

Semi-homomorphic Encryption and Multiparty Computation

Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias

Department of Computer Science, Aarhus University and CFEM*

Abstract. An additively-homomorphic encryption scheme enables us to compute linear functions of an encrypted input by manipulating only the ciphertexts. We define the relaxed notion of a *semi-homomorphic* encryption scheme, where the plaintext can be recovered as long as the computed function does not increase the size of the input “too much”. We show that a number of existing cryptosystems are captured by our relaxed notion. In particular, we give examples of semi-homomorphic encryption schemes based on lattices, subset sum and factoring. We then demonstrate how semi-homomorphic encryption schemes allow us to construct an *efficient* multiparty computation protocol for arithmetic circuits, UC-secure against a dishonest majority. The protocol consists of a preprocessing phase and an online phase. Neither the inputs nor the function to be computed have to be known during preprocessing. Moreover, the online phase is extremely efficient as it requires *no cryptographic operations*: the parties only need to exchange additive shares and verify information theoretic MACs. Our contribution is therefore twofold: from a theoretical point of view, we can base multiparty computation on a variety of different assumptions, while on the practical side we offer a protocol with better efficiency than any previous solution.

1 Introduction

The fascinating idea of computing on encrypted data can be traced back at least to a seminal paper by Rivest, Adleman and Dertouzos [RAD78] under the name of *privacy homomorphism*. A privacy homomorphism, or *homomorphic encryption scheme* in more modern terminology, is a public-key encryption scheme (G, E, D) for which it holds that $D(E(a) \otimes E(b)) = a \oplus b$, where (\otimes, \oplus) are some group operation in the ciphertext and plaintext space respectively. For instance, if \oplus represents modular addition in some ring, we call such a scheme *additively-homomorphic*. Intuitively a homomorphic encryption scheme enables two parties, say Alice and Bob, to perform *secure computation*: as an example, Alice could encrypt her input a under her public key, send the ciphertext $E(a)$ to Bob; now by the homomorphic property, Bob can compute a ciphertext containing, e.g., $E(a \cdot b + c)$ and send it back to Alice, who can decrypt and learn the

* Center for Research in the Foundations of Electronic Markets, supported by the Danish Strategic Research Council.

result. Thus, Bob has computed a non trivial function of the input a . However, Bob only sees an encryption of a which leaks no information on a itself, assuming that the encryption scheme is secure. Informally we will say that a set of parties P_1, \dots, P_n holding private inputs x_1, \dots, x_n *securely compute* a function of their inputs $y = f(x_1, \dots, x_n)$ if, by running some cryptographic protocol, the honest parties learn the correct output of the function y . In addition, even if (up to) $n - 1$ parties are corrupt and cooperate, they are not able to learn any information about the honest parties' inputs, no matter how they deviate from the specifications of the protocol.

Building *secure multiparty computation (MPC) protocols* for this case of *dishonest majority* is essential for several reasons: First, it is notoriously hard to handle dishonest majority efficiently and it is well known that unconditionally secure solutions do not exist. Therefore, we cannot avoid using some form of public-key technology which is typically much more expensive than the standard primitives used for honest majority (such as secret sharing). Secondly, security against dishonest majority is often the most natural to shoot for in applications, and is of course the only meaningful goal in the significant 2-party case. Thus, finding practical solutions for dishonest majority under reasonable assumptions is arguably the most important research goal with respect to applications of multiparty computation.

While *fully-homomorphic encryption* [Gen09] allows for significant improvement in communication complexity, it would incur a huge computational overhead with current state of the art. In this paper we take a different road: in a nutshell, we relax the requirements of homomorphic encryption so that we can implement it under a variety of assumptions, and we show how this weaker primitive is sufficient for efficient MPC. Our main contributions are:

A framework for semi-homomorphic encryption: we define the notion of a *semi-homomorphic encryption modulo p* , for a modulus p that is input to the key generation. Abstracting from the details, the encryption function is additively homomorphic and will accept any integer x as input plaintext. However, in contrast to what we usually require from a homomorphic cryptosystem, decryption returns the correct result modulo p only if x is numerically small enough. We demonstrate the generality of the framework by giving several examples of known cryptosystems that are semi-homomorphic or can be modified to be so by trivial adjustments. These include: the Okamoto-Uchiyama cryptosystem [OU98]; Paillier cryptosystem [Pai99] and its generalization by Damgård and Jurik [DJ01]; Regev's LWE based cryptosystem [Reg05]; the scheme of Damgård, Geisler and Krøigaard [DGK09] based on a subgroup-decision problem; the subset-sum based scheme by Lyubashevsky, Palacio and Segev [LPS10]; Gentry, Halevi and Vaikuntanathan's scheme [GHV10] based on LWE, and van Dijk, Gentry, Halevi and Vaikuntanathan's scheme [DGHV10] based on the approximate gcd problem. We also show a zero-knowledge protocol for any semi-homomorphic cryptosystem, where a prover, given ciphertext C and public key pk , demonstrates that he knows plaintext x and randomness r such that $C = E_{pk}(x, r)$, and that x furthermore is numerically less than a given bound. We show that using a twist

of the amortization technique of Cramer and Damgård [CD09], one can give u such proofs in parallel where the soundness error is 2^{-u} and the cost per instance proved is essentially 2 encryption operations for both parties. The application of the technique from [CD09] to prove that a plaintext is bounded in size is new and of independent interest.

Information-theoretic “online” MPC: we propose a UC secure [Can01] protocol for arithmetic multiparty computation that, in the presence of a trusted dealer who does not know the inputs, offers information-theoretic security against an adaptive, malicious adversary that corrupts any dishonest majority of the parties. The main idea of the protocol is that the parties will be given additive sharing of multiplicative triples [Bea91], together with information theoretic MACs of their shares – forcing the parties to use the correct shares during the protocol. This online phase is essentially optimal, as no symmetric or public-key cryptography is used, matching the efficiency of passive protocols for honest majority like [BOGW88, CCD88]. Concretely, each party performs $O(n^2)$ multiplications modulo p to evaluate a secure multiplication. This improves on the previous protocol of Damgård and Orlandi (DO) [DO10] where a Pedersen commitment was published for every shared value. Getting rid of the commitments we improve on efficiency (a factor of $\Omega(\kappa)$, where κ is the security parameter) and security (information theoretic against computational). Implementation results for the two-party case indicate about 6 msec per multiplication (see the full version [BDOZ10]), at least an order of magnitude faster than that of DO on the same platform. Moreover, in DO the modulus p of the computation had to match the prime order of the group where the commitments live. Here, we can, however, choose p freely to match the application which typically allows much smaller values of p .

An efficient implementation of the offline phase: we show how to replace the share dealer for the online phase by a protocol based solely on semi-homomorphic encryption¹. Our offline phase is UC-secure against any dishonest majority, and it matches the lower bound for secure computation with dishonest majority of $O(n^2)$ public-key operations per multiplication gate [HIK07]. In the most efficient instantiation, the offline phase of DO requires security of Paillier encryption and hardness of discrete logarithms. Our offline phase only has to assume security of Paillier cryptosystem and achieves similar efficiency: A count of operations suggests that our offline phase is as efficient as DO up to a small constant factor (about 2-3). Preliminary implementation results indicate about 2-3 sec to prepare a multiplication. Since we generalize to any semi-homomorphic scheme including Regev’s scheme, we get the first potentially practical solution for dishonest majority that is believed to withstand a quantum attack. It is not possible to achieve UC security for dishonest majority without set-up assumptions, and our protocol works in the registered public-key model of [BCNP04] where we

¹ The trusted dealer could be implemented using any existing MPC protocol for dishonest majority, but we want to show how we can do it *efficiently* using semi-homomorphic encryption.

assume that public keys for all parties are known, and corrupted parties know their own secret keys.

Related Work: It was shown by Canetti, Lindell, Ostrovsky and Sahai [CLOS02] that secure computation is possible under general assumptions even when considering any corrupted number of parties in a concurrent setting (the UC framework). Their solution is, however, very far from being practical. For computation over Boolean circuits efficient solutions can be constructed from Yao’s garbled circuit technique, see e.g. Pinkas, Schneider, Smart and Williams [PSSW09]. However, our main interest here is arithmetic computation over larger fields or rings, which is a much more efficient approach for applications such as benchmarking or some auction variants. A more efficient solution for the arithmetic case was shown by Cramer, Damgård and Nielsen [CDN01], based on threshold homomorphic encryption. However, it requires distributed key generation and uses heavy public-key machinery throughout the protocol. More recently, Ishai, Prabhakaran and Sahai [IPS09] and the aforementioned DO protocol show more efficient solutions. Although the techniques used are completely different, the asymptotic complexities are similar, but the constants are significantly smaller in the DO solution, which was the most practical protocol proposed so far.

Notation: We let U_S denote the uniform distribution over the set S . We use $x \leftarrow X$ to denote the process of sampling x from the distribution X or, if X is a set, a uniform choice from it.

We say that a function $f : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if $\forall c, \exists n_c$ s.t. if $n > n_c$ then $f(n) < n^{-c}$. We will use $\varepsilon(\cdot)$ to denote an unspecified negligible function.

For $p \in \mathbb{N}$, we represent \mathbb{Z}_p by the numbers $\{-(p-1)/2, \dots, (p-1)/2\}$. If \mathbf{x} is an m -dimensional vector, $\|\mathbf{x}\|_\infty := \max(|x_1|, \dots, |x_m|)$. Unless differently specified, all the logarithms are in base 2.

As a general convention: lowercase letters a, b, c, \dots represent integers and capital letters A, B, C, \dots ciphertexts. Bold lowercase letters $\mathbf{r}, \mathbf{s}, \dots$ are vectors and bold capitals $\mathbf{M}, \mathbf{A}, \dots$ are matrices. We call κ the computational security parameter and u the statistical security parameter. In practice u can be set to be much smaller than κ , as it does not depend on the computing power of the adversary.

2 The Framework for Semi-homomorphic Encryption

In this section we introduce a framework for public-key cryptosystems, that satisfy a relaxed version of the *additive homomorphic property*. Let $\text{PKE} = (G, E, D)$ be a tuple of algorithms where:

$G(1^\kappa, p)$ is a randomized algorithm that takes as input a security parameter κ and a modulus p^2 ; It outputs a public/secret key pair (pk, sk) and a set of parameters $\mathbb{P} = (p, M, R, \mathcal{D}_\sigma^d, \mathbb{G})$. Here, M, R are integers, \mathcal{D}_σ^d is the description

² In the framework there are no restrictions for the choice of p ; however in the next sections p will always be chosen to be a prime.

of a randomized algorithm producing as output d -vectors with integer entries (to be used as randomness for encryption). We require that except with negligible probability, \mathcal{D}_σ^d will always output \mathbf{r} with $\|\mathbf{r}\|_\infty \leq \sigma$, for some $\sigma < R$ that may depend on κ . Finally, \mathbb{G} is the abelian group where the ciphertexts belong (written in additive notation). For practical purposes one can think of M and R to be of size super-polynomial in κ , and p and σ as being much smaller than M and R respectively. We will assume that every other algorithm takes as input the parameters \mathbb{P} , without specifying this explicitly.

$E_{pk}(x, \mathbf{r})$ is a deterministic algorithm that takes as input an integer $x \in \mathbb{Z}$ and a vector $\mathbf{r} \in \mathbb{Z}^d$ and outputs a ciphertext $C \in \mathbb{G}$. We sometimes write $E_{pk}(x)$ when it is not important to specify the randomness explicitly. Given $C_1 = E_{pk}(x_1, \mathbf{r}_1)$, $C_2 = E_{pk}(x_2, \mathbf{r}_2)$ in \mathbb{G} , we have $C_1 + C_2 = E_{pk}(x_1 + x_2, \mathbf{r}_1 + \mathbf{r}_2)$. In other words, $E_{pk}(\cdot, \cdot)$ is a homomorphism from $(\mathbb{Z}^{d+1}, +)$ to $(\mathbb{G}, +)$. Given some τ and ρ we call C a (τ, ρ) -ciphertext if there exists x, \mathbf{r} with $|x| \leq \tau$ and $\|\mathbf{r}\|_\infty \leq \rho$ such that $C = E_{pk}(x, \mathbf{r})$. Note that given a ciphertext τ and ρ are not unique. When we refer to a (τ, ρ) -ciphertext, τ and ρ should be interpreted as an upper limit to the size of the message and randomness contained in the ciphertext.

$D_{sk}(C)$ is a deterministic algorithm that takes as input a ciphertext $C \in \mathbb{G}$ and outputs $x' \in \mathbb{Z}_p \cup \{\perp\}$.

We say that a semi-homomorphic encryption scheme PKE is *correct* if, $\forall p$:

$$\Pr[(pk, sk, \mathbb{P}) \leftarrow \mathbb{G}(1^\kappa, p), x \in \mathbb{Z}, |x| \leq M; \mathbf{r} \in \mathbb{Z}^d, \|\mathbf{r}\|_\infty \leq R : D_{sk}(E_{pk}(x, \mathbf{r})) \neq x \bmod p] < \varepsilon(\kappa)$$

where the probabilities are taken over the random coins of \mathbb{G} and \mathbb{E} .

We now define the IND-CPA security game for a semi-homomorphic cryptosystem. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a PPT TM, then we run the following experiment:

$$\begin{aligned} (pk, sk, \mathbb{P}) &\leftarrow \mathbb{G}(1^\kappa, p) \\ (m_0, m_1, \text{state}) &\leftarrow \mathcal{A}_1(1^\kappa, pk) \text{ with } m_0, m_1 \in \mathbb{Z}_p \\ b &\leftarrow \{0, 1\}, C \leftarrow E_{pk}(m_b), b' \leftarrow \mathcal{A}_2(1^\kappa, \text{state}, C) \end{aligned}$$

We define the advantage of \mathcal{A} as $\text{Adv}^{\text{CPA}}(\mathcal{A}, \kappa) = |\Pr[b = b'] - 1/2|$, where the probabilities are taken over the random choices of $\mathbb{G}, \mathbb{E}, \mathcal{A}$ in the above experiment. We say that PKE is IND-CPA *secure* if \forall PPT \mathcal{A} , $\text{Adv}^{\text{CPA}}(\mathcal{A}, \kappa) < \varepsilon(\kappa)$.

Next, we discuss the motivation for the way this framework is put together: when in the following, honest players encrypt data, plaintext x will be chosen in \mathbb{Z}_p and the randomness \mathbf{r} according to \mathcal{D}_σ^d . This ensures IND-CPA security and also that such data can be decrypted correctly, since by assumption on \mathcal{D}_σ^d , $\|\mathbf{r}\|_\infty \leq \sigma \leq R$. However, we also want that a (possibly dishonest) player P_i is committed to x by publishing $C = E_{pk}(x, \mathbf{r})$. We are not able to force a player to choose x in \mathbb{Z}_p , nor that \mathbf{r} is sampled with the correct distribution. But our zero-knowledge protocols *can* ensure that C is a (τ, ρ) -ciphertext, for concrete values of τ, ρ . If $\tau < M, \rho < R$, then correctness implies that C commits P_i to $x \bmod p$, even if x, \mathbf{r} may not be uniquely determined from C .

2.1 Examples of Semi-homomorphic Encryption

Regev’s cryptosystem [Reg05] is parametrized by p, q, m and α , and is given by (G, E, D) . A variant of the system was also given in [BD10], where parameters are chosen slightly differently than in the original. In both [Reg05] and [BD10] only a single bit was encrypted, it is quite easy, though, to extend it to elements of a bigger ring. It is this generalized version of the variant in [BD10] that we describe here. All calculations are done in \mathbb{Z}_q . Key generation $G(1^\kappa)$ is done by sampling $\mathbf{s} \in \mathbb{Z}_q^n$ and $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ uniformly at random and $\mathbf{x} \in \mathbb{Z}_q^m$ from a discrete Gaussian distribution with mean 0 and standard deviation $\frac{q\alpha}{\sqrt{2\pi}}$. We then have the key pair $(pk, sk) = ((\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{x}), \mathbf{s})$. Encryption of a message $\gamma \in \mathbb{Z}_p$ is done by sampling a uniformly random vector $\mathbf{r} \in \{-1, 0, 1\}^m$. A ciphertext C is then given by $C = E_{pk}(\gamma, \mathbf{r}) = (\mathbf{a}, b) = (\mathbf{A}^T \mathbf{r}, (\mathbf{A}\mathbf{s} + \mathbf{x})^T \mathbf{r} + \gamma \lfloor q/p \rfloor)$. Decryption is given by $D_{sk}(C) = \lfloor (b - \mathbf{s}^T \mathbf{a}) \cdot p/q \rfloor$. Regev’s cryptosystem works with a decryption error, which can, however, be made negligibly small when choosing the parameters.

Fitting the cryptosystem to the framework is quite straight forward. The group $\mathbb{G} = \mathbb{Z}_q^n \times \mathbb{Z}_q$ and p is just the same. The distribution \mathcal{D}_σ^d from which the randomness \mathbf{r} is taken is the uniform distribution over $\{-1, 0, 1\}^m$, that is $d = m$ and $\sigma = 1$. Given two ciphertexts (\mathbf{a}, b) and (\mathbf{a}', b') we define addition to be $(\mathbf{a} + \mathbf{a}', b + b')$. With this definition it follows quite easily that the homomorphic property holds. Due to the choices of message space and randomness distribution in Regev’s cryptosystem, we will always have that the relation $M = Rp/2$ should hold. How M can be chosen, and thereby also R , depends on all the original parameters of the cryptosystem. First assume that $q \cdot \alpha = \sqrt[d]{q}$ with $d > 1$. Furthermore we will need that $p \leq q/(4\sqrt[d]{q})$ for some constant $c < d$. Then to bound M we should have first that $M < q/(4p)$ and secondly that $M < p\sqrt[d]{q}/(2m)$ for some $s > cd/(d-c)$. Obtaining these bounds requires some tedious computation which we leave out here.

In *Paillier’s cryptosystem* [Pai99] the secret key is two large primes p_1, p_2 , the public key is $N = p_1 p_2$, and the encryption function is $E_{pk}(x, r) = (N + 1)^x r^N \bmod N^2$ where $x \in \mathbb{Z}_N$ and r is random in $\mathbb{Z}_{N^2}^*$. The decryption function D'_{sk} reconstructs correctly any plaintext in \mathbb{Z}_N , and to get a semi-homomorphic scheme modulo p , we simply redefine the decryption as $D(c) = D'(c) \bmod p$. It is not hard to see that we get a semi-homomorphic scheme with $M = (N - 1)/2, R = \infty, d = 1, \mathcal{D}_\sigma^d = U_{\mathbb{Z}_{N^2}^*}, \sigma = \infty$ and $\mathbb{G} = \mathbb{Z}_{N^2}^*$. In particular, note that we do not need to bound the size of the randomness, hence we set $\sigma = R = \infty$.

The cryptosystem looks syntactically a bit different from our definition which writes \mathbb{G} additively, while $\mathbb{Z}_{N^2}^*$ is usually written with multiplicative notation; also for Paillier we have $E_{pk}(x, r) + E_{pk}(x', r) = E_{pk}(x + x', r \cdot r')$ and not $E_{pk}(x + x', r + r')$. However, this makes no difference in the following, except that it actually makes some of the zero-knowledge protocols simpler (more details in Section 2.2). It is easy to see that the generalization of Paillier in [DJ01] can be modified in a similar way to be semi-homomorphic.

In the full paper [BDOZ10] we show how several other cryptosystems are semi-homomorphic.

2.2 Zero-Knowledge Proofs

We present two zero-knowledge protocols, Π_{PoPK} , Π_{PoCM} where a prover P proves to a verifier V that some ciphertexts are correctly computed and that some ciphertexts satisfy a multiplicative relation respectively. Π_{PoPK} has (amortized) complexity $O(\kappa + u)$ bits per instance proved, where the soundness error is 2^{-u} . Π_{PoCM} has complexity $O(\kappa u)$. We also show a more efficient version of Π_{PoCM} that works only for Paillier encryption, with complexity $O(\kappa + u)$. Finally, in the full paper [BDOZ10], we define the *multiplication security* property that we conjecture is satisfied for all our example cryptosystems after applying a simple modification. We show that assuming this property, Π_{PoCM} can be replaced by a different check that has complexity $O(\kappa + u)$.

Π_{PoPK} and Π_{PoCM} will both be of the standard 3-move form with a random u -bit challenge, and so they are honest verifier zero-knowledge. To achieve zero-knowledge against an arbitrary verifier standard techniques can be used. In particular, in our MPC protocol we will assume – only for the sake of simplicity – a functionality $\mathcal{F}_{\text{RAND}}$ that generates random challenges on demand. The $\mathcal{F}_{\text{RAND}}$ functionality is specified in detail in the full paper [BDOZ10] and can be implemented in our key registration model using only semi-homomorphic encryption. In the protocols both prover and verifier will have public keys pk_P and pk_V . By $E_P(a, \mathbf{r})$ we denote an encryption under pk_P , similarly for $E_V(a, \mathbf{r})$.

We emphasize that the zero-knowledge property of our protocols does not depend on IND-CPA security of the cryptosystem, instead it follows from the homomorphic property and the fact that the honest prover creates, for the purpose of the protocol, some auxiliary ciphertexts containing enough randomness to hide the prover’s secrets.

Proof of Plaintext Knowledge. Π_{PoPK} takes as common input u ciphertexts $C_k, k = 1, \dots, u$. If these are (τ, ρ) -ciphertexts, the protocol is complete and statistical zero-knowledge. The protocol is sound in the following sense: assuming that pk_P is well-formed, if P is corrupt and can make V accept with probability larger than 2^{-u} , then all the C_k are $(2^{2u+\log u}\tau, 2^{2u+\log u}\rho)$ -ciphertexts. The protocol is also a proof of knowledge with knowledge error 2^{-u} that P knows correctly formed plaintexts and randomness for all the C_k ’s.

In other words, Π_{PoPK} is a ZKPoK for the following relation, *except* that zero-knowledge and completeness only hold if the C_k ’s satisfy the stronger condition of being (τ, ρ) -ciphertexts. However, this is no problem in the following: the prover will always create the C_k ’s himself and can therefore ensure that they are correctly formed if he is honest.

$$\begin{aligned}
 R_{\text{PoPK}}^{(u, \tau, \rho)} = \{ (x, w) \mid & \quad x = (pk_P, C_1, \dots, C_u); \\
 & \quad w = ((x_1, \mathbf{r}_1), \dots, (x_u, \mathbf{r}_u)) : C_k = E_P(x_k, \mathbf{r}_k), \\
 & \quad |x_k| \leq 2^{2u+\log u}\tau, \|\mathbf{r}_k\|_\infty \leq 2^{2u+\log u}\rho \}
 \end{aligned}$$

We use the approach of [CD09] to get small amortized complexity of the zero-knowledge proofs, and thereby gaining efficiency by performing the proofs on u

simultaneous instances. In the following we define $m = 2u - 1$, furthermore \mathbf{M}_e is an $m \times u$ matrix constructed given a uniformly random vector $\mathbf{e} = (e_1, \dots, e_u) \in \{0, 1\}^u$. Specifically the (i, k) -th entry $\mathbf{M}_{e,i,k}$ is given by $\mathbf{M}_{e,i,k} = e_{i-k+1}$ for $1 \leq i - k + 1 \leq u$ and 0 otherwise. By $\mathbf{M}_{e,i}$ we denote the i -th row of \mathbf{M}_e . The protocol can be seen in Figure 1. Completeness and zero-knowledge follow by standard arguments that can be found in the full paper [BDOZ10]. Here we argue soundness which is the more interesting case: Assume we are given any prover P^* , and consider the case where P^* can make V accept for both \mathbf{e} and \mathbf{e}' , $\mathbf{e} \neq \mathbf{e}'$, by sending \mathbf{z} , \mathbf{z}' , \mathbf{T} and \mathbf{T}' respectively. We now have the following equation:

$$(\mathbf{M}_e - \mathbf{M}_{e'})\mathbf{c} = (\mathbf{d} - \mathbf{d}') \tag{1}$$

What we would like is to find $\mathbf{x} = (x_1, \dots, x_u)$ and $\mathbf{R} = (\mathbf{r}_1, \dots, \mathbf{r}_u)$ such that $C_k = E_P(x_k, \mathbf{r}_k)$. We can do this by viewing (1) as a system of linear equations. First let j be the biggest index such that $\mathbf{e}_j \neq \mathbf{e}'_j$. Now look at the $u \times u$ submatrix of $\mathbf{M}_e - \mathbf{M}_{e'}$ given by the rows j through $j + u$ both included. This is an upper triangular matrix with entries in $\{-1, 0, 1\}$ and $\mathbf{e}_j - \mathbf{e}'_j \neq 0$ on a diagonal. Now remember the form of the entries in the vectors \mathbf{c} , \mathbf{d} and \mathbf{d}' , we have $C_k = E_P(x_k, \mathbf{r}_k)$, $D_k = E_P(z_k, \mathbf{t}_k)$, $D'_k = E_P(z'_k, \mathbf{t}'_k)$. We can now directly solve the equations for the x_k 's and the \mathbf{r}_k 's by starting with C_u and going up. We give examples of the first few equations (remember we are going bottom up). For simplicity we will assume that all entries in $\mathbf{M}_e - \mathbf{M}_{e'}$ will be 1.

$$\begin{aligned} E_P(x_u, \mathbf{r}_u) &= E_P(z_{u+j} - z'_{u+j}, \mathbf{t}_{u+j} - \mathbf{t}'_{u+j}) \\ E_P(x_{u-1}, \mathbf{r}_{u-1}) + E_P(x_u, \mathbf{r}_u) &= E_P(z_{u+j-1} - z'_{u+j-1}, \mathbf{t}_{u+j-1} - \mathbf{t}'_{u+j-1}) \\ &\vdots \end{aligned}$$

Since we know all values used on the right hand sides and since the cryptosystem used is additively homomorphic, it should now be clear that we can find x_k and \mathbf{r}_k such that $C_k = E_P(x_k, \mathbf{r}_k)$. A final note should be said about what we can guarantee about the sizes of x_k and \mathbf{r}_k . Knowing that $|z_i| \leq 2^{u-1+\log u \tau}$, $|z'_i| \leq 2^{u-1+\log u \tau}$, $\|\mathbf{t}_i\|_\infty \leq 2^{u-1+\log u \rho}$ and $\|\mathbf{t}'_i\|_\infty \leq 2^{u-1+\log u \rho}$ we could potentially have that C_1 would become a $(2^{2u+\log u \tau}, 2^{2u+\log u \rho})$ ciphertext. Thus this is what we can guarantee.

Proof of Correct Multiplication. $\Pi_{\text{POCM}}(u, \tau, \rho)$ takes as common input u triples of ciphertexts (A_k, B_k, C_k) for $k = 1, \dots, u$, where A_k is under pk_P and B_k and C_k are under pk_V (and so are in the group \mathbb{G}_V). If P is honest, he will know a_k and $a_k \leq \tau$. Furthermore P has created C_k as $C_k = a_k B_k + E_V(r_k, \mathbf{t}_k)$, where $E_V(r_k, \mathbf{t}_k)$ is a random $(2^{3u+\log u \tau^2}, 2^{3u+\log u \tau \rho})$ -ciphertext. Under these assumptions the protocol is zero-knowledge.

Jumping ahead, we note that in the context where the protocol will be used, it will always be known that B_k in every triple is a $(2^{2u+\log u \tau}, 2^{2u+\log u \rho})$ -ciphertext, as a result of executing Π_{POPK} . The choice of sizes for $E_V(r_k, \mathbf{t}_k)$ then ensures that C_k is statistically close to a random $(2^{3u+\log u \tau^2}, 2^{3u+\log u \tau \rho})$ -ciphertext, and so reveals no information on a_k to V .

Subprotocol Π_{PoPK} : Proof of Plaintext Knowledge**PoPK**(u, τ, ρ):

1. The input is u ciphertexts $\{C_k = \mathbf{E}_P(x_k, \mathbf{r}_k)\}_{k=1}^u$. We define the vectors $\mathbf{c} = (C_1, \dots, C_u)$ and $\mathbf{x} = (x_1, \dots, x_u)$ and the matrix $\mathbf{R} = (\mathbf{r}_1, \dots, \mathbf{r}_u)$, where the \mathbf{r}_k 's are rows.
2. P constructs m ($2^{u-1+\log u} \tau, 2^{u-1+\log u} \rho$)-ciphertexts $\{A_i = \mathbf{E}_P(y_i, \mathbf{s}_i)\}_{i=1}^m$, and sends them to V . We define vectors \mathbf{a} and \mathbf{y} and matrix \mathbf{S} as above.
3. V chooses a uniformly random vector $\mathbf{e} = (e_1, \dots, e_u) \in \{0, 1\}^u$, and sends it to P .
4. Finally P computes and sends $\mathbf{z} = \mathbf{y} + \mathbf{M}_e \cdot \mathbf{x}$ and $\mathbf{T} = \mathbf{S} + \mathbf{M}_e \cdot \mathbf{R}$ to V .
5. V checks that $\mathbf{d} = \mathbf{a} + \mathbf{M}_e \cdot \mathbf{c}$ where $\mathbf{d} = (\mathbf{E}_P(z_1, \mathbf{t}_1), \dots, \mathbf{E}_P(z_m, \mathbf{t}_m))$. Furthermore, V checks that $|z_i| \leq 2^{u-1+\log u} \tau$ and $\|\mathbf{t}_i\|_\infty \leq 2^{u-1+\log u} \rho$.

Fig. 1. Proof of Plaintext KnowledgeSubprotocol Π_{PoCM} : Proof of Correct Multiplication**PoCM**(u, τ, ρ):

1. The input is u triples of ciphertexts $\{(A_k, B_k, C_k)\}_{k=1}^u$, where $A_k = \mathbf{E}_P(a_k, \mathbf{h}_k)$ and $C_k = a_k B_k + \mathbf{E}_V(r_k, \mathbf{t}_k)$.
2. P constructs u uniformly random ($2^{3u-1+\log u} \tau, 2^{3u-1+\log u} \rho$)-ciphertexts $D_k = \mathbf{E}_P(d_k, \mathbf{s}_k)$ and u ciphertexts $F_k = d_k B_k + \mathbf{E}_V(f_k, \mathbf{y}_k)$, where $\mathbf{E}_V(f_k, \mathbf{y}_k)$ are uniformly random ($2^{4u-1+\log u} \tau^2, 2^{4u-1+\log u} \tau \rho$)-ciphertexts.
3. V chooses u uniformly random bits e_k and sends them to P .
4. P returns $\{(z_k, \mathbf{v}_k)\}_{k=1}^u$ and $\{(x_k, \mathbf{w}_k)\}_{k=1}^u$ to V . Here $z_k = d_k + e_k a_k$, $\mathbf{v}_k = \mathbf{s}_k + e_k \mathbf{h}_k$, $x_k = f_k + e_k r_k$ and $\mathbf{w}_k = \mathbf{y}_k + e_k \mathbf{t}_k$.
5. V checks that $D_k + e_k A_k = \mathbf{E}_P(z_k, \mathbf{v}_k)$ and that $F_k + e_k C_k = z_k B_k + \mathbf{E}_V(x_k, \mathbf{w}_k)$. Furthermore, he checks that $|z_k| \leq 2^{3u-1+\log u} \tau$, $\|\mathbf{v}_k\|_\infty \leq 2^{3u-1+\log u} \rho$, $|x_k| \leq 2^{4u-1+\log u} \tau^2$ and $\|\mathbf{w}_k\|_\infty \leq 2^{4u-1+\log u} \tau \rho$.
6. Step 2-5 is repeated in parallel u times.

Fig. 2. Proof of Correct Multiplication

Summarizing, Π_{PoCM} is a ZKPoK for the relation (under the assumption that pk_P, pk_V are well-formed):

$$\begin{aligned}
 R_{\text{PoCM}}^{(u, \tau, \rho)} = \{(x, w) \mid & x = (pk_P, pk_V, (A_1, B_1, C_1), \dots, (A_u, B_u, C_u)); \\
 & w = ((a_1, \mathbf{h}_1, r_1, \mathbf{t}_1), \dots, (a_u, \mathbf{h}_u, r_u, \mathbf{t}_u)) : \\
 & A_k = \mathbf{E}_P(a_k, \mathbf{h}_k), B_k \in \mathbb{G}_V, C_k = a_k B_k + \mathbf{E}_V(r_k, \mathbf{t}_k), \\
 & |a_k| \leq 2^{3u+\log u} \tau, \|\mathbf{h}_k\|_\infty \leq 2^{3u+\log u} \rho, \\
 & |r_k| \leq 2^{4u+\log u} \tau^2, \|\mathbf{t}_k\|_\infty \leq 2^{4u+\log u} \tau \rho\}
 \end{aligned}$$

The protocol can be seen in Figure 2. Note that Step 6 could also be interpreted as choosing e_k as a u -bit vector instead, thereby only calling $\mathcal{F}_{\text{RAND}}$ once. Completeness, soundness and zero-knowledge follow by standard arguments that can be found in the full paper [BDOZ10].

Zero-Knowledge Protocols for Paillier. For the particular case of Paillier encryption, Π_{PoPK} can be used as it is, except that there is no bound required on the randomness, instead all random values used in encryptions are expected to be in $\mathbb{Z}_{N^2}^*$. Thus, the relations to prove will only require that the random values are in $\mathbb{Z}_{N^2}^*$ and this is also what the verifier should check in the protocol.

For Π_{PoCM} we sketch a version that is more efficient than the above, using special properties of Paillier encryption. In order to improve readability, we depart here from the additive notation for operations on ciphertexts, since multiplicative notation is usually used for Paillier. In the following, let $pk_V = N$. Note first that based on such a public key, one can define an unconditionally hiding commitment scheme with public key $g = E_V(0)$. To commit to $a \in \mathbb{Z}_N$, one sends $\text{com}(a, r) = g^a r^N \bmod N$, for random $r \in \mathbb{Z}_{N^2}^*$. One can show that the scheme is binding assuming it is hard to extract N -th roots modulo N^2 (which must be the case if Paillier encryption is secure).

We restate the relation $R_{\text{PoCM}}^{(u, \tau, \rho)}$ from above as it will look for the Paillier case, in multiplicative notation and without bounds on the randomness:

$$\begin{aligned}
 R_{\text{PoCM}, \text{Paillier}}^{(\tau, \rho)} = \{ (x, w) \mid & x = (pk_P, pk_V, (A_1, B_1, C_1), \dots, (A_u, B_u, C_u)); \\
 & w = ((a_1, h_1, r_1, t_1), \dots, (a_u, h_u, r_u, t_u)) : \\
 & A_k = E_P(a_k, h_k), B_k \in \mathbb{Z}_{N^2}, C_k = B_k^{a_k} \cdot E_V(r_k, t_k), \\
 & |a_k| \leq 2^{2u + \log u \tau}, |r_k| \leq 2^{5u + 2 \log u \tau^2} \}
 \end{aligned}$$

The idea for the proof of knowledge for this relation is now to ask the prover to also send commitments $\Psi_k = \text{com}(a_k, \alpha_k), \Phi_k = \text{com}(r_k, \beta_k), k = 1 \dots u$ to the r_k 's and a_k 's. Now, the prover must first provide a proof of knowledge that for each k : 1) the same bounded size value is contained in both A_k and Ψ_k , and that 2) a bounded size value is contained in Φ_k . The proof for $\{\Phi_k\}$ is simply Π_{PoPK} since a commitment has the same form as an encryption (with $(N + 1)$ replaced by g). The proof for $\{\Psi_k, A_k\}$ is made of two instances of Π_{PoPK} run in parallel, using the same challenge e and responses z_i in both instances. Finally, the prover must show that C_k can be written as $C_k = B_k^{a_k} \cdot E_V(r_k, t_k)$, where a_k is the value contained in Ψ_k and r_k is the value in Φ_k . Since all commitments and ciphertexts live in the same group $\mathbb{Z}_{N^2}^*$, where $pk_V = N$, we can do this efficiently using a variant of a protocol from [CDN01]. The resulting protocol is shown in Figure 3.

Completeness of the protocol in steps 1-4 of Figure 3 is straightforward by inspection. Honest verifier zero-knowledge follows by the standard argument: choose e and the prover's responses uniformly in their respective domains and use the equations checked by the verifier to compute a matching first message D, X, Y . This implies completeness and honest verifier zero-knowledge for the overall protocol, since the subprotocols in steps 2 and 3 have these properties as well.

Subprotocol Π_{PoCM} : Proof of Correct Multiplication (only for Paillier)

1. P sends $\Psi_k = \text{com}(a_k, \alpha_k), \Phi_k = \text{com}(r_k, \beta_k), k = 1, \dots, u$ to the verifier.
2. P uses Π_{PoPK} on Φ_k to prove that, even if P is corrupted, each Φ_k contains a value r_k with $|r_k| \leq 2^{5u+2 \log u} \tau^2$.
3. P uses Π_{PoPK} in parallel on (A_k, Ψ_k) (where V uses the same \mathbf{e} in both runs) to prove that, even if P is corrupted, Ψ_k and A_k contains the same value a_k and $|a_k| \leq 2^{2u+1 \log u} \tau$.
4. To show that the C_k 's are well-formed, we do the following for each k :
 - (a) P picks random $x, y, v, \gamma, \delta \leftarrow \mathbb{Z}_{N^2}^*$ and sends $D = B_k^x \text{E}_V(y, v), X = \text{com}(x, \gamma_x), Y = \text{com}(y, \gamma_y)$ to V .
 - (b) V sends a random u -bit challenge e .
 - (c) P computes $z_a = x + ea_k \bmod N, z_r = y + er_k \bmod N$.
He also computes q_a, q_r , where $x + ea = q_a N + z_a, y + er_k = q_r N + z_r$ ^a.
 P sends $z_a, z_r, w = v s_k^e B_k^{q_a} \bmod N^2, \delta_a = \gamma_x \alpha_k^e g^{q_a} \bmod N^2$, and $\delta_r = \gamma_y \beta_k^e g^{q_r} \bmod N^2$ to V .
 - (d) V accepts if $DC_k^e = B_k^{z_a} \text{E}_V(z_r, w) \bmod N^2 \wedge X \Psi_k^e = \text{com}(z_a, \delta_a) \bmod N^2 \wedge Y \Phi_k^e = \text{com}(z_r, \delta_r) \bmod N^2$.

^a Since g and B_k do not have order N , we need to explicitly handle the quotients q_a and q_r , in order to move the “excess multiples” of N into the randomness parts of the commitments and ciphertext.

Fig. 3. Proof of Correct Multiplication for Paillier encryption

Finally, soundness follows by assuming we are given correct responses in step 7 to two different challenges. From the equations checked by the verifier, we can then easily compute $a_k, \alpha_k, r_k, \beta_k, s_k$ such that $\Psi_k = \text{com}(a_k, \alpha_k), \Phi_k = \text{com}(r_k, \beta_k), C_k = B_k^{a_k} \text{E}_V(r_k, s_k)$. Now, by soundness of the protocols in steps 2 and 3, we can also compute bounded size values a'_k, r'_k that are contained in Ψ_k, Φ_k . By the binding property of the commitment scheme, we have $r'_k = r_k, a'_k = a_k$ except with negligible probability, so we have a witness as required in the specification of the relation.

3 The Online Phase

Our goal is to implement reactive arithmetic multiparty computation over \mathbb{Z}_p for a prime p of size super-polynomial in the statistical security parameter u . The (standard) ideal functionality $\mathcal{F}_{\text{AMPC}}$ that we implement can be seen in Figure 6. We assume here that the parties already have a functionality for synchronous³, secure communication and broadcast.

³ A malicious adversary can always stop sending messages and, in any protocol for dishonest majority, all parties are required for the computation to terminate. Without synchronous channels the honest parties might wait forever for the adversary to send his messages. Synchronous channels guarantee that the honest parties can detect that the adversary is not participating anymore and therefore they can abort the protocol. If termination is not required, the protocol can be implemented over an asynchronous network instead.

We first present a protocol for an *online phase* that assumes access to a functionality $\mathcal{F}_{\text{TRIP}}$ which we later show how to implement using an *offline protocol*. The online phase is based on a representation of values in \mathbb{Z}_p that are shared additively where shares are authenticated using information theoretic message authentication codes (MACs). Before presenting the protocol we introduce how the MACs work and how they are included in the representation of a value in \mathbb{Z}_p . Furthermore, we argue how one can compute with these representations as we do with simple values, and in particular how the relation to the MACs are maintained.

In the rest of this section, all additions and multiplications are to be read modulo p , even if not specified. The number of parties is denoted by n , and we call the parties P_1, \dots, P_n .

3.1 The MACs

A key K in this system is a random pair $K = (\alpha, \beta) \in \mathbb{Z}_p^2$, and the authentication code for a value $a \in \mathbb{Z}_p$ is $\text{MAC}_K(a) = \alpha a + \beta \pmod p$.

We will apply the MACs by having one party P_i hold a , $\text{MAC}_K(a)$ and another party P_j holding K . The idea is to use the MAC to prevent P_i from lying about a when he is supposed to reveal it to P_j . It will be very important in the following that if we keep α constant over several different MAC keys, then one can add two MACs and get a valid authentication code for the sum of the two corresponding messages. More concretely, two keys $K = (\alpha, \beta), K' = (\alpha', \beta')$ are said to be *consistent* if $\alpha = \alpha'$. For consistent keys, we define $K + K' = (\alpha, \beta + \beta')$ so that it holds that $\text{MAC}_K(a) + \text{MAC}_{K'}(a') = \text{MAC}_{K+K'}(a + a')$.

The MACs will be used as follows: we give to P_i several different values m_1, m_2, \dots with corresponding MACs $\gamma_1, \gamma_2, \dots$ computed using keys $K_i = (\alpha, \beta_i)$ that are random but consistent. It is then easy to see that if P_i claims a false value for any of the m_i 's (or a linear combination of them) he can guess an acceptable MAC for such a value with probability at most $1/p$.

3.2 The Representation and Linear Computation

To represent a value $a \in \mathbb{Z}_p$, we will give a share a_i to each party P_i . In addition, P_i will hold MAC keys $K_{a_1}^i, \dots, K_{a_n}^i$. He will use key $K_{a_j}^i$ to check the share of P_j , if we decide to make a public. Finally, P_i also holds a set of authentication codes $\text{MAC}_{K_{a_i}^j}(a_i)$. We will denote $\text{MAC}_{K_{a_i}^j}(a_i)$ by $m_j(a_i)$ from now on. Party P_i will use $m_j(a_i)$ to convince P_j that a_i is correct, if we decide to make a public. Summing up, we have the following way of representing a :

$$[a] = [\{a_i, \{K_{a_j}^i, m_j(a_i)\}_{j=1}^n\}_{i=1}^n]$$

where $\{a_i, \{K_{a_j}^i, m_j(a_i)\}_{j=1}^n\}$ is the information held privately by P_i , and where we use $[a]$ as shorthand when it is not needed to explicitly talk about the shares and MACs. We say that $[a] = [\{a_i, \{K_{a_j}^i, m_j(a_i)\}_{j=1}^n\}_{i=1}^n]$ is *consistent*, with $a = \sum_i a_i$, if $m_j(a_i) = \text{MAC}_{K_{a_i}^j}(a_i)$ for all i, j . Two representations

$$[a] = [\{a_i, \{K_{a_j}^i, m_j(a_i)\}_{j=1}^n\}_{i=1}^n], \quad [a'] = [\{a'_i, \{K_{a'_j}^i, m_j(a'_i)\}_{j=1}^n\}_{i=1}^n]$$

Opening: We can reliably open a consistent representation to P_j : each P_i sends $a_i, m_j(a_i)$ to P_j . P_j checks that $m_j(a_i) = \text{MAC}_{K_{a_i}^j}(a_i)$ and broadcasts *OK* or *fail* accordingly. If all is OK, P_j computes $a = \sum_i a_i$, else we abort. We can modify this to opening a value $[a]$ to all parties, by opening as above to every P_j .

Addition: Given two key-consistent representations as above we get that

$$[a + a'] = [\{a_i + a'_i, \{K_{a_j}^i + K_{a'_j}^i, m_j(a_i) + m_j(a'_i)\}_{j=1}^n\}_{i=1}^n]$$

is a consistent representation of $a+a'$. This new representation can be computed only by local operations.

Multiplication by constants: In a similar way, we can multiply a public constant δ “into” a representation. This is written $\delta[a]$ and is taken to mean that all parties multiply their shares, keys and MACs by δ . This gives a consistent representation $[\delta a]$.

Addition of constants: We can add a public constant δ into a representation. This is written $\delta + [a]$ and is taken to mean that P_1 will add δ to his share a_1 . Also, each P_j will replace his key $K_{a_1}^j = (\alpha_1^j, \beta_{a_1}^j)$ by $K_{a_1+\delta}^j = (\alpha_1^j, \beta_{a_1}^j - \delta\alpha_1^j)$. This will ensure that the MACs held by P_1 will now be valid for the new share $a_1 + \delta$, so we now have a consistent representation $[a + \delta]$.

Fig. 4. Operations on $[\cdot]$ -representations

are said to be *key-consistent* if they are both consistent, and if for all i, j the keys $K_{a_j}^i, K_{a'_j}^i$ are consistent. We will want *all* representations in the following to be key-consistent: this is ensured by letting P_i use the same α_j -value in keys towards P_j throughout. Therefore the notation $K_{a_j}^i = (\alpha_j^i, \beta_{a_j}^i)$ makes sense and we can compute with the representations, as detailed in Figure 4.

3.3 Triples and Multiplication

For multiplication and input sharing we will need both random single values $[a]$ and triples $[a], [b], [c]$ where a, b are random and $c = ab \bmod p$. Also, we assume that all singles and triples we ever produce are key consistent, so that we can freely add them together. More precisely, we assume we have access to an ideal functionality $\mathcal{F}_{\text{TRIP}}$ providing us with the above. This is presented in Figure 5.

The principle in the specification of the functionality is that the environment is allowed to specify all the data that the corrupted parties should hold, including all shares of secrets, keys and MACs. Then, the functionality chooses the secrets to be shared and constructs the data for honest parties so it is consistent with the secrets and the data specified by the environment.

Thanks to this functionality we are also able to compute multiplications in the following way: If the parties hold two key-consistent representations $[x], [y]$, we can use one precomputed key-consistent triple $[a], [b], [c]$ (with $c = ab$) to compute a new representation of $[xy]$.

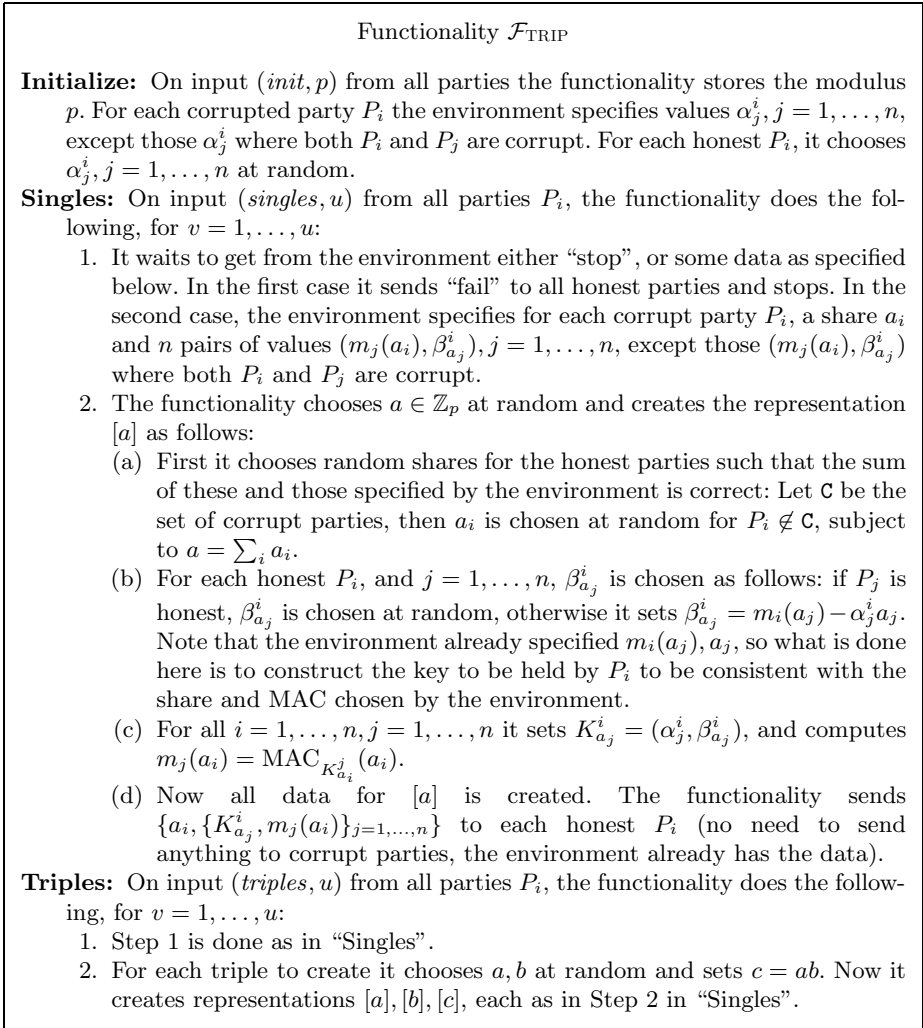


Fig. 5. The ideal functionality for making singles $[a]$ and triples $[a], [b], [c]$

To compute $[xy]$ we first open $[x] - [a]$ to get a value ε , and $[y] - [b]$ to get δ . Then, we have $xy = (a + \varepsilon)(b + \delta) = c + \varepsilon b + \delta a + \varepsilon \delta$. Therefore, we get a new representation of xy as follows:

$$[xy] = [c] + \varepsilon[b] + \delta[a] + \varepsilon\delta.$$

Using the tools from the previous sections we can now construct a protocol Π_{AMPC} that securely implements the MPC functionality $\mathcal{F}_{\text{AMPC}}$ in the UC security framework. $\mathcal{F}_{\text{AMPC}}$ and Π_{AMPC} are presented in Figure 6 and Figure 7 respectively. The proof of Theorem 1 can be found in the full paper [BDOZ10].

Functionality $\mathcal{F}_{\text{AMPC}}$

Initialize: On input $(init, p)$ from all parties, the functionality activates and stores the modulus p .

Rand: On input $(rand, P_i, varid)$ from all parties P_i , with $varid$ a fresh identifier, the functionality picks $r \leftarrow \mathbb{Z}_p$ and stores $(varid, r)$.

Input: On input $(input, P_i, varid, x)$ from P_i and $(input, P_i, varid, ?)$ from all other parties, with $varid$ a fresh identifier, the functionality stores $(varid, x)$.

Add: On command $(add, varid_1, varid_2, varid_3)$ from all parties (if $varid_1, varid_2$ are present in memory and $varid_3$ is not), the functionality retrieves $(varid_1, x)$, $(varid_2, y)$ and stores $(varid_3, x + y \bmod p)$.

Multiply: On input $(multiply, varid_1, varid_2, varid_3)$ from all parties (if $varid_1, varid_2$ are present in memory and $varid_3$ is not), the functionality retrieves $(varid_1, x)$, $(varid_2, y)$ and stores $(varid_3, x \cdot y \bmod p)$.

Output: On input $(output, P_i, varid)$ from all parties (if $varid$ is present in memory), the functionality retrieves $(varid, x)$ and outputs it to P_i .

Fig. 6. The ideal functionality for arithmetic MPC

Protocol Π_{AMPC}

Initialize: The parties first invoke $\mathcal{F}_{\text{TRIP}}(init, p)$. Then, they invoke $\mathcal{F}_{\text{TRIP}}(triples, u)$ and $\mathcal{F}_{\text{TRIP}}(singles, u)$ a sufficient number of times to create enough singles and triples.

Input: To share P_i 's input $[x_i]$ with identifier $varid$, P_i takes a single $[a]$ from the set of available ones. Then, the following is performed:

1. $[a]$ is opened to P_i .
2. P_i broadcasts $\delta = x_i - a$.
3. The parties compute $[x_i] = [a] + \delta$.

Rand: The parties take an available single $[a]$ and store with identifier $varid$.

Add: To add $[x], [y]$ with identifiers $varid_1, varid_2$ the parties compute $[z] = [x] + [y]$ and assign $[z]$ the identifier $varid_3$.

Multiply: To multiply $[x], [y]$ with identifiers $varid_1, varid_2$ the parties do the following:

1. They take a triple $([a], [b], [c])$ from the set of the available ones.
2. $[x] - [a] = \varepsilon$ and $[y] - [b] = \delta$ are opened.
3. They compute $[z] = [c] + \varepsilon[b] + \delta[a] + \varepsilon\delta$
4. They assign $[z]$ the identifier $varid_3$ and remove $([a], [b], [c])$ from the set of the available triples.

Output: To output $[x]$ with identifier $varid$ to P_i the parties do an opening of $[x]$ to P_i .

Fig. 7. The protocol for arithmetic MPC

Theorem 1. *In the $\mathcal{F}_{\text{TRIP}}$ -hybrid model, the protocol Π_{AMPC} implements $\mathcal{F}_{\text{AMPC}}$ with statistical security against any static⁴, active adversary corrupting up to $n - 1$ parties.*

4 The Offline Phase

In this section we describe the protocol Π_{TRIP} which securely implements the functionality $\mathcal{F}_{\text{TRIP}}$ described in Section 3 in the presence of two standard functionalities: a key registration functionality $\mathcal{F}_{\text{KEYREG}}$ and a functionality that generates random challenges $\mathcal{F}_{\text{RAND}}$ ⁵. Detailed specifications of these functionalities can be found in the full paper [BDOZ10].

4.1 $\langle \cdot \rangle$ -Representation

Throughout the description of the offline phase, E_i will denote E_{pk_i} where pk_i is the public key of party P_i , as established by $\mathcal{F}_{\text{KEYREG}}$. We assume the cryptosystem used is semi-homomorphic modulo p , as defined in Section 2. In the following, we will always set $\tau = p/2$ and $\rho = \sigma$. Thus, if P_i generates a ciphertext $C = E_i(x, \mathbf{r})$ where $x \in \mathbb{Z}_p$ and \mathbf{r} is generated by \mathcal{D}_σ^d , C will be a (τ, ρ) -ciphertext. We will use the zero-knowledge protocols from Section 2.2. They depend on an “information theoretic” security parameter u controlling, e.g., the soundness error. We will say that a semi-homomorphic cryptosystem is *admissible* if it allows correct decryption of ciphertext produced in those protocols, that is, if $M \geq 2^{5u+2 \log u \tau^2}$ and $R \geq 2^{4u+\log u \tau \rho}$.

In the following $\langle x_k \rangle$ will stand for the following representation of $x_k \in \mathbb{Z}_p$: each P_i has published $E_i(x_{k,i})$ and holds $x_{k,i}$ privately, such that $x_k = \sum_i x_{k,i} \bmod p$. For the protocol to be secure, it will be necessary to ensure that the parties encrypt small enough plaintexts. For this purpose we use the Π_{POPK} described in Section 2.2, and we get the protocol in Figure 8 to establish a set $\langle x_k \rangle, k = 1, \dots, u$ of such random representations.

4.2 $\langle \cdot \rangle$ -Multiplication

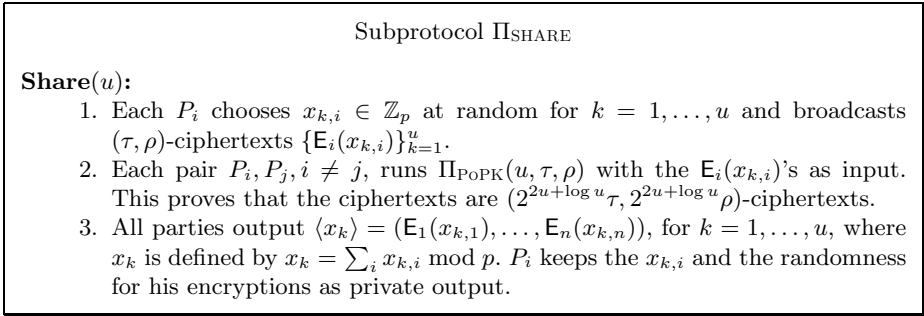
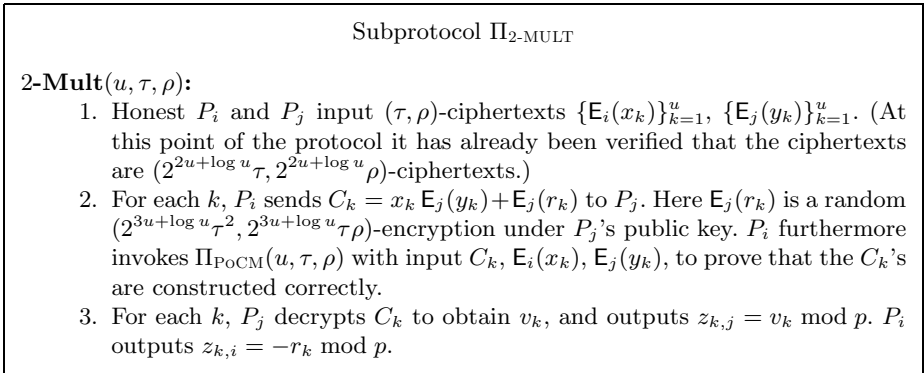
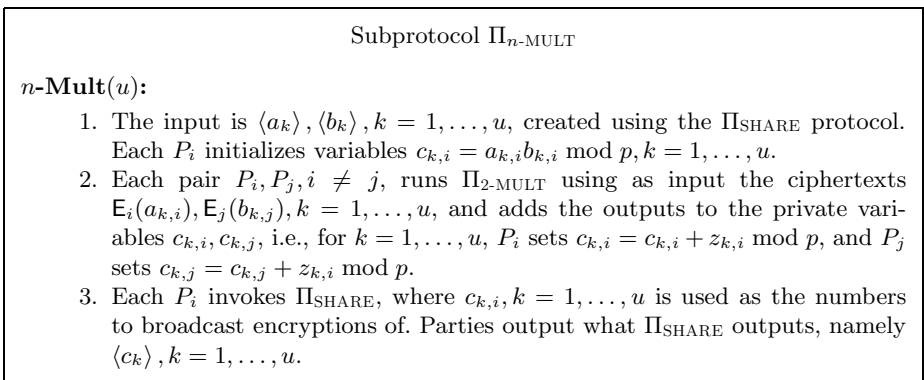
The final goal of the Π_{TRIP} protocol is to produce triples $[a_k], [b_k], [c_k]$ with $a_k b_k = c_k \bmod p$ in the $[\cdot]$ -representation, but for now we will disregard the MACs and construct a protocol $\Pi_{n\text{-MULT}}$ which produces triples $\langle a_k \rangle, \langle b_k \rangle, \langle c_k \rangle$ in the $\langle \cdot \rangle$ -representation⁶.

We will start by describing a two-party protocol. Assume P_i is holding a set of u (τ, ρ) -encryptions $E_i(x_k)$ under his public key and likewise P_j is holding u

⁴ Π_{AMPC} can actually be shown to adaptively secure, but our implementation of $\mathcal{F}_{\text{TRIP}}$ will only be statically secure.

⁵ $\mathcal{F}_{\text{RAND}}$ is only introduced for the sake of a cleaner presentation, and it could easily be implemented in the $\mathcal{F}_{\text{KEYREG}}$ model using semi-homomorphic encryption only.

⁶ In fact, due to the nature of the MACs, the same protocol that is used to compute two-party multiplications will be used later in order to construct the MACs as well.

**Fig. 8.** Subprotocol allowing parties to create random additively shared values**Fig. 9.** Subprotocol allowing two parties to obtain encrypted sharings of the product of their inputs**Fig. 10.** Protocol allowing the parties to construct $\langle c_k = a_k b_k \bmod p \rangle$ from $\langle a_k \rangle, \langle b_k \rangle$




Fig. 11. Subprotocol constructing $[a_k]$ from $\langle a_k \rangle$

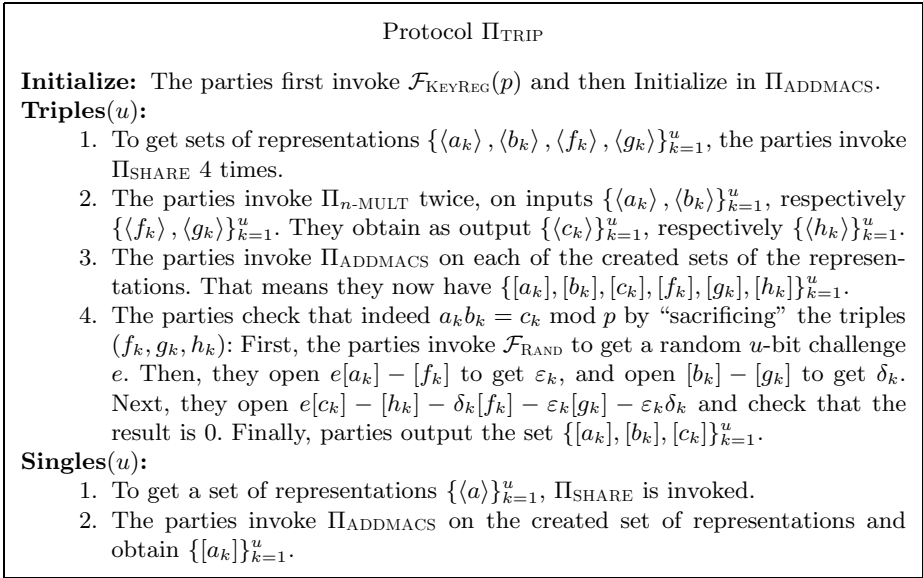


Fig. 12. The protocol for the offline phase

(τ, ρ) -encryptions $E_j(y_k)$ under his public key. For each k , we want the protocol to output $z_{k,i}, z_{k,j}$ to P_i, P_j , respectively, such that $x_k y_k = z_{k,i} + z_{k,j} \pmod p$. Such a protocol can be seen in Figure 9. This protocol does not commit parties to their output, so there is no guarantee that corrupt parties will later use their output correctly – however, the protocol ensures that malicious parties *know* which shares they ought to continue with. To build the protocol $\Pi_{n\text{-MULT}}$, the

first thing to notice is that given $\langle a_k \rangle$ and $\langle b_k \rangle$ we have that $c_k = a_k b_k = \sum_i \sum_j a_{k,i} b_{k,j}$. Constructing each of the terms in this sum in shared form is exactly what $\Pi_{2\text{-MULT}}$ allows us to do. The $\Pi_{n\text{-MULT}}$ protocol can now be seen in Figure 10. Note that it does not guarantee that the multiplicative relation in the triples holds, we will check for this later.

4.3 From $\langle \cdot \rangle$ -Triples to $[\cdot]$ -Triples

We first describe a protocol that allows us to add MACs to the $\langle \cdot \rangle$ -representation. This consists essentially of invoking the $\Pi_{2\text{-MULT}}$ a number of times. The protocol is shown in Figure 11. The full protocol Π_{TRIP} , which also includes the possibility of creating a set of single values, is now a straightforward application of the subprotocols we have defined now. This is shown in Figure 12. The proof of Theorem 2 can be found in the full paper [BDOZ10].

Theorem 2. *If the underlying cryptosystem is semi-homomorphic modulo p , admissible and IND-CPA secure, then Π_{TRIP} implements $\mathcal{F}_{\text{TRIP}}$ with computational security against any static, active adversary corrupting up to $n-1$ parties, in the $(\mathcal{F}_{\text{KEYREG}}, \mathcal{F}_{\text{RAND}})$ -hybrid model.*

References

- [BCNP04] Barak, B., Canetti, R., Nielsen, J.B., Pass, R.: Universally composable protocols with relaxed set-up assumptions. In: FOCS, pp. 186–195 (2004)
- [BD10] Bendlin, R., Damgård, I.: Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 201–218. Springer, Heidelberg (2010)
- [BDOZ10] Bendlin, R., Damgård, I., Orlandi, C., Zakarias, S.: Semi-homomorphic encryption and multiparty computation (full version). In: The Eprint Archive, report 2010/514 (2010)
- [Bea91] Beaver, D.: Efficient Multiparty Protocols Using Circuit Randomization. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 420–432. Springer, Heidelberg (1992)
- [BOGW88] Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: STOC, pp. 1–10 (1988)
- [Can01] Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: FOCS, pp. 136–145 (2001)
- [CCD88] Chaum, D., Crépeau, C., Damgård, I.: Multiparty unconditionally secure protocols (extended abstract). In: STOC, pp. 11–19 (1988)
- [CD09] Cramer, R., Damgård, I.: On the amortized complexity of zero-knowledge protocols. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 177–191. Springer, Heidelberg (2009)
- [CDN01] Cramer, R., Damgård, I., Nielsen, J.B.: Multiparty computation from threshold homomorphic encryption. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 280–299. Springer, Heidelberg (2001)

- [CLOS02] Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party and multi-party secure computation. In: STOC, pp. 494–503 (2002)
- [DGK09] Damgård, I., Geisler, M., Krøigaard, M.: A correction to efficient and secure comparison for on-line auctions. *IJACT* 1(4), 323–324 (2009)
- [DGHV10] van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 24–43. Springer, Heidelberg (2010)
- [DJ01] Damgård, I., Jurik, M.: A generalisation, a simplification and some applications of paillier’s probabilistic public-key system. In: Kim, K.-c. (ed.) PKC 2001. LNCS, vol. 1992, pp. 119–136. Springer, Heidelberg (2001)
- [DO10] Damgård, I., Orlandi, C.: Multiparty computation for dishonest majority: From passive to active security at low cost. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 558–576. Springer, Heidelberg (2010)
- [Gen09] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC, pp. 169–178 (2009)
- [GHV10] Gentry, C., Halevi, S., Vaikuntanathan, V.: A simple bgn-type cryptosystem from lwe. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 506–522. Springer, Heidelberg (2010)
- [HIK07] Harnik, D., Ishai, Y., Kushilevitz, E.: How Many Oblivious Transfers Are Needed for Secure Multiparty Computation? In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 284–302. Springer, Heidelberg (2007)
- [IPS09] Ishai, Y., Prabhakaran, M., Sahai, A.: Secure arithmetic computation with no honest majority. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 294–314. Springer, Heidelberg (2009)
- [LPS10] Lyubashevsky, V., Palacio, A., Segev, G.: Public-key cryptographic primitives provably as secure as subset sum. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 382–400. Springer, Heidelberg (2010)
- [OU98] Okamoto, T., Uchiyama, S.: A new public-key cryptosystem as secure as factoring. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 308–318. Springer, Heidelberg (1998)
- [Pai99] Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999)
- [PSSW09] Pinkas, B., Schneider, T., Smart, N.P., Williams, S.C.: Secure two-party computation is practical. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 250–267. Springer, Heidelberg (2009)
- [RAD78] Rivest, R., Adleman, L., Dertouzos, M.: On data banks and privacy homomorphisms. In: Foundations of Secure Computation, pp. 169–178 (1978)
- [Reg05] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC, pp. 84–93 (2005)