

Constant-Round Multiparty Computation Using a Black-Box Pseudorandom Generator

Ivan Damgård^{1*} and Yuval Ishai^{2**}

¹ Aarhus University (ivan@daimi.au.dk)

² Technion (yuvali@cs.technion.ac.il)

Abstract. We present a constant-round protocol for general secure multiparty computation which makes a *black-box* use of a pseudorandom generator. In particular, the protocol does not require expensive zero-knowledge proofs and its communication complexity does not depend on the computational complexity of the underlying cryptographic primitive. Our protocol withstands an active, adaptive adversary corrupting a minority of the parties. Previous constant-round protocols of this type were only known in the semi-honest model or for restricted classes of functionalities.

1 Introduction

General secure computation is often perceived as being inherently impractical. One valid reason for this perception is the fact that all current protocols either require many rounds of interaction (e.g., [20, 5, 27, 35, 13, 22]), or alternatively require only a constant number of rounds but make use of expensive zero-knowledge proofs for each gate of the circuit being computed (e.g., [40, 4, 28, 6, 26, 34, 25]). Indeed, in all constant-round protocols from the literature, players need to provide zero-knowledge proofs for statements that involve the computation of a pseudorandom generator or other cryptographic primitives on which the “semi-honest” version of the protocol relies. Thus, these protocols make a *non-black-box* use of their underlying cryptographic primitives. We stress that this holds for all settings of secure computation with security against malicious parties, both in the two-party and in the multi-party case. The only exceptions to this general state of affairs are unconditionally secure protocols that apply to restricted classes of functionalities such as NC^1 [2] or protocols that require an exponential amount of computation [3].

In this work we consider the setting of multiparty computation with an honest majority, and present a general constant-round protocol that makes a *black-box* use of a pseudorandom generator.³ Similarly to all general constant-round protocols from the literature, our protocol relies on Yao’s garbled circuit technique [40],

* Supported by BRICS, Basic research in Computer Science, Center of the Danish National Research Foundation and FICS, Foundations in Cryptography and Security, funded by the Danish Natural Sciences Research Council.

** Research supported by Israel Science Foundation grant 36/03.

³ Here we assume the standard model of secure point-to-point channels.

which was later adapted to the multi-party setting by Beaver, Micali, and Rogaway [4]. The latter “BMR protocol” requires players to verifiably secret-share seeds to a PRG as well as the outputs of the PRG on these seeds. To ensure that this is done correctly, the protocol makes a non-black-box use of the PRG by requiring players to prove via (distributed) zero-knowledge that the shared seeds are consistent with the shared PRG outputs. We get around this problem by modifying the basic structure of the BMR protocol, using “distributed symmetric encryption” and error-correction to replace the zero-knowledge proofs. Before providing a more detailed account of our results we give some more background to put them in context.

Black-box reductions in cryptography. Most reductions between cryptographic primitives are (fully) black box, in the sense that they implement a primitive A by using some other primitive B as an oracle, without depending on the implementation details of B . Moreover, the security proof of such reductions is also black-box in the sense that an adversary breaking A can be used as an oracle in order to break the underlying primitive B . (See [36] for a more detailed definition and discussion.) In contrast, a non-black-box reduction can use the “code” of B when implementing A . Most examples for non-black-box reductions in cryptography are ones in which the construction of A requires parties to prove in zero-knowledge statements that involve the computation of the underlying primitive B . For instance, the construction of an identification scheme from a one-way function [14] makes a non-black-box use of the one-way function. A rich line of work, originating from [24], uses oracle separations to rule out the existence of (various forms of) black-box reductions. Most notably, it is shown in [24] that there is no black-box reduction from key agreement to a one-way function. A common interpretation of such results is that they rule out the existence of “practically feasible” reductions. Indeed, all known examples for non-black-box reductions in cryptography involve a considerable overhead. In the context of cryptographic protocols, this overhead typically involves not only local computation but also communication: the *communication* complexity of the protocol B depends on the *computational* complexity of the underlying primitive A .⁴ Our work provides further demonstration for the usefulness of distinguishing between the two types of reductions.

Constant-round secure computation. The question of implementing secure computation in a constant number of rounds has attracted a considerable amount of attention. The first general constant-round protocol for secure two-party computation was given by Yao [40]. Yao’s original protocol considered only the case of semi-honest parties; an extension to the case of malicious parties (equivalently, an *active* adversary) was given by Lindell [28]. While Yao’s original protocol makes a black-box use of the underlying primitives (a pseudorandom generator and oblivious transfer), the protocol from [28] relies on the methodology of Gol-

⁴ In some cases it is possible to reduce the communication overhead by using communication-efficient zero-knowledge *arguments* (cf. [31]). However, this approach would make the computational overhead even higher.

dreich, Micali, and Wigderson [20] and thus makes a non-black-box use of these primitives. Recently, Katz and Ostrovsky obtained a two-party protocol with an optimal exact round complexity [25].

An extension of Yao’s protocol to the case of multiparty computation with an honest majority was given by Beaver, Micali, and Rogaway [4] (see also [37, 39]). Similarly to the two-party case, the BMR protocol makes a black-box use of a PRG in the semi-honest case and a non-black-box use of a PRG in the malicious case. (Because of the honest majority assumption, the protocol does not need to rely on OT.) Constant-round multiparty protocols withstanding a *dishonest* majority were recently obtained by Katz et al. [26] and by Pass [33]. (The latter protocol also achieves bounded concurrent security, extending a previous two-party protocol of Pass and Rosen [34].) These protocols are only proved secure with respect to a non-adaptive adversary and allow the adversary to prevent honest parties from receiving any output (even when corrupting a minority of the parties). Like the two-party protocols, these protocols follow the GMW methodology and thus make a non-black-box use of the underlying primitives.

Finally, there has also been a considerable amount of work on *unconditionally* secure constant-round multiparty computation in the case of an honest majority (e.g. [2, 15, 23, 9]). Unfortunately, all known protocols in this setting can only be efficiently applied to restricted classes of functions such as NC^1 or non-deterministic logspace.

1.1 Our Results

We consider the model of computationally secure multiparty computation against an active, adaptive adversary corrupting up to $t < n/2$ players. Our default network model assumes secure point-to-point channels and the availability of broadcast (see more on that later). As stated above, our main result is a new “black-box feasibility” result. Specifically, we construct the first general constant-round protocol which makes a black-box use of a PRG (equivalently, using [21], a black-box use of a one-way function). Since much of our motivation comes from the goal of making secure computation more efficient, we also attempt to minimize the amount of interaction and communication required by our protocols. To this end, it is convenient to cast the protocols in the following “client-server” framework.

The client-server model. We divide the players into “input clients” who provide inputs, “output clients” who receive outputs, and “servers” who perform the actual computation. The security of the protocol should hold as long as at most t servers are corrupted, regardless of the number of corrupted clients. These three sets of players need not be disjoint, hence this is a strict generalization of the standard MPC framework in which all parties play all three roles. It also represents a likely scenario for applying MPC in practice, using specialized (but untrusted) servers to perform the bulk of the work. We stress again that this is just a refinement of the standard model. The main advantage of this refinement, besides conceptual clarity, is that it allows to decouple the number

of “consumers” from the required “level of security”. (The latter depends on the number of servers and the security threshold.) For instance, we can have just two clients and many servers (which may be viewed a distributed implementation of two-party computation), or a very large number of clients and only few servers. The latter might be the most realistic setting for secure computations involving inputs from many players.

Linear preprocessing. We present our main protocol in two stages. First, we present a protocol in what we call the “linear preprocessing model”. In this model, it is assumed that there is a trusted setup phase where a dealer can provide clients and servers with linearly-correlated resources, e.g., Shamir-shares of random secrets. Then we use standard subprotocols for emulating the trusted setup in the plain model. The linear preprocessing model is motivated by the pseudorandom secret-sharing technique of [11]: when the number of servers is small, linear preprocessing can be emulated using a “once and for all” setup phase in which (roughly $\binom{n}{t}$) replicated and independent seeds are given to the players. Following this setup, the players can locally generate the required correlated shares without further interaction.⁵

Our protocols. Our main protocol in the linear preprocessing model requires only two communication rounds when $t < n/5$. In the first round each input client broadcasts its masked inputs to the servers, and in the second round the servers send to each output client a total of $O(n^2|C|k)$ bits, where $|C|$ is the size of the circuit being computed and k is a security parameter. In the plain model, one can obtain similar protocols at the cost of a higher communication complexity and additional rounds of interaction. When $t < n/5$, it suffices to use 3 rounds of interaction by relying on a VSS protocol from [17]. Alternatively, it is possible to tolerate $t < n/3$ or even $t < n/2$ malicious servers at the expense of further increasing the communication and the (constant) number of rounds.

In the case of computing a randomized functionality which has no inputs, the 2-round protocol in the linear preprocessing model becomes totally non-interactive when combined with pseudorandom secret-sharing. That is, to securely compute such a functionality it suffices for each server to send a single message to each output client, without using any broadcasts. Such non-interactive protocols can be used to obtain efficient distributed implementations of a trusted dealer in a wide range of applications.

On the use of broadcast. As in most of the MPC literature, our network model assumes the availability of broadcast as an atomic primitive. However, using the (expected) constant-round broadcast protocol of Feldman and Micali [16, 29], our protocols can be turned into (expected) constant-round protocols also in the point to point model. Concerning the communication complexity, since it is possible to implement our protocol so that the number of broadcasts involved is

⁵ The method of [11] was only proved to be secure in the case of a non-adaptive adversary. Thus, when relying on pseudorandom secret sharing our protocol loses its provable adaptive security.

independent of $|C|$ (using the techniques of [22]), one can get the same (amortized) communication complexity in the point-to-point model. Moreover, in the typical scenario where the number of servers is small, even a “brute-force” simulation of the broadcasts will not have a major impact on efficiency.

Organization. The remainder of the paper is organized as follows. In Section 2 we define our security model and preprocessing models, and present some standard subprotocols in these models. Our main protocol and its variations are presented in Section 3, where the underlying distributed encryption idea is highlighted in Section 3.3. Due to our elaborate use of techniques from previous works (mostly in the context of *information-theoretic* multiparty computation), we omit some of the low-level details and assume the reader’s familiarity with standard MPC techniques from the literature. Some discussions and extensions (e.g., the case $t < n/2$) were omitted for lack of space and can be found in the full version.

2 Preliminaries

The Model. We consider a system consisting of several players, who interact in synchronous rounds via authenticated secure point-to-point channels and a broadcast medium. Players can be designated three different roles: *input clients* who hold inputs, *output clients* who receive outputs, and *servers* who may be involved in the actual computation. As discussed in Section 1.1, this is just a generalization of the standard model, where each party can play all three roles. We denote the number of servers by n . The functionalities we wish to compute only receive inputs from input clients and only provide outputs to output clients. For simplicity we will only explicitly consider deterministic functionalities providing all output clients with the same output, though an extension of our results to the general case is straightforward.⁶ (In contrast, we will employ sub-protocols that compute randomized functionalities and provide servers and input clients with outputs as well.)

We assume by default an active, adaptive, rushing adversary corrupting at most t servers. (There is no restriction on the number of corrupted clients.) We refer the reader to, e.g., [7] for the standard definition of security in this model.

Our protocols will employ secret-sharing over a finite field $K = \text{GF}(2^k)$, where k is a security parameter that will be used as the length of a seed to a PRG. Slightly abusing notation, each server P_i is assigned a unique nonzero value $i \in K$.

2.1 Linear Preprocessing

As discussed in Section 1.1, it will be convenient to describe and analyze our protocols in the *linear preprocessing model*, where we allow some restricted trusted setup as described below. The protocols can then be converted to the *plain*

⁶ Standard reductions from the general case to this special case involve interaction between output clients. This can be avoided by directly generalizing the protocol to the randomized multi-output case.

model, where no setup assumptions or preprocessing are allowed, at the price of some efficiency loss.

In the linear preprocessing model, we assume a dealer who initially gives to each player a set of values in K or in its subfield $\text{GF}(2)$. The values distributed by the dealer are restricted to be “linearly correlated”. Specifically, the dealer picks a random codeword in a linear code defined over K or over $\text{GF}(2)$, and then hands to each player a subset of the coordinates in the codeword. It is public which subsets are used, but the values themselves are private. This procedure can be repeated multiple times, possibly using different linear codes.

Of course, we do not expect that such a dealer would exist in practice. This is only a convenient abstraction, that can be formalized as an ideal functionality. We later separately look at how such a dealer may be implemented.

Note that in Shamir’s secret sharing scheme, shares are computed as linear functions of the secret and random elements chosen by the dealer. We may therefore assume that the dealer can give to players Shamir shares of a random secret or of 0. More concretely, we assume that the following subroutines are available. When they are called in our protocol descriptions that follow, this should be taken to mean that the players retrieve from their preprocessed material values as specified below.

RandSS(t) Each server P_i obtains $f(i)$, where f is a random polynomial over K of degree at most t .

RandSS₀(t) Same as **RandSS**(t), except that f is subject to $f(0) = 0$.

RandSS_{bin}(t) Same as **RandSS**(t), except that f is subject to the restriction that $f(0)$ is either 0 or 1. Note that this correlation pattern is linear over $\text{GF}(2)$.

RandSS ^{P} (t) Same as **RandSS**(t), except that player P additionally receives f .

RandSS_{bin} ^{P} (t) Same as **RandSS**_{bin}(t), except that P additionally receives f .

As discussed in Section 1.1, the linear preprocessing model is motivated by the pseudorandom secret-sharing technique from [11] (see also [19]). When the number of servers is small, a “once and for all” setup is sufficient for enabling players to execute any number of calls to the above subroutines without having to communicate.

One can emulate the linear preprocessing model in the plain model using constant-round interaction between players. This is trivial for a passive adversary, and can be done for an active adversary based on standard verifiable secret sharing schemes from the literature (e.g., [5, 12, 10]). In particular, [12] shows how to build VSS from any linear secret sharing scheme, and this can conveniently be used to implement **RandSS**_{bin}(t) using VSS over $\text{GF}(2)$.

Our protocols will invoke the following variants of VSS as subroutines.

VSS ^{P} (t) Player P has a value $s \in K$ as private input. The goal is for each server to receive a Shamir share of s . Using linear preprocessing, this only requires a single round of broadcast: in the setup phase we invoke **RandSS** ^{P} (t). Then the value $r = f(0)$ can be computed by P , and P broadcasts $z = s - r$. Each server P_i takes $z + f(i)$ to be his private output.

$VSS_{bin}^P(t)$ Player P has private input a value $b \in \text{GF}(2)$. The goal is for each server to receive a Shamir share of b (computed over K). This can be implemented similarly to $VSS^P(t)$, replacing $\text{RandSS}^P(t)$ by $\text{RandSS}_{bin}^P(t)$.

2.2 Secure Computation of Low-Degree Polynomials

In the linear preprocessing model, we now show how to securely compute the following functionality, which will be useful later.

The functionality is defined by $Q()$, a degree d polynomial over K in l variables x_1, \dots, x_l . Each input client is to supply values for some of the variables, the others are to be chosen at random by the functionality. For each server P_i we define an index set D_i ; these sets are mutually disjoint and designate subsets of the random inputs to $Q()$. If $j \in D_i$, the functionality will output x_j to P_i . Finally, the functionality will output $Q(x_1, \dots, x_l)$ to the output clients, as well as Shamir shares of this value to the servers.

The functionality is specifically designed to fit into our protocol for computing Yao-garbled circuits to be presented later. In particular, some of the random values x_j will be used as encryptions keys. Each such key has to be known to exactly one server, and this is the reason why the functionality outputs some of the x_j 's to the servers. This functionality, denoted by F_{Q, D_1, \dots, D_n} , is more precisely defined as follows:

1. In the first round, it receives from each honest input client the x_j 's this client supplies. In addition, it receives from the honest input clients and servers a set of values of the form produced in the linear preprocessing model. More precisely, these additional inputs take the following form:
 - For each x_j that is supplied by input client I , a set of values for $\text{RandSS}^I(t)$ (as determined by a polynomial f_j).
 - For each x_j that is random, a set of values for $\text{RandSS}(t)$ (as determined by a polynomial f_j).
 - For each x_j that is random and where $j \in D_i$, a set of values for $\text{RandSS}^{P_i}(t)$ (as determined by a polynomial f_j).
 - A set of values for $\text{RandSS}_0(dt)$ (as determined by a polynomial f_0).

The functionality computes the shares and polynomials that all the above results in for the corrupt players and outputs this to the adversary. For instance, for every x_j supplied by corrupt input client I , this will be the polynomial $f_j()$. Note that, due to our assumed constraint on the number of corrupt servers, the honest players' information is enough to determine this information for the corrupt players. Also, for each x_j supplied by an honest input client, it sends $x_j - f_j(0)$ to the adversary. (No information is sent to honest players in this round.)

2. In round 2, the functionality receives from each corrupt input client the x_j 's that it is responsible for. For each $j \in D_i$, the functionality will output x_j to P_i . It then outputs to each server a Shamir share of the value $Q(x_1, \dots, x_l)$ generated by the polynomial $Q(f_1(), f_2(), \dots, f_l()) + f_0()$ (this will be a univariate polynomial of degree at most dt). Finally, the functionality outputs $Q(x_1, \dots, x_l)$ to all output clients.

We securely implement the above functionality in the linear preprocessing model using the following standard protocol:⁷

1. We do the following for each x_j : if x_j is supplied by input client I , execute $VSS^I(t)$ where I uses x_j as his private input. The communication implied by this is the only communication in the first round.
 If x_j is random, execute $\text{RandSS}(t)$. If x_j is random and $j \in D_i$, execute $\text{RandSS}^{P_i}(t)$. In all cases, a set of shares of x_j is obtained. Let $x_{j,i}$ be the share of x_j obtained by server P_i . We execute $\text{RandSS}_0(dt)$, creating shares of a degree dt polynomial that evaluates to 0 in 0. Let z_i be the share obtained by server P_i .
2. In the second round, each server P_i sends $Q(x_{1,i}, \dots, x_{l,i}) + z_i$ to each output client. Each output client considers the values he receives as points on a degree dt polynomial f , reconstructs the polynomial (applying error-correction in the active adversary case) and outputs $f(0)$. Each server P_i outputs $Q(x_{1,i}, \dots, x_{l,i}) + z_i$, and the values x_j for which $j \in D_i$.

We now show the security of this protocol using Canetti's UC framework [8]. We only show this for environments that supply inputs of correct form as specified above. This is sufficient, since we will only use the functionality in conjunction with the linear preprocessing, which is assumed to produce values of the right form.

Theorem 1. *There exists a 2-round protocol computing F_{Q,D_1,\dots,D_n} , the protocol is secure for all environments that supply inputs for honest players as specified in the description of F_{Q,D_1,\dots,D_n} . Furthermore, the protocol is secure for an adaptive adversary corrupting at most t servers and an arbitrary number of clients. For a passive adversary, we assume $dt < n$, for an active adversary we need $(d+2)t < n$. The communication complexity involves each output client receiving n field elements and each input client broadcasting its (masked) inputs.*

Proof. To prove security, we describe the required simulator (ideal model adversary), which as usual works by running an internal copy of the real-life adversary. In the first round, we receive a set of polynomials and shares from the ideal functionality, which we pass on to the adversary. In particular, this includes, for each corrupt input client, a random polynomial $f_j(\cdot)$ of degree at most t , for each x_j this client supplies. When the adversary broadcasts a values r_j on behalf of the client, we compute $x_j = f_j(0) - r_j$ and give x_j as input to the ideal functionality. We also received from the ideal functionality $f_j(0) - x_j$ for each x_j supplied by an honest client. We use these values to simulate the broadcasts of honest clients. Note that for each corrupt server P_i , the information received from the functionality now defines a share $x_{j,i}$ of each x_j , and that it also defines a share z_i of the degree dt sharing of 0.

⁷ For our purposes the degree d of Q will be no larger than 3, hence we will not consider optimizations that apply to a larger d .

In the second round, the ideal functionality sends $Q(x_1, \dots, x_l)$ to the simulator (we assume the adversary has corrupted at least one output client, otherwise the simulation becomes trivial). For each corrupt output client, we use the value $Q(x_1, \dots, x_l)$ and the shares of it known by corrupt servers to interpolate a random polynomial f of degree at most dt with $f(0) = Q(x_1, \dots, x_l)$ and $f(i) = Q(x_{1,i}, \dots, x_{l,i}) + z_i$, for each corrupt P_i , where the $x_{j,i}, z_i$ are the previously defined values for P_i . Then for each honest server P_i , we claim $f(i)$ as the value sent by P_i .

To establish adaptive security, we now show how to reconstruct the history of the players if they are corrupted after the protocol. Earlier corruptions are handled by truncating the reconstruction procedure.

If an input client or a server is corrupted after the protocol, we learn all his input, and pass this on to the adversary. This already determines his view of the protocol, and is consistent with what the adversary already knows by definition of the ideal functionality.

If an output client is corrupted after the protocol, we need not produce new values, as all output clients receive the same set of messages from honest servers, and these were already produced earlier.

It is straightforward to verify that this simulation leads to perfect indistinguishability between the real and ideal process. Namely, the only item that is produced using different algorithms in the two is the polynomial that determines $Q(x_1, \dots, x_l)$. However, it is in both cases a random polynomial of degree at most dt under the constraints that it is consistent with corrupt servers' shares and its value at 0 is $Q(x_1, \dots, x_l)$. \square

The protocol can be easily extended to secure computation in parallel of several low-degree polynomials, on some set of inputs, where some inputs may go to several polynomials. While it is not clear that this is implied by the composition theorem (because of the overlapping inputs) it can be shown by a trivial extension of the above proof. One can also modify the protocol so that only shares of $Q(x_1, \dots, x_l)$ are computed for the servers, and not the value itself, by simply not sending the shares to the output clients.

3 Constant-Round MPC Using a Black-Box PRG

In this section we present our main protocol. We start in Section 3.1 by describing a variant of Yao's garbled circuit technique on which we rely. Then, in Section 3.2 we sketch the BMR approach for computing a garbled circuit in a distributed way. Finally, in Sections 3.3 and 3.4 we describe our modified approach.

3.1 The Basic Garbled Circuit Technique

Loosely speaking, the garbled circuit technique allows to represent a circuit C on ℓ input bits by an "encrypted" circuit $E(C)$ along with ℓ pairs of random keys, such that given $E(C)$ and the ℓ keys corresponding to a specific input

b_1, \dots, b_ℓ , one can efficiently compute the output $C(b_1, \dots, b_\ell)$, but this is the only information about the inputs that can be learned. Yao designed a method for generating such encrypted circuits, and used it to obtain a general constant-round two-party protocol for semi-honest parties. (See [30] for a formal proof of security of Yao's protocol and [32, 1] for other variants of this technique.) This protocol was generalized to the multi-party case by Beaver, Micali and Rogaway [4]. The circuit encryption process can be done in parallel for every gate in C , yielding a constant-round protocol for secure function evaluation.

We now describe how the basic technique for garbling a circuit works, by specifying how a trusted functionality could prepare an encrypted circuit and input keys as above. We will later show various protocols for implementing this functionality.

Without loss of generality, we assume the function to be computed is described as a Boolean circuit C with 2-input NAND gates. Let the number of wires in the circuit be W . We number the wires from 0 to $W - 1$. For simplicity, we assume the circuit produces just a single output bit, to be learnt by all output clients, where this bit corresponds to the last wire, number $W - 1$. Each input wire w has a bit b_w assigned to it, where each such bit is supplied by an input client. We will be using an index $0 \leq j \leq 2W - 1$, where the values $2w, 2w + 1$ are assigned to wire w .

We assume we have available a secure secret-key encryption scheme $E_S()$, where S is a k -bit key. We need to assume that the cryptosystem is semantically secure as long as each key is used on at most $2z$ messages, of length $k + 1$ bits each, where z is the maximal fan-out in C .

To compute the garbled circuit, we get as input the bits b_w for each input wire, and then proceed as follows (see below for intuition):

- For every wire w , choose a random bit λ_w (masking the true value of the wire) and random keys S_{2w}, S_{2w+1} .
- For every gate g in C , do the following: suppose g has input wires α, β and output wire γ . Define the following values

$$\begin{aligned} a_g^{00} &= S_{2\gamma+\delta_g^{00}} ; & \delta_g^{00} &= (\lambda_\alpha \text{ nand } \lambda_\beta) \oplus \lambda_\gamma \\ a_g^{01} &= S_{2\gamma+\delta_g^{01}} ; & \delta_g^{01} &= (\lambda_\alpha \text{ nand } \bar{\lambda}_\beta) \oplus \lambda_\gamma \\ a_g^{10} &= S_{2\gamma+\delta_g^{10}} ; & \delta_g^{10} &= (\bar{\lambda}_\alpha \text{ nand } \lambda_\beta) \oplus \lambda_\gamma \\ a_g^{11} &= S_{2\gamma+\delta_g^{11}} ; & \delta_g^{11} &= (\bar{\lambda}_\alpha \text{ nand } \bar{\lambda}_\beta) \oplus \lambda_\gamma \end{aligned}$$

Define $A_g^{cd} = (a_g^{cd}, \delta_g^{cd})$, for $c, d \in \{0, 1\}$. We compute the encryptions

$$\begin{aligned} &E_{S_{2\alpha}}(E_{S_{2\beta}}(A_g^{00})) \\ &E_{S_{2\alpha}}(E_{S_{2\beta+1}}(A_g^{01})) \\ &E_{S_{2\alpha+1}}(E_{S_{2\beta}}(A_g^{10})) \\ &E_{S_{2\alpha+1}}(E_{S_{2\beta+1}}(A_g^{11})) \end{aligned}$$

- We output, for each g , the 4 encryptions as above along with the mask λ_{W-1} of the output wire - this is the encrypted circuit which we denote by $E(C)$. We also output, for each input wire w , the values $b_w \oplus \lambda_w$ and $S_{2w+(b_w \oplus \lambda_w)}$ - these are the encrypted inputs.

A word about the underlying intuition behind this: assume we knew all the inputs and did an ordinary computation of the circuit. This would result in assigning a bit b_w to every wire w . Instead, we get to know exactly one of the two encryption keys that are assigned to each wire, namely the key $S_{2w+(b_w \oplus \lambda_w)}$ and the bit $b_w \oplus \lambda_w$, and this is ensured for all input wires initially. We can think of this information as an encrypted representation of the bit b_w . This also means that by making λ_{W-1} public, we reveal the output bit, and only that bit.

The idea behind making the individual gates work in this scenario is to encrypt the keys and bits that might be assigned to the gate's output wire under keys assigned to input wires in such a way that players will be able to decrypt the "correct" key and bits for the output wire, and only this information. For instance, suppose that some gate g in the circuit has input wires α, β and output wire γ . If the information known for the input wires is $S_{2\alpha+c}, c$ and $S_{2\beta+d}, d$ for bits c, d , then the bit that should be revealed for output wire is $\delta_g^{cd} = ((c \oplus \lambda_\alpha) \text{ nand } (d \oplus \lambda_\beta)) \oplus \lambda_\gamma$, and so the key that should be revealed is $S_{2\gamma+\delta_g^{cd}}$. The idea is therefore to encrypt these two values under $S_{2\alpha+c}$ and $S_{2\beta+d}$, for all 4 values of c, d .

Anyone who is given encrypted circuit and inputs can compute the output by the following *local circuit evaluation procedure*: for each input wire w , the key $S_{2w+(b_w \oplus \lambda_w)}$ and the bit $b_w \oplus \lambda_w$ are given. There will now be a number of gates, for which a key and a bit are known for both input wires. Let g be such a gate, say with input wires α, β and output wire γ . Since we know $b_\alpha \oplus \lambda_\alpha$ and $b_\beta \oplus \lambda_\beta$, we know which of the encryptions associated to g we can decrypt, namely those where both involved keys are known. We decrypt and obtain as result a key and a bit, which are easily seen to be $S_{2\gamma+(b_\gamma \oplus \lambda_\gamma)}$ and the bit $b_\gamma \oplus \lambda_\gamma$. Continuing this way, we will obtain a key and a bit $b_{W-1} \oplus \lambda_{W-1}$ for the output wire $W-1$. Since we also know λ_{W-1} , we can compute the output bit b_{W-1} .

3.2 The BMR Protocol

The protocol from [4] can be seen as a concrete proposal for an encryption scheme E as required for the garbled circuit technique and a protocol for computing $E(C)$. Their scheme assumes a pseudorandom generator G taking as input a k -bit seed (such a generator can be constructed from any one-way function [21]). For a seed s , the output of $G(s)$ is split into k -bit blocks where the j 'th block is denoted by $G(s)_j$.

A key S in the encryption scheme consists of n subkeys $S = (s^1, s^2, \dots, s^n)$, each of which is k bits long, and where initially s^i is known only to P_i . An element $m \in K$ is encrypted under S as $E_S(m) = m \oplus G(s^1)_j \oplus \dots \oplus G(s^n)_j$, assuming m is the j 'th k -bit string we encrypt under S .

Assuming m and each s^i have been secret shared among the servers, we can securely compute the encryption by having server i locally compute and secret share $G(s^i)_j$. We can then use linearity of the secret sharing to get shares of $E_S(m)$, and send these shares to the output clients. This will work, and makes only a black-box use of G , if the adversary is passive. But if he is active, each server needs to prove in zero-knowledge that he has computed and secret

shared $G(s^i)_j$ correctly. In general, this requires generic zero-knowledge techniques, which means we no longer make a black-box use of G , and also leads to a major loss of efficiency.

3.3 Our Distributed Encryption Scheme

We now suggest a different encryption scheme for the garbled circuit technique, allowing to avoid the use of zero-knowledge proofs in the case of an active adversary. We assume as before a pseudorandom generator G which expands a k bit seed. A key S consists again of n subkeys $S = (s^1, s^2, \dots, s^n)$ where initially s^i is known only to P_i .

Consider now a situation where a message $m \in K$ has been secret shared among the n servers using a polynomial of degree d , where $t \leq d < n$. Let m_i denote the share of m given to P_i . To encrypt such a message under a key S , we will let each server encrypt the share he knows under his part of the key (expanded by G).

We define $E_S^j(m) = (G(s^1)_j \oplus m_1, \dots, G(s^n)_j \oplus m_n)$. Having received the parts of the ciphertext $E_S^j(m)$ from the servers and given the key S , one can decrypt each share and reconstruct m from the shares, where error correction⁸ is used to recover m if the adversary has actively corrupted some of the servers. The following lemma is straightforward.

Lemma 1. *The above distributed encryption scheme has the following properties:*

- If an adversary is given up to t of the s^i 's, and E_S^j is used on at most one message m , the encryption keeps m semantically secure.
- If the adversary is passive, and an honest output client is given S and receives $E_S^j(m)$ from the servers, he can decrypt correctly if $d < n$.
- If the adversary is active, and an honest output client is given S and receives $E_S^j(m)$ from the servers, he can decrypt correctly if $d + 2t < n$.

This generalizes in a straightforward way to cases where two keys U, V are used. We write $E_U^i(E_V^j(m)) = (G(u^1)_i \oplus G(v^1)_j \oplus m_1, \dots, G(u^n)_i \oplus G(v^n)_j \oplus m_n)$.

In the following, we will be encrypting several elements in K under the same key. This is done in the natural way, by using a fresh part of the output from G for each new element.

3.4 Distributed Computation of a Garbled Circuit

We now apply the distributed encryption idea for securely computing a garbled circuit using a black-box PRG. The protocol takes place in the linear preprocessing model, and will use the subroutines and protocols described in Sections 2.1 and 2.2.

⁸ In the case $t < n/2$ the subkeys and the message will be distributed using authenticated shares, in which case the decryption will involve a correction of *erasures* rather than errors.

1. In round 1, for each wire $w = 0..W - 1$ the servers execute $\text{RandSS}_{bin}(t)$ to create shares of the secret wire masks λ_w 's. Also, for $i = 1..n, j = 0..2W - 1$, they execute $\text{RandSS}^{P_i}(t)$ to create shares of the subkeys s_j^i , such that s_j^i is known to P_i . Finally, for each input bit b_w held by input client I_j , the players execute $\text{VSS}^{I_j}(t)$ (i.e., with I_j as dealer) and b_w as shared secret. Thus the only communication in round 1 consists of broadcasts done in the VSS subroutines.
2. In round 2, the servers first do some local computation.
 - For each input wire w and $i = 1..n$, each server locally computes a random share of the value $s_{2w+(b_w \oplus \lambda_w)}^i$. Note that since we work over a field of characteristic 2, this value can be written as a degree 2 polynomial, namely $(1 + b_w + \lambda_w)s_{2w}^i + (b_w + \lambda_w)s_{2w+1}^i$. We can therefore compute shares of the value $s_{2w+(b_w \oplus \lambda_w)}^i$ defined by a random degree $2t$ polynomial and send all these shares to the output clients, using the protocol from Section 2.2.
 - For each input wire w , the servers compute shares of the value $b_w \oplus \lambda_w$ and send them to the output clients.
 - For each gate g in the circuit, suppose the two inputs and output wire are wires α, β, γ , respectively. Then for each $i = 1..n$, the servers locally compute random shares of the values

$$\begin{aligned}
a_g^{00,i} &= s_{2\gamma+\delta_g^{00}}^i ; & \delta_g^{00} &= (\lambda_\alpha \text{ nand } \lambda_\beta) \oplus \lambda_\gamma \\
a_g^{01,i} &= s_{2\gamma+\delta_g^{01}}^i ; & \delta_g^{01} &= (\lambda_\alpha \text{ nand } \bar{\lambda}_\beta) \oplus \lambda_\gamma \\
a_g^{10,i} &= s_{2\gamma+\delta_g^{10}}^i ; & \delta_g^{10} &= (\bar{\lambda}_\alpha \text{ nand } \lambda_\beta) \oplus \lambda_\gamma \\
a_g^{11,i} &= s_{2\gamma+\delta_g^{11}}^i ; & \delta_g^{11} &= (\bar{\lambda}_\alpha \text{ nand } \bar{\lambda}_\beta) \oplus \lambda_\gamma
\end{aligned}$$

Note that these values can be written as degree 3 polynomials in the already shared values, for instance, $a_g^{00,i} = (\lambda_\alpha \lambda_\beta + \lambda_\gamma)s_{2\gamma}^i + (1 + \lambda_\alpha \lambda_\beta + \lambda_\gamma)s_{2\gamma+1}^i$. We can therefore use the protocol from Section 2.2 to locally compute these random shares (without sending them to the output clients).

3. Let $a_g^{cd} = (a_g^{cd,1}, \dots, a_g^{cd,n})$, for $c, d \in \{0, 1\}$. (This vector of n subkeys replaces the single key a_g^{cd} in the basic garbled circuit construction from Section 3.1). Define $A_g^{cd} = (A_g^{cd}, \delta_g^{cd})$. The servers can now reveal to the output clients encryptions of the form $E_{S_{2\alpha+c}}(E_{S_{2\beta+d}}(A_g^{cd}))$ using the distributed encryption scheme from Section 3.3. (Note that both the data we need to encrypt and the encryption subkeys are already shared in the required form.)
4. The output clients now apply the local circuit evaluation procedure described in Section 3.1, replacing ordinary decryption with distributed decryption (Section 3.3).

Theorem 2 (Black-box constant-round protocol in linear preprocessing model). *In the linear preprocessing model, there is a general 2-round MPC protocol making a black-box use of a pseudorandom generator. The protocol tolerates an active, adaptive adversary corrupting $t < n/5$ servers and an arbitrary*

number of clients.⁹ The communication complexity for computing a circuit C involves $O(n^2|C|k)$ bits sent to each output client, and each input client must broadcast its (masked) inputs to the n servers.

Proof sketch: Formally speaking, we want to prove that the above protocol realizes a functionality F_C that accepts inputs b_1, \dots, b_ℓ from the input clients and then outputs $C(b_1, \dots, b_\ell)$ to all output clients.

Let $F_{low-degree}$ be an extended version of F_{Q, D_1, \dots, D_n} , computing all polynomials of degree 2, 3 and shares we compute in the protocol π as described above (see the remarks following Theorem 1). Let $\pi^{F_{low-degree}}$ be the protocol we obtain by replacing in the natural way steps 1 and 2 in π by a call to $F_{low-degree}$. By Theorem 1 and the composition theorem, to show security of π , it is sufficient to show security of $\pi^{F_{low-degree}}$.

We now describe a black-box simulator for this protocol. The simulator proceeds by running internally copies of the linear preprocessing functionality, $F_{low-degree}$ and the (initially) honest players. The internal copies of honest players are called *virtual honest players*. They will be given 0's as input instead of the real values of their b_i 's (which are unknown to the simulator). Otherwise all these internal entities proceed according to the protocol. There are only two differences between this and the real process: When the adversary specifies input bits to $F_{low-degree}$, this in particular fixes values of the b_i 's for the corrupt players. The simulator then sends these bits to F_C . Second, when we get $C(b_1, \dots, b_\ell)$ from F_C and construct the encrypted circuit, we assign a number of ciphertexts to the output gate g , exactly one of which will be decrypted in the local evaluation procedure. The simulator will put the bit $\lambda_{W-1} \oplus C(b_1, \dots, b_\ell)$ as plaintext inside this encryption. This is done as follows: the degree $3t$ polynomial that defines this bit is of the form $g() + z()$ where $z()$ is a random degree $3t$ polynomial with $z(0) = 0$. We then change $z()$ to a random $z'()$ of the same degree, but such that $g(0) + z'(0)$ is that value we want and $z'()$ is consistent with the shares of corrupt players. This change introduces no inconsistencies in the view of the virtual players, since $z()$ is used for nothing else than randomizing g . The simulated execution now results in output $C(b_1, \dots, b_\ell)$ which is consistent with what the ideal functionality gives to the honest clients.

It remains to be described how the simulator handles corruptions. We describe how the simulator will reconstruct the view of a newly corrupted player. We assume the player is corrupted after the protocol terminates. For earlier corruptions, the reconstruction procedure is truncated appropriately. The general idea is that the simulator already has the views of the virtual honest players, including the one the adversary now wants to corrupt. We then modify this information so it becomes consistent with what we learn as a result of the corruption, without changing what the adversary already knows.

If an input client is corrupted, we learn his input bit(s), say b_w . We already broadcasted a value r_w related to this, and the virtual client has from the preprocessing a polynomial f_w with $f_w(0) = r_w$. We then change f_w to f'_w , so f'_w is

⁹ In the full paper, we give an optimized version of Theorem 1 for the case of a passive adversary, allowing us to prove the above theorem for $t < n/2$ in this case.

random of degree at most t , subject to $f'_w(0) = b_w \oplus r_w$, and $f_w(c) = f'_w(c)$ for all corrupt P_c . We now want to claim that the virtual honest players used the polynomial $g'_w() = f'_w() + r_w$ in the further computation, instead of $g_w() = f_w() + r_w$ we used so far. Note that, to compute the garbled inputs, the protocol computes a set of degree-2 polynomials. Consider one of them, say $Q(b_w, \lambda_w, s_{2w}, s_{2w+1})$. Say the last 3 variables are shared using polynomials $g_2(), g_3(), g_4()$. Before corruption, we had $Q(0, \lambda_w, s_{2w}, s_{2w+1})$ shared using a univariate polynomial of the form $Q(g_w(), g_2(), g_3(), g_4()) + z_w()$, where $z_w()$ is random of degree $2t$ and $z_w(0) = 0$. Define $z'_w()$ by

$$Q(g_w(), g_2(), g_3(), g_4()) + z_w() = Q(g'_w(), g_2(), g_3(), g_4()) + z'_w()$$

and change $z_w()$ to $z'_w()$. This will change honest virtual server's shares, but not the corrupted server's shares, by construction of $f'_w(), g'_w()$. It will also preserve the data sent to output clients. We now give to the adversary the updated view of the virtual client.

If a server or an output client is corrupted, note that this does not result in any new data learnt from F_C . We can therefore give the current view of the virtual output client or server to the adversary.

This concludes the description of the simulator. The intuition of the analysis of the simulation is that all plaintext data are identically distributed in simulation as in real execution, the difference lies in the data that remain encrypted, and this cannot be detected efficiently by semantic security of the encryption. \square

3.5 The Plain Model

To implement the above protocol in the plain model, we need to emulate the procedures for generating random shared secrets. The cost of the resulting protocol is dominated by the cost of emulating $O(n|C|)$ invocations of (different variants of) $\text{RandSS}(t)$. In the semi-honest case, each invocation of $\text{RandSS}(t)$ can be implemented in a straightforward way by letting each player distribute a random secret and output the sum of the shares it received. (In fact, it suffices that $t + 1$ players share secrets.) In the malicious case, one could use a similar procedure based on any standard constant-round VSS protocol from the literature (e.g., the one from [5]). In fact, using the 2-round VSS protocol from [17], one can obtain a 3-round protocol in the plain model (assuming $t < n/5$).

We can weaken the assumption on t in the active case to $t < n/3$ by replacing the non-interactive polynomial evaluation protocol from Theorem 1 by an interactive one (e.g., using [5]). The resulting protocol will have a larger (but still constant) number of rounds and a higher communication complexity. In the full version, we sketch how to use the VSS and multiplication protocol from [10] to further extend the feasibility result to the case $t < n/2$. This is based on two observations: first, a variant of our distributed encryption scheme can be used to encrypt values that have been shared under any VSS with a non-interactive reconstruction protocol. Second, by requiring that all values in the computation

as well as shares of these values are VSS'ed, we can obtain a multiplication protocol that is guaranteed to terminate in a constant number of rounds, even for the case of $t < n/2$. Thus, our main feasibility result in the plain model is the following:

Theorem 3 (Black-box constant-round protocol in plain model). *In the plain model, there is a general constant-round MPC protocol making a black-box use of a pseudorandom generator. The protocol tolerates an active, adaptive adversary corrupting $t < n/2$ servers and an arbitrary number of clients.*

References

- [1] B. Applebaum, Y. Ishai, and E. Kushilevitz. Computationally private randomizing polynomials and their applications. In *Proc. 20th Conference on Computational Complexity*, 2005.
- [2] J. Bar-Ilan and D. Beaver. Non-cryptographic fault-tolerant computing in a constant number of rounds. In *Proc. 8th ACM PODC*, pages 201–209, 1989.
- [3] D. Beaver, J. Feigenbaum, J. Kilian, and P. Rogaway. Security with low communication overhead (extended abstract). In *Proc. of CRYPTO '90*.
- [4] D. Beaver, S. Micali, and P. Rogaway. The round complexity of secure protocols (extended abstract). In *Proc. of 22nd STOC*, pages 503–513, 1990.
- [5] M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proc. of 20th STOC*, pages 1–10, 1988.
- [6] C. Cachin, J. Camenisch, J. Kilian, and J. Muller. One-round secure computation and secure autonomous mobile agents. In *Proceedings of ICALP' 00*, 2000.
- [7] R. Canetti. Security and composition of multiparty cryptographic protocols. In *J. of Cryptology*, 13(1), 2000.
- [8] R. Canetti. Ran Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. *FOCS 2001*: 136-145.
- [9] R. Cramer and I. Damgård. Secure distributed linear algebra in a constant number of rounds. In *Proc. Crypto 2001*.
- [10] R. Cramer, I. Damgård, S. Dziembowski, M. Hirt, and T. Rabin. Efficient Multiparty Computations Secure Against an Adaptive Adversary. In *Proc. EUROCRYPT 1999*, pages 311-326.
- [11] R. Cramer, I. Damgård, and Y. Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In *Proc. of second TCC*, 2005.
- [12] R. Cramer, I. Damgård, and U. Maurer. General secure multi-party computation from any linear secret-sharing scheme. In *Proc. of EUROCRYPT '00*, LNCS 1807, pp. 316-334, 2000.
- [13] R. Cramer, I. Damgård, and J. Nielsen. Multiparty computation from threshold homomorphic encryption. In *Proc. of EUROCRYPT '01*, LNCS 2045, pp. 280-299, 2001.
- [14] U. Feige, A. Fiat, and A. Shamir. Zero-Knowledge Proofs of Identity. *J. Cryptology* 1(2): 77-94 (1988).
- [15] Uri Feige, Joe Kilian, and Moni Naor. A minimal model for secure computation (extended abstract). In *Proc. 26th STOC*, pages 554–563. ACM, 1994.
- [16] P. Feldman and S. Micali. An Optimal Algorithm for Synchronous Byzantine Agreement. *SIAM. J. Computing*, 26(2):873–933, 1997.

- [17] R. Gennaro, Y. Ishai, E. Kushilevitz and T. Rabin. The Round Complexity of Verifiable Secret Sharing and Secure Multicast. In *Proceedings of the 33rd ACM Symp. on Theory of Computing (STOC '01)*, pages 580-589, 2001.
- [18] R. Gennaro, Y. Ishai, E. Kushilevitz and T. Rabin. On 2-round secure multiparty computation. In *Proc. Crypto '02*.
- [19] N. Gilboa and Y. Ishai. Compressing cryptographic resources. In *Proc. of CRYPTO '99*.
- [20] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game (extended abstract). In *Proc. of 19th STOC*, pages 218-229, 1987.
- [21] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364-1396, 1999.
- [22] M. Hirt and U. M. Maurer. Robustness for Free in Unconditional Multi-party Computation. *CRYPTO 2001*: 101-118.
- [23] Y. Ishai and E. Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *Proc. 41st FOCS*, pp. 294-304, 2000.
- [24] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. *Proceedings of 21st Annual ACM Symposium on the Theory of Computing*, 1989, pp. 44 - 61.
- [25] J. Katz and R. Ostrovsky. Round-Optimal Secure Two-Party Computation. In *CRYPTO 2004*, pages 335-354.
- [26] J. Katz, R. Ostrovsky, and A. Smith. Round Efficiency of Multi-party Computation with a Dishonest Majority. In *EUROCRYPT 2003*, pages 578-595.
- [27] J. Kilian. Founding cryptography on oblivious transfer. In *Proc. 20th STOC*, pages 20-31, 1988.
- [28] Y. Lindell. Parallel Coin-Tossing and Constant-Round Secure Two-Party Computation. *J. Cryptology* 16(3): 143-184 (2003). Preliminary version in *Crypto 2001*.
- [29] Y. Lindell, A. Lysyanskaya, and T. Rabin. Sequential composition of protocols without simultaneous termination. In *Proc. PODC 2002*, pages 203-212.
- [30] Y. Lindell and B. Pinkas. A Proof of Yao's Protocol for Secure Two-Party Computation. *Cryptology ePrint Archive*, Report 2004/175, 2004.
- [31] M. Naor and K. Nissim. Communication preserving protocols for secure function evaluation. In *Proc. STOC 2001*, pages 590-599.
- [32] M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *Proc. 1st ACM Conference on Electronic Commerce*, pages 129-139, 1999.
- [33] R. Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In *Proc. STOC 2004*, pages 232-241.
- [34] R. Pass and A. Rosen. Bounded-Concurrent Secure Two-Party Computation in a Constant Number of Rounds. *FOCS 2003*.
- [35] T. Rabin and M. Ben-Or. Verifiable Secret Sharing and Multiparty Protocols with Honest Majority. In *Proc. 21st STOC*, pages 73-85. ACM, 1989.
- [36] O. Reingold, L. Trevisan, and S. P. Vadhan. Notions of Reducibility between Cryptographic Primitives. *TCC 2004*: 1-20.
- [37] P. Rogaway. *The Round Complexity of Secure Protocols*. PhD thesis, MIT, June 1991.
- [38] A. Shamir. How to share a secret. *Commun. ACM*, 22(6):612-613, June 1979.
- [39] S. R. Tate and K. Xu. On garbled circuits and constant round secure function evaluation. CoPS Lab Technical Report 2003-02, University of North Texas, 2003.
- [40] A. C. Yao. How to generate and exchange secrets. In *Proc. 27th FOCS*, pp. 162-167, 1986.