OT for transferring  $P_B$ 's input-wire labels when garbling is done using the free-XOR approach [14]. This optimization is compatible with our protocol in a straightforward manner.

**Algorithm Judge**

**Inputs:** A public key  $pk$ , a circuit  $\mathcal{C}$ , and a certificate  $\text{cert}$ .

1. Parse  $\text{cert}$  as  $(j, \text{trans}_j, \mathcal{H}_j, \mathbf{c}_j, \sigma_j, \text{seed}_j^B, \text{decom}_j^B)$ . Compute  $h_j := \text{Com}(\text{seed}_j^B; \text{decom}_j^B)$ .
2. If  $\text{Vrfy}_{pk}(\mathcal{C}, j, h_j, \text{trans}_j, \mathcal{H}_j, \mathbf{c}_j, \sigma_j) = 0$ , output 0.
3. Simulate an execution of  $\Pi_{\text{OT}}$  by  $\mathsf{P}_B$ , where  $\mathsf{P}_B$ 's input is 0, its randomness is derived from  $\text{seed}_j^B$ , and  $\mathsf{P}_A$ 's messages are those included in  $\text{trans}_j$ . Check that all of  $\mathsf{P}_B$ 's messages generated in this simulation are consistent with  $\text{trans}_j$ ; terminate with output 0 if not. Otherwise, let  $\text{seed}_j^A$  denote the output of  $\mathsf{P}_B$  from the simulated execution of  $\Pi_{\text{OT}}$ .
4. Use  $\text{seed}_j^A$  and  $\text{seed}_j^B$  to simulate an honest execution of steps 3 and 4 of the protocol, and in particular compute  $\hat{\mathcal{H}}_j, \hat{\mathbf{c}}_j$ .
5. Do:
  - (a) If  $(\hat{\mathcal{H}}_j, \hat{\mathbf{c}}_j) = (\mathcal{H}_j, \mathbf{c}_j)$  then output 0.
  - (b) If  $\hat{\mathbf{c}}_j \neq \mathbf{c}_j$  then output 1.
  - (c) Find the first message for which  $\hat{\mathcal{H}}_j \neq \mathcal{H}_j$ . If this corresponds to a message sent by  $\mathsf{P}_A$ , output 1; otherwise, output 0.

**Fig. 5.** The Judge algorithm.

**Avoiding committing to the input-wire labels.** In our protocol, we have  $\mathsf{P}_A$  commit to its input-wire labels (along with the rest of the garbled circuit). This is done to prevent  $\mathsf{P}_A$  from sending incorrect input-wire labels in the final step. We observe that this is unnecessary if the garbling scheme has the additional property that it is infeasible to generate a garbled circuit along with incorrect input-wire labels that result in a valid output when evaluated. (We omit a formal definition.) Many standard garbling schemes have this property.

## 5 Proof of Security

The remainder of this section is devoted to a proof of the following result:

**Theorem 1.** *Assume  $\text{Com}$  is computationally hiding/binding,  $H$  is collision-resistant, the garbling scheme is secure,  $\Pi_{\text{OT}}$  realizes  $\mathcal{F}_{\text{OT}}$ , and the signature scheme is existentially unforgeable under a chosen-message attack. Then protocol  $\Pi_{\text{pvc}}$  along with Blame as in step 6 and Judge as in Figure 5 is publicly verifiable covert with deterrence  $\epsilon = 1 - \frac{1}{\lambda}$ .*

*Proof.* We separately prove covert security with  $\epsilon$ -deterrence (handling the cases where either  $\mathsf{P}_A$  or  $\mathsf{P}_B$  is corrupted), public verifiability, and defamation freeness.

### Covert Security—Malicious $P_A$

Let  $\mathcal{A}$  be an adversary corrupting  $P_A$ . We construct the following simulator  $\mathcal{S}$  that holds  $pk$  and runs  $\mathcal{A}$  as a subroutine, while playing the role of  $P_A$  in the ideal world interacting with  $\mathcal{F}$ :

1. Choose uniform  $\kappa$ -bit strings  $\{\text{seed}_j^B\}_{j \in [\lambda]}$ , set  $h_j \leftarrow \text{Com}(\text{seed}_j^B)$  for all  $j$ , and send  $\{h_j\}_{j \in [\lambda]}$  to  $\mathcal{A}$ .
2. For all  $j \in [\lambda]$ , run  $\Pi_{OT}$  with  $\mathcal{A}$ , using input 0 and randomness derived from  $\text{seed}_j^B$ . In this way,  $\mathcal{S}$  obtains  $\{\text{seed}_j^A\}_{j \in [\lambda]}$ . Let  $\text{trans}_j$  denote the transcript of the  $j$ th execution.
3. For  $j \in [\lambda]$ , run an execution of  $\Pi_{OT}$  with  $\mathcal{A}$ , using input  $0^{n_2}$  and randomness derived from  $\text{seed}_j^B$ . Let  $\mathcal{H}_j$  denote the transcript hash of the  $j$ th execution.
4. Receive  $\{c_j\}_{j \in [\lambda]}$  from  $\mathcal{A}$ .
5. Receive  $\{\sigma_j\}$  from  $\mathcal{A}$ . If any of the signatures are invalid, send  $\perp$  to  $\mathcal{F}$  and halt.
6. For all  $j \in [\lambda]$ , use  $\text{seed}_j^A$  and the messages sent previously to simulate the computation of an honest  $P_A$  in steps 3 and 4, and in particular compute  $\hat{\mathcal{H}}_j, \hat{c}_j$ . Let  $J$  be the set of indices for which  $(\hat{\mathcal{H}}_j, \hat{c}_j) \neq (\mathcal{H}_j, c_j)$ .  
There are now three cases:
  - If  $|J| \geq 2$  then send `blatantCheat` to  $\mathcal{F}$ , send `cert` :=  $(j, \text{trans}_j, \mathcal{H}_j, c_j, \sigma_j, \text{seed}_j^B, \text{decom}_j^B)$  to  $\mathcal{A}$  (for uniform  $j \in J$ ), and halt.
  - If  $|J| = 1$  then send `cheat` to  $\mathcal{F}$ . If  $\mathcal{F}$  returns `corrupted` then set `caught` := `true`; if  $\mathcal{F}$  returns `(undetected, y)`, set `caught` := `false`. In either case, continue below.
  - If  $|J| = 0$  then set `caught` :=  $\perp$  and continue below.
- 0'. Rewind  $\mathcal{A}$  and run steps 1'–6' below until<sup>5</sup>  $|J'| = |J|$  and `caught'` = `caught`.
  - 1'. Choose uniform  $\hat{j} \in [\lambda]$ . For  $j \neq \hat{j}$ , choose uniform  $\kappa$ -bit strings  $\{\text{seed}_j^B\}$  and set  $h_j \leftarrow \text{Com}(\text{seed}_j^B)$ . Set  $h_{\hat{j}} \leftarrow \text{Com}(0^\kappa)$ . Send  $\{h_j\}_{j \in [\lambda]}$  to  $\mathcal{A}$ .
  - 2'. For all  $j \neq \hat{j}$ , run  $\Pi_{OT}$  with  $\mathcal{A}$ , using input 0 and randomness derived from  $\text{seed}_j^B$ . In this way,  $\mathcal{S}$  obtains  $\{\text{seed}_j^A\}_{j \neq \hat{j}}$ . For the  $\hat{j}$ th execution, use the simulator  $\mathcal{S}_{OT}$  for protocol  $\Pi_{OT}$ , thus extracting both  $\text{seed}_{\hat{j}}^A$  and  $\text{witness}_{\hat{j}}$ . Let  $\text{trans}_{\hat{j}}$  denote the transcript of the  $\hat{j}$ th execution.
  - 3'. For all  $j \neq \hat{j}$ , run  $\Pi_{OT}$  with  $\mathcal{A}$ , using input  $0^{n_2}$  and randomness derived from  $\text{seed}_j^B$ . For  $j = \hat{j}$ , use the simulator  $\mathcal{S}_{OT}$  for protocol  $\Pi_{OT}$ , thus extracting  $\{B_{\hat{j}, i, b}\}_{i \in [n_2], b \in \{0,1\}}$ . Let  $\mathcal{H}_j$  denote the transcript hash of the  $j$ th execution.
  - 4'. Receive  $\{c_j\}_{j \in [\lambda]}$  from  $\mathcal{A}$ .
  - 5'. Receive  $\{\sigma_j\}$  from  $\mathcal{A}$ . If any of the signatures are invalid, then return to step 1'.
  - 6'. For all  $j \in [\lambda]$ , use  $\text{seed}_j^A$  and the messages sent previously to simulate the computation of an honest  $P_A$  in steps 3' and 4', and in particular compute  $\hat{\mathcal{H}}_j, \hat{c}_j$ . Let  $J'$  be the set of indices for which  $(\hat{\mathcal{H}}_j, \hat{c}_j) \neq (\mathcal{H}_j, c_j)$ . If  $|J'| = 1$  and  $\hat{j} \notin J'$  then set `caught'` := `true`. If  $|J'| = 1$  and  $\hat{j} \in J'$  then set `caught'` := `false`. If  $|J'| = 0$  then set `caught'` :=  $\perp$ .

<sup>5</sup> We use standard techniques [8, 16] to ensure that  $\mathcal{S}$  runs in expected polynomial time; details are omitted for the sake of the exposition.

7. If  $|J'| = 1$  and  $\text{caught}' = \text{true}$ , then send  $\text{cert} := (j, \text{trans}_j, \mathcal{H}_j, \mathbf{c}_j, \sigma_j, \text{seed}_j^B)$  to  $\mathcal{A}$  (where  $j$  is the unique index in  $J'$ ) and halt.  
Otherwise, send  $(\hat{j}, \{\text{seed}_j^A\}_{j \neq \hat{j}}, \text{witness}_{\hat{j}})$  to  $\mathcal{A}$ .
8. Receive  $\text{GC}$ ,  $\{A_i\}_{i \in [n_1]}$ ,  $\{h_{i,b}^A\}_{i \in [n_1], b \in \{0,1\}}$ ,  $\{Z_{i,b}\}_{i \in [n_3], b \in \{0,1\}}$ , and the corresponding decommitments from  $\mathcal{A}$ . If any of the decommitments are incorrect, send  $\perp$  to  $\mathcal{F}$  and halt.  
Otherwise, there are two possibilities:
  - If  $|J'| = 1$  and  $\text{caught}' = \text{false}$ , then use  $\{B_{j,i,b}\}_{i \in [n_2], b \in \{0,1\}}$  and the value  $y$  received from  $\mathcal{F}$  to compute an output  $z$  exactly as an honest  $\text{P}_B$  would. Send  $z$  to  $\mathcal{F}$  and halt.
  - If  $|J'| = 0$ , then compute an effective input  $x \in \{0,1\}^{n_1}$  using  $\text{seed}_j^A$  and the input-wire labels  $\{A_i\}_{i \in [n_1]}$ . Send  $x$  to  $\mathcal{F}$  and halt.

We now show that the joint distribution of the view of  $\mathcal{A}$  and the output of  $\text{P}_B$  in the ideal world is computationally indistinguishable from the joint distribution of the view of  $\mathcal{A}$  and the output of  $\text{P}_B$  in a real protocol execution. We prove this by considering a sequence of experiments, where the output of each is defined to be the view of  $\mathcal{A}$  and the output of  $\text{P}_B$ , and showing that the output of each is computationally indistinguishable from the output of the next one.

**Expt<sub>0</sub>**. This is the ideal-world execution between  $\mathcal{S}$  (as described above) and the honest  $\text{P}_B$  holding some input  $y$ , both interacting with functionality  $\mathcal{F}$ .

By inlining the actions of  $\mathcal{S}$ ,  $\mathcal{F}$ , and  $\text{P}_B$ , we may rewrite the experiment as follows:

1. Choose uniform  $\kappa$ -bit strings  $\{\text{seed}_j^B\}_{j \in [\lambda]}$ , set  $h_j \leftarrow \text{Com}(\text{seed}_j^B)$  for all  $j$ , and send  $\{h_j\}_{j \in [\lambda]}$  to  $\mathcal{A}$ .
2. For all  $j \in [\lambda]$ , run  $\Pi_{\text{OT}}$  with  $\mathcal{A}$ , using input 0 and randomness derived from  $\text{seed}_j^B$ . Obtain  $\{\text{seed}_j^A\}_{j \in [\lambda]}$  as the outputs. Let  $\text{trans}_j$  denote the transcript of the  $j$ th execution.
3. For  $j \in [\lambda]$ , run an execution of  $\Pi_{\text{OT}}$  with  $\mathcal{A}$ , using input  $0^{n_2}$  and randomness derived from  $\text{seed}_j^B$ . Let  $\mathcal{H}_j$  denote the transcript hash of the  $j$ th execution.
4. Receive  $\{\mathbf{c}_j\}_{j \in [\lambda]}$  from  $\mathcal{A}$ .
5. Receive  $\{\sigma_j\}$  from  $\mathcal{A}$ . If any of the signatures are invalid, then  $\text{P}_B$  outputs  $\perp$  and the experiment halts.
6. For all  $j \in [\lambda]$ , use  $\text{seed}_j^A$  and the messages sent previously to  $\mathcal{A}$  to simulate the computation of an honest  $\text{P}_A$  in steps 3 and 4, and in particular compute  $\hat{\mathcal{H}}_j, \hat{\mathbf{c}}_j$ . Let  $J$  be the set of indices for which  $(\hat{\mathcal{H}}_j, \hat{\mathbf{c}}_j) \neq (\mathcal{H}_j, \mathbf{c}_j)$ .

There are now three cases:

- If  $|J| \geq 2$ , send  $\text{cert} := (j, \text{trans}_j, \mathcal{H}_j, \mathbf{c}_j, \sigma_j, \text{seed}_j^B)$  to  $\mathcal{A}$  (for uniform  $j \in J$ ). Then  $\text{P}_B$  outputs **corrupted** and the experiment halts.
  - If  $|J| = 1$  then with probability  $\epsilon$  set  $\text{caught} := \text{true}$  and with the remaining probability set  $\text{caught} := \text{false}$ . If  $\text{caught} = \text{true}$  then  $\text{P}_B$  outputs **corrupted** (but the experiment continues below in either case).
  - If  $|J| = 0$  then set  $\text{caught} := \perp$  and continue below.
- 0'. Rewind  $\mathcal{A}$  and run steps 1'–6' below until  $|J'| = |J|$  and  $\text{caught}' = \text{caught}$  (using standard techniques [8, 16] to ensure the experiment runs in expected polynomial time).

- 1'. Choose uniform  $\hat{j} \in [\lambda]$ . For  $j \neq \hat{j}$ , choose uniform  $\kappa$ -bit strings  $\{\text{seed}_j^B\}$  and set  $h_j \leftarrow \text{Com}(\text{seed}_j^B)$ . Set  $h_{\hat{j}} \leftarrow \text{Com}(0^\kappa)$ . Send  $\{h_j\}_{j \in [\lambda]}$  to  $\mathcal{A}$ .
- 2'. For all  $j \neq \hat{j}$ , run  $\Pi_{\text{OT}}$  with  $\mathcal{A}$ , using input 0 and randomness derived from  $\text{seed}_j^B$ . Obtain  $\{\text{seed}_j^A\}_{j \neq \hat{j}}$  as the outputs of these executions. For the  $\hat{j}$ th execution, use the simulator  $\mathcal{S}_{\text{OT}}$  for protocol  $\Pi_{\text{OT}}$ , thus extracting both  $\text{seed}_{\hat{j}}^A$  and  $\text{witness}_{\hat{j}}$ . Let  $\text{trans}_{\hat{j}}$  denote the transcript of the  $\hat{j}$ th execution.
- 3'. For  $j \neq \hat{j}$ , run an execution of  $\Pi_{\text{OT}}$  with  $\mathcal{A}$  using input  $0^{n_2}$  and randomness derived from  $\text{seed}_j^B$ . For  $j = \hat{j}$ , use the simulator  $\mathcal{S}_{\text{OT}}$  for protocol  $\Pi_{\text{OT}}$ , thus extracting  $\{B_{\hat{j},i,b}\}_{i \in [n_2], b \in \{0,1\}}$ . Let  $\mathcal{H}_j$  denote the transcript hash of the  $j$ th execution.
- 4'. Receive  $\{c_j\}_{j \in [\lambda]}$  from  $\mathcal{A}$ .
- 5'. Receive  $\{\sigma_j\}$  from  $\mathcal{A}$ . If any of the signatures are invalid, then return to step 1'.
- 6'. For all  $j \in [\lambda]$ , use  $\text{seed}_j^A$  and the messages sent previously to simulate the computation of an honest  $\text{P}_A$  in steps 3' and 4', and in particular compute  $\hat{\mathcal{H}}_j, \hat{c}_j$ . Let  $J'$  be the set of indices for which  $(\hat{\mathcal{H}}_j, \hat{c}_j) \neq (\mathcal{H}_j, c_j)$ . If  $|J'| = 1$  and  $\hat{j} \notin J'$  then set  $\text{caught}' := \text{true}$ . If  $|J'| = 1$  and  $\hat{j} \in J'$  then set  $\text{caught}' := \text{false}$ . If  $|J'| = 0$  then set  $\text{caught}' := \perp$ .
7. If  $|J'| = 1$  and  $\text{caught}' = \text{true}$ , then send  $\text{cert} := (j, \text{trans}_j, \mathcal{H}_j, c_j, \sigma_j, \text{seed}_j^B)$  to  $\mathcal{A}$  (where  $j$  is the unique index in  $J'$ ) and halt. Otherwise, send  $(\hat{j}, \{\text{seed}_j^A\}_{j \neq \hat{j}}, \text{witness}_{\hat{j}})$  to  $\mathcal{A}$ .
8. Receive  $\text{GC}$ ,  $\{A_i\}_{i \in [n_1]}$ ,  $\{h_{i,b}^A\}_{i \in [n_1], b \in \{0,1\}}$ ,  $\{Z_{i,b}\}_{i \in [n_3], b \in \{0,1\}}$ , and the corresponding decommitments from  $\mathcal{A}$ . If any of the decommitments are incorrect, then  $\text{P}_B$  outputs  $\perp$  and the experiment halts. Otherwise, there are two possibilities:
  - If  $|J'| = 1$  and  $\text{caught}' = \text{false}$  then use  $\{B_{\hat{j},i,b}\}_{i \in [n_2], b \in \{0,1\}}$  and  $y$  to compute  $z$  exactly as in the protocol.  $\text{P}_B$  outputs  $z$  and the experiment halts.
  - If  $|J'| = |J| = 0$ , compute an effective input  $x \in \{0,1\}^{n_1}$  using  $\text{seed}_{\hat{j}}^A$  and the input-wire labels  $\{A_i\}_{i \in [n_1]}$ . Then  $\text{P}_B$  outputs  $f(x, y)$  and the experiment halts.

**Expt<sub>1</sub>**. Here we modify the previous experiment in the following way: Choose a uniform  $\hat{j} \in [\lambda]$  at the outset of the experiment. Then in step 6:

- If  $|J| \geq 2$  then send  $\text{cert} := (j, \text{trans}_j, \mathcal{H}_j, c_j, \sigma_j, \text{seed}_j^B)$  to  $\mathcal{A}$  for uniform  $j \in J \setminus \{\hat{j}\}$ . Then  $\text{P}_B$  outputs **corrupted** and the experiment halts.
- if  $|J| = 1$  set  $\text{caught} := \text{true}$  if  $\hat{j} \notin J$  and set  $\text{caught} := \text{false}$  if  $\hat{j} \in J$ .

Since  $\hat{j} \notin J$  with probability  $\epsilon$  when  $|J| = 1$ , the outputs of **Expt<sub>1</sub>** and **Expt<sub>0</sub>** are identically distributed.

**Expt<sub>2</sub>**. The previous experiment is modified as follows: In step 1, do not choose  $\text{seed}_{\hat{j}}^B$ . Instead, in step 1 set  $h_{\hat{j}} \leftarrow \text{Com}(0^\kappa)$ , and in steps 2 and 4 use true randomness in the  $\hat{j}$ th execution of  $\Pi_{\text{OT}}$ .

It is immediate that the distribution of the output of **Expt<sub>2</sub>** is computationally indistinguishable from the distribution of the output of **Expt<sub>1</sub>**.

**Expt<sub>3</sub>**. We change the previous experiment in the following way: In steps 2 and 4, use  $\mathcal{S}_{\text{OT}}$  to run the  $\hat{j}$ th instances of  $\Pi_{\text{OT}}$ . In doing so, extract all of  $\mathcal{A}$ 's inputs in those executions.

It follows from security of  $\Pi_{\text{OT}}$  that the distribution of the output of **Expt<sub>3</sub>** is computationally indistinguishable from the distribution of the output of **Expt<sub>2</sub>**.

**Expt<sub>3a</sub>**. Because steps 1'–4' in **Expt<sub>3</sub>** are identical to steps 1–4, we can “collapse” the rewinding and thus obtain the following experiment **Expt<sub>3a</sub>** that is statistically indistinguishable from **Expt<sub>3</sub>** (with the only difference occurring in case of an aborted rewinding in the latter):

1. Choose uniform  $\hat{j} \in [\lambda]$ . For  $j \neq \hat{j}$ , choose uniform  $\kappa$ -bit strings  $\{\text{seed}_j^B\}$  and set  $h_j \leftarrow \text{Com}(\text{seed}_j^B)$ . Set  $h_{\hat{j}} \leftarrow \text{Com}(0^\kappa)$ . Send  $\{h_j\}_{j \in [\lambda]}$  to  $\mathcal{A}$ .
2. For all  $j \neq \hat{j}$ , run  $\Pi_{\text{OT}}$  with  $\mathcal{A}$ , using input 0 and randomness derived from  $\text{seed}_j^B$ . Obtain  $\{\text{seed}_j^A\}_{j \neq \hat{j}}$  as the outputs of these executions. For the  $\hat{j}$ th execution, use the simulator  $\mathcal{S}_{\text{OT}}$  for protocol  $\Pi_{\text{OT}}$ , thus extracting both  $\text{seed}_{\hat{j}}^A$  and  $\text{witness}_{\hat{j}}$ . Let  $\text{trans}_j$  denote the transcript of the  $j$ th execution.
3. For all  $j \neq \hat{j}$ , run  $\Pi_{\text{OT}}$  with  $\mathcal{A}$  using input  $0^{n_2}$  and randomness derived from  $\text{seed}_j^B$ . For  $j = \hat{j}$ , use the simulator  $\mathcal{S}_{\text{OT}}$  for protocol  $\Pi_{\text{OT}}$ , thus extracting  $\{B_{j,i,b}\}_{i \in [n_2], b \in \{0,1\}}$ . Let  $\mathcal{H}_j$  denote the transcript hash of the  $j$ th execution.
4. Receive  $\{c_j\}_{j \in [\lambda]}$  from  $\mathcal{A}$ .
5. Receive  $\{\sigma_j\}$  from  $\mathcal{A}$ . If any of the signatures are invalid, then  $\text{P}_B$  outputs  $\perp$  and the experiment halts.
6. For all  $j \in [\lambda]$ , use  $\text{seed}_j^A$  and the messages sent previously to  $\mathcal{A}$  to simulate the computation of an honest  $\text{P}_A$  in steps 3 and 4, and in particular compute  $\hat{\mathcal{H}}_j, \hat{c}_j$ . Let  $J$  be the set of indices for which  $(\hat{\mathcal{H}}_j, \hat{c}_j) \neq (\mathcal{H}_j, c_j)$ .

There are now two cases:

- If  $|J| \geq 2$ , or if  $|J| = 1$  and  $\hat{j} \notin J$ , then choose uniform  $j \in J \setminus \{\hat{j}\}$  and send  $\text{cert} := (j, \text{trans}_j, \mathcal{H}_j, c_j, \sigma_j, \text{seed}_j^B)$  to  $\mathcal{A}$ . Then  $\text{P}_B$  outputs corrupted and the experiment halts.
  - If  $|J| = 1$  and  $\hat{j} \in J$ , or if  $|J| = 0$ , then continue below.
7. Send  $(\hat{j}, \{\text{seed}_j^A\}_{j \neq \hat{j}}, \text{witness}_{\hat{j}})$  to  $\mathcal{A}$ .
  8. Receive GC,  $\{A_i\}_{i \in [n_1]}$ ,  $\{h_{i,b}^A\}_{i \in [n_1], b \in \{0,1\}}$ ,  $\{Z_{i,b}\}_{i \in [n_3], b \in \{0,1\}}$ , and the corresponding decommitments from  $\mathcal{A}$ . If any of the decommitments are incorrect, then  $\text{P}_B$  outputs  $\perp$  and the experiment halts.

Otherwise, there are two possibilities:

- If  $|J| = 1$  then  $\text{P}_B$  uses  $\{B_{j,i,b}\}_{i \in [n_2], b \in \{0,1\}}$  and  $y$  to compute  $z$  exactly as in the protocol.  $\text{P}_B$  outputs  $z$  and the experiment halts.
- If  $|J| = 0$ , then compute an effective input  $x \in \{0,1\}^{n_1}$  using  $\text{seed}_{\hat{j}}^A$  and the input-wire labels  $\{A_i\}_{i \in [n_1]}$ . Then  $\text{P}_B$  outputs  $f(x, y)$  and the experiment halts.

**Expt<sub>4</sub>**. We modify the previous experiment as follows: In step 8, if  $|J| = 0$  (and  $\text{P}_B$  has not already output  $\perp$  in that step), use  $y$  to compute  $z$  exactly as in the protocol. Then  $\text{P}_B$  outputs  $z$  and the experiment halts.

Since  $|J| = 0$ , we know that  $c_j$  is a commitment to a correctly computed garbled circuit along with commitments to (correctly permuted) input-wire labels



$\{A_{j,i,b}\}$  and output-wire labels. Thus—unless  $\mathcal{A}$  has managed to violate the commitment property of **Com**—if  $P_B$  does not output  $\perp$  in this step it must be the case that the values **GC**,  $\{A_i\}_{i \in [n_1]}$ ,  $\{h_{i,b}^A\}_{i \in [n_1], b \in \{0,1\}}$ , and  $\{Z_{i,b}\}_{i \in [n_3], b \in \{0,1\}}$  sent by  $\mathcal{A}$  in step 8 are correct. Moreover, since  $|J| = 0$  the execution of  $\Pi_{OT}$  in step 4 was run honestly by  $\mathcal{A}$  using correct input-wire labels  $\{B_{j,i,b}\}$ . Thus, evaluating **GC** using  $\{A_i\}_{i \in [n_1]}$  and  $\{B_{j,i,y[i]}\}$  yields a result that is equal to  $f(x, y)$  as computed in **Expt<sub>3</sub>**.

Since **Com** is computationally binding, this means that the distribution of the output of **Expt<sub>4</sub>** is computationally indistinguishable from the distribution of the output of **Expt<sub>3a</sub>**.

**Expt<sub>5</sub>**. Here we change the previous experiment in the following way: The computation in step 6 is done only for  $j \in [\lambda] \setminus \{\hat{j}\}$ ; let  $\hat{J} \subseteq [\lambda] \setminus \{\hat{j}\}$  be the set of indices for which  $(\hat{\mathcal{H}}_j, \hat{c}_j) \neq (\mathcal{H}_j, c_j)$ . Then:

- If  $\hat{J} \neq \emptyset$  choose uniform  $j \in \hat{J}$  and send  $\text{cert} := (j, \text{trans}_j, \mathcal{H}_j, c_j, \sigma_j, \text{seed}_j^B)$  to  $\mathcal{A}$ . Then  $P_B$  outputs **corrupted** and the experiment halts.
- If  $\hat{J} = \emptyset$  then run steps 7 and 8 as in **Expt<sub>4</sub>**.

Letting  $J$  be defined as in **Expt<sub>4</sub>**, note that

$$|J| \geq 2 \text{ or } |J| = 1; \hat{j} \notin J \iff \hat{J} \neq \emptyset$$

and

$$|J| = 1, \hat{j} \in J \text{ or } |J| = 0 \iff \hat{J} = \emptyset.$$

Thus, the outputs of **Expt<sub>4</sub>** and **Expt<sub>5</sub>** are identically distributed.

**Expt<sub>6</sub>**. We now modify the previous experiment by running the  $\hat{j}$ th instances of  $\Pi_{OT}$  honestly in steps 2 and 4, using input 1 in step 2 and input  $y$  in step 4.

It follows from security of  $\Pi_{OT}$  that the distribution of the output of **Expt<sub>6</sub>** is computationally indistinguishable from the distribution of the output of **Expt<sub>5</sub>**.

**Expt<sub>7</sub>**. Finally, we modify the previous experiment so the  $\hat{j}$ th instance of  $\Pi_{OT}$  in steps 2 and 4 uses pseudorandomness derived from a uniform seed  $\text{seed}_j^B$ , and we compute  $h_j \leftarrow \text{Com}(\text{seed}_j^B)$ .

It is immediate that the distribution of the output of **Expt<sub>7</sub>** is computationally indistinguishable from the distribution of the output of **Expt<sub>6</sub>**.

Since **Expt<sub>7</sub>** corresponds to a real-world execution of the protocol between  $\mathcal{A}$  and  $P_B$  holding input  $y$ , this completes the proof.

### Covert Security—Malicious $P_B$

Let  $\mathcal{A}$  be an adversary corrupting  $P_B$ . We construct the following simulator  $\mathcal{S}$  that runs  $\mathcal{A}$  as a subroutine while playing the role of  $P_B$  in the ideal world interacting with  $\mathcal{F}$ :

0. Run **Gen** to generate keys  $(pk, sk)$ , and send  $pk$  to  $\mathcal{A}$ .
1. Receive  $\{h_j\}_{j \in [\lambda]}$  from  $\mathcal{A}$ .

2. Use the simulator  $\mathcal{S}_{\text{OT}}$  for protocol  $\Pi_{\text{OT}}$  to interact with  $\mathcal{A}$ . In this way,  $\mathcal{S}$  extracts  $\mathcal{A}$ 's inputs  $\{b_j\}_{j \in [\lambda]}$ ; let  $J := \{j : b_j = 1\}$ . As part of the simulation, return uniform  $\kappa$ -bit strings  $\{\text{seed}_j^A\}_{j \notin J}$  and  $\{\text{witness}_j\}_{j \in J}$  as output to  $\mathcal{A}$ .
3. For each  $j \notin J$ , run this step exactly as an honest  $\text{P}_A$  would. For each  $j \in J$  do:
  - If  $|J| = 1$  then let  $\hat{j}$  be the unique index in  $J$ . Use  $\mathcal{S}_{\text{OT}}$  to interact with  $\mathcal{A}$  in the  $\hat{j}$ th execution of  $\Pi_{\text{OT}}$ . In this way,  $\mathcal{S}$  extracts  $\mathcal{A}$ 's input  $y$  for that execution. Send  $y$  to  $\mathcal{F}$ , and receive in return a value  $z$ . Compute

$$(\{A_{\hat{j},i}\}, \{B_{\hat{j},i}\}, \text{GC}_{\hat{j}}, \{Z_{\hat{j},i,b}\}) \leftarrow \mathcal{S}_{\text{Gb}}(1^\kappa, \mathcal{C}, z),$$

where we let  $\{A_{\hat{j},i}\}$  correspond to input wires of  $\text{P}_A$  and  $\{B_{\hat{j},i}\}$  correspond to input wires of  $\text{P}_B$ . Return  $\{B_{\hat{j},i}\}$  as output to  $\mathcal{A}$  from this execution of  $\Pi_{\text{OT}}$ .

- If  $|J| > 1$  then act as an honest  $\text{P}_A$  would but using true randomness.
4. For each  $j \notin J$ , compute  $c_j$  exactly as an honest  $\text{P}_A$  would. For each  $j \in J$  do:
    - If  $|J| = 1$  then compute  $h_{j,i,0}^A \leftarrow \text{Com}(A_{j,i})$  and let  $h_{j,i,1}^A$  be a commitment to the 0-string. Compute  $c_j \leftarrow \text{Com}(\text{GC}_j, \{h_{j,i,b}^A\}, \{Z_{j,i,b}\})$ , where each pair  $(h_{j,i,0}^A, h_{j,i,1}^A)$  is in random permuted order.
    - If  $|J| > 1$  then compute  $c_j$  exactly as an honest  $\text{P}_A$  would but using true randomness.

Send  $\{c_j\}_{j \in [\lambda]}$  to  $\mathcal{A}$ .

- 5–6. Compute signatures  $\{\sigma_j\}$  as an honest  $\text{P}_A$  would, and send them to  $\mathcal{A}$ .
7. If  $|J| \neq 1$  then abort. Otherwise, receive  $(\hat{j}, \{\text{seed}_j\}_{j \neq \hat{j}}, \text{witness}_{\hat{j}})$  from  $\mathcal{A}$  and verify these as an honest  $\text{P}_A$  would. (If verification fails, then abort.)
8. Send  $\text{GC}_{\hat{j}}, \{A_{\hat{j},i}\}, \{h_{\hat{j},i,b}^A\}$  (in the same permuted order as before), and  $\{Z_{\hat{j},i,b}\}$  to  $\mathcal{A}$ , along with the corresponding decommitments. Then halt.

We show that the distribution of the view of  $\mathcal{A}$  in the ideal world is computationally indistinguishable from its view in a real protocol execution. (Note that  $\text{P}_A$  has no output.) Let  $\mathbf{Expt}_0$  be the ideal-world execution between  $\mathcal{S}$  (as described above) and the honest  $\text{P}_A$  holding some input  $x$ , both interacting with functionality  $\mathcal{F}$ .

**Expt<sub>1</sub>**. Here we modify the previous experiment when  $|J| = 1$  as follows. In step 3, compute

$$(\{A_{\hat{j},i,b}\}, \{B_{\hat{j},i,b}\}, \text{GC}_{\hat{j}}, \{Z_{\hat{j},i,b}\}) \leftarrow \text{Gb}(1^\kappa, \mathcal{C}),$$

and return the values  $\{B_{\hat{j},i,y[i]}\}$  as output to  $\mathcal{A}$  from the simulated execution of  $\Pi_{\text{OT}}$  in that step. In steps 4 and 8, the values  $A_{\hat{j},i,x[i]}$  are used in place of  $A_{\hat{j},i}$ .

It follows from security of the garbling scheme that the view of  $\mathcal{A}$  in **Expt<sub>1</sub>** is computationally indistinguishable from its view in **Expt<sub>0</sub>**.

**Expt<sub>2</sub>**. Now we change the previous experiment when  $|J| = 1$  as follows: In step 3, compute  $h_{j,i,b}^A \leftarrow \text{Com}(A_{j,i,b})$  for all  $i, b$ . It follows from the hiding property of the commitment scheme that the view of  $\mathcal{A}$  in **Expt<sub>2</sub>** is computationally indistinguishable from its view in **Expt<sub>1</sub>**.

**Expt<sub>3</sub>**. This time, the previous experiment is modified by executing protocol  $\Pi_{\text{OT}}$  with  $\mathcal{A}$  when  $|J| = 1$  in step 3. Security of  $\Pi_{\text{OT}}$  implies that the view of  $\mathcal{A}$  in **Expt<sub>3</sub>** is computationally indistinguishable from its view in **Expt<sub>2</sub>**.

**Expt<sub>4</sub>**. The previous experiment is now modified in the following way. In step 2, also choose uniform  $\{\text{seed}_j^A\}_{j \in J}$  and  $\{\text{witness}_j^A\}_{j \notin J}$ , and use pseudorandomness derived from  $\{\text{seed}_j^A\}_{j \in J}$  in steps 3 and 4 in place of true randomness. Also, in step 7 continue to run the protocol as an honest  $P_A$  would even in the case that  $|J| \neq 1$ .

It is not hard to show that when  $|J| \neq 1$  then  $P_A$  aborts in **Expt<sub>4</sub>** with all but negligible probability. Computational indistinguishability of  $\mathcal{A}$ 's view in **Expt<sub>4</sub>** and **Expt<sub>3</sub>** follows.

**Expt<sub>5</sub>**. Finally, we change the last experiment by executing protocol  $\Pi_{\text{OT}}$  in step 2. It follows from the security of  $\Pi_{\text{OT}}$  that the view of  $\mathcal{A}$  in **Expt<sub>5</sub>** is computationally indistinguishable from its view in **Expt<sub>4</sub>**.

Since **Expt<sub>5</sub>** corresponds to a real-world execution of the protocol, this completes the proof.

## Public Verifiability and Defamation Freeness

It is easy to check (by inspecting the protocol) that whenever an honest  $P_B$  outputs **corrupted** then it also outputs a valid certificate. Thus our protocol satisfies public verifiability. It is similarly easy to verify defamation freeness under the assumptions of the theorem.


## 6 Implementation and Evaluation

We implemented our PVC protocol using the optimizations from Section 4.3 and state-of-the-art techniques for garbling [4, 20], oblivious transfer [5], and OT extension [12]. Our implementation uses SHA-256 for the hash function (modeled as a random oracle) and the standard ECDSA implementation provided by `openssl` as the signature scheme. We target  $\kappa = 128$  in our implementation.

We evaluate our protocol in both LAN and WAN settings. In the LAN setting, the network bandwidth is 1 Gbps and the latency is less than 1 ms; in the WAN setting, the bandwidth is 200 Mbps and the latency is 75 ms. In either setting, the machines running the protocol have 32 cores, each running at 3.0GHz. Due to pipelining, we never observe any issues with memory usage. All reported timing results are computed as the average of 10 executions.

### 6.1 Certificate Size

The size of the certificate in our protocol is independent of the circuit size or the lengths of the parties' inputs. The following figure gives a graphical decomposition of the certificate. (Note that since we instantiate `Com` by a random oracle as discussed in Section 4.3, we do not need to include an extra decommitment in the certificate.) In total, a certificate requires 354 bytes.



$\mathcal{H}_j$  contains 4 hash values, corresponding to a 4-round OT protocol obtained by piggybacking a 2-round OT-extension protocol with a 3-round base-OT protocol. The signature size varies from 70–72 bytes; we allocate 72 bytes for the signature so the total length of a certificate is fixed.

## 6.2 Comparison to Prior PVC Protocols

Because it enables signed-OT extension, the PVC protocol by Kolesnikov and Malozemoff [13] (the *KM15 protocol*) would be strictly more efficient than the original PVC protocol by Asharov and Orlandi [2]. We therefore focus our attention on the KM15 protocol. We compare our protocol to theirs in three respects.

**Parameters.** We briefly discuss the overhead needed to achieve deterrence factors larger than  $\frac{1}{2}$  for each protocol. Recall that in the KM15 protocol the overall deterrence factor  $\epsilon$  depends on both the garbled-circuit replication factor  $\lambda$  and the XOR-tree expansion factor  $\nu$  as  $\epsilon = (1 - \frac{1}{\lambda}) \cdot (1 - 2^{-\nu+1})$ . For deterrence  $\epsilon \approx 1 - \frac{1}{2^k}$ , setting  $\lambda = 2^{k+1}, \nu = k + 2$  gives the best efficiency. In contrast, our protocol achieves this deterrence with  $\lambda = 2^k, \nu = 1$ , which means garbling half as many circuits and avoiding the XOR-tree approach altogether. For example, to achieve deterrence  $\epsilon = 7/8$ , our protocol garbles 8 circuits, whereas prior work would need to garble 16 circuits. Additionally, prior work would need to execute  $5\times$  as many OTs. (Plus, in prior work each OT is actually a *signed* OT, which is more expensive than standard OT; see next.)

**Signed OT vs. standard OT.** Signed OT induces higher costs than standard OT in terms of both communication and computation. As an illustration, fix the deterrence factor to  $1/2$ . In that case our protocol runs OT extension twice, where each is used for  $n_2$  OTs on  $\kappa$ -bit strings. Compared to this, the KM15 protocol needs to run  $3n_2$  OTs on  $2\kappa$ -bit strings. The total communication complexity of the OT step (for the input-wire labels) is  $4\kappa n_2$  bits in our protocol, while in the KM15 protocol it is  $3 * 2 * 3\kappa n_2 + 3 * 2.6\kappa n_2 = 25.8\kappa n_2$  bits, more than  $6\times$  higher.

Moreover, signed OT also has a very high computational overhead:

- Signed-OT extension needs to use a wider matrix (by a factor of roughly  $2.6\times$ ) compared to standard OT extension. Besides the direct penalty this incurs, a wider matrix means that the correlation-robust hash  $H$  cannot be based on fixed-key AES but must instead be based on a hash function like SHA-256. This impacts performance significantly.
- As part of signed-OT extension,  $P_B$  needs to reveal  $\kappa$  random columns in the matrix. Even with AVX operations, this incurs significant computational overhead.

Circuit	$n_1$	$n_2$	$n_3$	$ \mathcal{C} $
AES-128	128	128	128	6800
SHA-128	256	256	160	37300
SHA-256	256	256	256	90825
Sorting	131072	131072	131072	10223K
Integer mult.	2048	2048	2048	4192K
Hamming dist.	1048K	1048K	22	2097K

**Table 1.** Circuits used in our evaluation. The parties’ input lengths are  $n_1$  and  $n_2$ , and the output length is  $n_3$ . The number of AND gates in the circuit is denoted by  $|\mathcal{C}|$ .

Signed-OT extension [13] is complex, and we did not implement it in its entirety. However, we modified an existing (standard) OT-extension protocol to match the matrix width required by signed-OT extension; this can be used to give a conservative lower bound on the performance of signed-OT extension. Our results indicate that signed-OT extension requires roughly  $5\times$  more computation than state-of-the-art OT extension.

**Certificate size.** In the KM15 protocol, the certificate size is at least  $2\kappa \cdot n_2$  bits. Even for AES (with only 128-bit input length), this gives a certificate roughly  $10\times$  larger than ours.

### 6.3 Comparing to Semi-Honest and Malicious Protocols

We believe our PVC protocol provides an excellent performance/security tradeoff that makes it the best choice for many applications of secure computation.

**Performance.** Our protocol is not much less efficient than the best known semi-honest protocols, and is significantly faster than the best known malicious protocols.

**Security.** The PVC model provides much more meaningful guarantees than the notion of semi-honest security, and may be appropriate for many (even if not all) applications of secure computation where full malicious security is overkill.

To support the first point, we compare the performance of our PVC protocol against state-of-the-art two-party computation protocols. The semi-honest protocol we compare against is a garbled-circuit protocol including all existing optimizations; for the malicious protocol we use the recent implementation of Wang et al. [19]. Our comparison uses the circuits listed in Table 1.

**Running time.** In Table 2 we compare the running time of our protocol to that of a semi-honest protocol. From the table, we see that over a LAN our protocol adds at most 36% overhead except in two cases: AES and Hamming-distance computation. For AES, the reason is that the circuit is small and so the overall time is dominated by the base OTs. For Hamming distance, the total input

Circuit	LAN setting			WAN setting		
	Our PVC	Semi-honest	Slowdown	Our PVC	Semi-honest	Slowdown
AES-128	25 ms	15 ms	1.60×	960 ms	821 ms	1.17×
SHA-128	34 ms	25 ms	1.36×	1146 ms	977 ms	1.17×
SHA-256	48 ms	38 ms	1.27×	1252 ms	1080 ms	1.16×
Sort.	3468 ms	2715 ms	1.28×	13130 ms	12270 ms	1.07×
Mult.	1285 ms	1110 ms	1.16×	5707 ms	5462 ms	1.04×
Hamming	2585 ms	1550 ms	1.67×	11850 ms	6317 ms	1.69×

**Table 2.** Comparing the running times of our protocol and a semi-honest protocol in the LAN and WAN settings.

size is equal to the number of AND gates in the circuit; therefore, the cost of processing the inputs becomes more significant.

In the WAN setting, our PVC protocol incurs only 17% overhead except for the Hamming-distance example (for a similar reason as above).

The comparison between our PVC protocol and the malicious protocol is shown in Table 3. As expected, our PVC protocol achieves much better performance, by a factor of 4–18×

Circuit	LAN setting			WAN setting		
	Our PVC	Malicious [19]	Speedup	Our PVC	Malicious [19]	Speedup
AES-128	25 ms	157 ms	6.41×	960 ms	11170 ms	11.6×
SHA-128	34 ms	319 ms	9.47×	1146 ms	13860 ms	12.1×
SHA-256	48 ms	612 ms	12.6×	1252 ms	17300 ms	13.8×
Sort.	3468 ms	45130 ms	13.0×	13130 ms	197900 ms	15.1×
Mult.	1285 ms	17860 ms	13.9×	5707 ms	99930 ms	17.5×
Hamming	2586ms	11380 ms	4.40×	11850 ms	76280 ms	6.44×

**Table 3.** Comparing the running times of our protocol and a malicious protocol in the LAN and WAN settings.

	AES-128	SHA-128	SHA-256	Sort.	Mult.	Hamming
Semi-honest	0.2218	1.165	2.800	313.1	128.0	96.01
Malicious [11]	3.545	17.69	42.95	2953	1228	662.7
Our PVC	0.2427	1.205	2.844	325.1	128.2	144.2

**Table 4.** Communication complexity in MB of our protocol with  $\lambda = 2$  and other protocols.

**Communication complexity.** We also compare the communication complexity of our protocol to other protocols in a similar way; see Table 4. In this comparison we use the same semi-honest protocol as above, but use the more communication-efficient protocol by Katz et al. [11] as the malicious protocol. We see that, with the exception of the Hamming-distance example, the communication of our protocol is very close to the semi-honest case.

## 6.4 Higher Deterrence Factors

Another important aspect of our protocol is how the performance is affected by the deterrence factor. Recall that the deterrence factor  $\epsilon$  is the probability that a cheating party is caught, and in our protocol  $\epsilon = 1 - \frac{1}{\lambda}$  where  $\lambda$  is the garbled-circuit replication factor. The performance of our protocol as a function of  $\epsilon$  is shown in Table 5. We see that when doubling the value of  $\lambda$ , the running time of the protocol increases by only  $\approx 20\%$  unless the circuit is very small (in which case the cost of the base OTs dominates the total running time). The running time when  $\epsilon = 3/4$  (i.e.,  $\lambda = 4$ ) is still less than twice the running time of a semi-honest protocol.

		AES-128	SHA-128	SHA-256	Sort.	Mult.	Hamming
LAN	$\lambda = 2$	25	34	49	3468	1285	2586
	$\lambda = 4$	36	46	59	3554	1308	3156
	$\lambda = 8$	47	57	71	3954	1396	4856
	$\lambda = 16$	101	127	152	6238	2355	7143
	$\lambda = 32$	175	228	229	7649	2995	12984
WAN	$\lambda = 2$	960	1146	1252	13130	5707	11850
	$\lambda = 4$	1112	1375	1700	14400	5952	12899
	$\lambda = 8$	1424	1912	2436	16130	6167	19840
	$\lambda = 16$	1920	2094	2191	19087	7801	36270
	$\lambda = 32$	3228	3434	3535	25197	9229	64468

**Table 5.** Running time in milliseconds of our protocol for different  $\lambda$ .  $\epsilon = 1 - \frac{1}{\lambda}$ .

## 6.5 Scalability

Our protocol scales linearly in all parameters, and so can easily handle large circuits. To demonstrate this, we benchmarked our protocol with different input lengths, output lengths, and circuit sizes. Initially, the input and output lengths are all 128 bits, and the circuit size is 1024 AND gates. We then gradually increase one of the input/output lengths or circuit size (while holding everything else constant) and record the running time. Since the dependence is linear in all cases, we report only the marginal cost (i.e., the slope), summarized in Table 6.

	$n_1$ ( $\mu\text{s}/\text{bit}$ )	$n_2$ ( $\mu\text{s}/\text{bit}$ )	$n_3$ ( $\mu\text{s}/\text{bit}$ )	$ \mathcal{C} $ ( $\mu\text{s}/\text{gate}$ )
LAN	0.20	0.88	0.23	0.29
WAN	0.61	3.13	0.62	1.10

**Table 6.** Scalability of our protocol.

## References

1. Gilad Asharov, Yehuda Lindell, Thomas Schneider, and Michael Zohner. More efficient oblivious transfer and extensions for faster secure computation. In *20th ACM Conf. on Computer and Communications Security (CCS)*, pages 535–548. ACM Press, 2013.
2. Gilad Asharov and Claudio Orlandi. Calling out cheaters: Covert security with public verifiability. In *Advances in Cryptology—Asiacrypt 2012*, volume 7658 of *LNCS*, pages 681–698. Springer, 2012.
3. Yonatan Aumann and Yehuda Lindell. Security against covert adversaries: Efficient protocols for realistic adversaries. *Journal of Cryptology*, 23(2):281–343, 2010.
4. Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. Efficient garbling from a fixed-key blockcipher. In *2013 IEEE Symposium on Security & Privacy*, pages 478–492. IEEE, 2013.
5. Tung Chou and Claudio Orlandi. The simplest protocol for oblivious transfer. In *Progress in Cryptology—Latincrypt 2015*, volume 9230 of *LNCS*, pages 40–58. Springer, 2015.
6. Ivan Damgård, Martin Geisler, and Jesper Buus Nielsen. From passive to covert security at low cost. In *7th Theory of Cryptography Conference—TCC 2010*, volume 5978 of *LNCS*, pages 128–145. Springer, 2010.
7. O. Goldreich. *Foundations of Cryptography, vol. 2: Basic Applications*. Cambridge University Press, Cambridge, UK, 2004.
8. Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *Journal of Cryptology*, 9(3):167–190, 1996.
9. Vipul Goyal, Payman Mohassel, and Adam Smith. Efficient two party and multi party computation against covert adversaries. In *Advances in Cryptology—Eurocrypt 2008*, volume 4965 of *LNCS*, pages 289–306. Springer, 2008.
10. Jonathan Katz and Rafail Ostrovsky. Round-optimal secure two-party computation. In *Advances in Cryptology—Crypto 2004*, volume 3152 of *LNCS*, pages 335–354. Springer, 2004.
11. Jonathan Katz, Samuel Ranellucci, Mike Rosulek, and Xiao Wang. Optimizing authenticated garbling for faster secure two-party computation. In *Advances in Cryptology—Crypto 2018, Part III*, volume 10993 of *LNCS*, pages 365–391. Springer, 2018.
12. Marcel Keller, Emmanuela Orsini, and Peter Scholl. Actively secure OT extension with optimal overhead. In *Advances in Cryptology—Crypto 2015, Part I*, volume 9215 of *LNCS*, pages 724–741. Springer, 2015.
13. Vladimir Kolesnikov and Alex J. Malozemoff. Public verifiability in the covert model (almost) for free. In *Advances in Cryptology—Asiacrypt 2015, Part II*, volume 9453 of *LNCS*, pages 210–235. Springer, 2015.
14. Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In *35th Intl. Colloquium on Automata, Languages, and Programming (ICALP), Part II*, volume 5126 of *LNCS*, pages 486–498. Springer, 2008.
15. Y. Lindell and B. Pinkas. A proof of security of Yao’s protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, 2009.
16. Yehuda Lindell. A note on constant-round zero-knowledge proofs of knowledge. *Journal of Cryptology*, 26(4):638–654, 2013.
17. Yehuda Lindell. Fast cut-and-choose-based protocols for malicious and covert adversaries. *Journal of Cryptology*, 29(2):456–490, 2016.



18. Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In *Advances in Cryptology—Eurocrypt 2007*, volume 4515 of *LNCS*, pages 52–78. Springer, 2007.
19. Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated garbling and efficient maliciously secure two-party computation. In *24th ACM Conf. on Computer and Communications Security (CCS)*, pages 21–37. ACM Press, 2017.
20. Samee Zahur, Mike Rosulek, and David Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In *Advances in Cryptology—Eurocrypt 2015, Part II*, volume 9057 of *LNCS*, pages 220–250. Springer, 2015.