

Maliciously Secure Oblivious Linear Function Evaluation with Constant Overhead

Satrajit Ghosh^{*}, Jesper Buus Nielsen^{**}, and Tobias Nilges^{*}

Aarhus University, Denmark

Abstract. In this work we consider the problem of oblivious linear function evaluation (OLE). OLE is a special case of oblivious polynomial evaluation (OPE) and deals with the oblivious evaluation of a linear function $f(x) = ax + b$. This problem is non-trivial in the sense that the sender chooses a, b and the receiver x , but the receiver may only learn $f(x)$. We present a highly efficient and UC-secure construction of OLE in the OT-hybrid model that requires only $O(1)$ OTs per OLE. The construction is based on noisy encodings introduced by Naor and Pinkas (STOC'99) and used for passive secure OLEs by Ishai, Prabhakaran and Sahai (TCC'09). A result asymptotically similar to ours is known by applying the IPS compiler to the mentioned passive secure OLE protocol, but our protocol provides better constants and would be considerably simpler to implement. Concretely we use only 16 OTs to generate one active secure OLE, and our protocol achieves active security by adding fairly simple checks to the passive secure protocol. We therefore believe our protocol takes an important step towards basing practical active-secure arithmetic computations on OLEs. Our result requires novel techniques that might be of independent interest. As an application we present the currently most efficient OPE construction.

1 Introduction

The oblivious evaluation of functions is an essential building block in cryptographic protocols. The first and arguably most famous result in the area is oblivious transfer (OT), which was introduced in the seminal work of Rabin [31]. Here, a sender can specify two bits s_0, s_1 , and a receiver can learn one of the bits s_c depending on his choice bit c . It is guaranteed that the sender does not learn c , while the receiver learns nothing about s_{1-c} . Kilian [27] subsequently showed that OT allows the (oblivious) evaluation of *any* function.

While there has been tremendous progress in the area of generic MPC over the last three decades, there are certain classes of functions that can be evaluated more efficiently by direct constructions instead of taking the detour via MPC. In this context, Naor and Pinkas [29] introduced oblivious polynomial

^{*} Supported by the European Union's Horizon 2020 research and innovation programme under grant agreement #669255 (MPCPRO).

^{**} Supported by the European Union's Horizon 2020 research and innovation programme under grant agreement #731583 (SODA).

evaluation (OPE) as an useful primitive. OPE deals with the problem of evaluating a polynomial P on an input α obliviously, i.e., in such a way that the sender specifies the polynomial P but does not learn α , while the receiver learns $P(\alpha)$ but nothing else about P . OPE has many applications, ranging from secure set intersection [30,18] over RSA key generation [16] to oblivious keyword search [13]. Due to its versatility, OPE has received considerable attention in recent years [28,14,7,34,26,19,18,33].

A special case of OPE, called oblivious linear function evaluation (OLE, sometimes also referred to as OLFE, or OAFE) has more recently been introduced as an essential building block in (the preprocessing of) MPC protocols for arithmetic circuits [8] or garbled arithmetic circuits [1]. Instead of evaluating an arbitrary polynomial P , the receiver wants to evaluate a linear or affine function $f(x) = ax + b$. Ishai et al. [22] propose a passively secure protocol for oblivious multiplication which uses a similar approach as [29], and can be easily modified to give a passively secure OLE. Their approach to achieve active security is to apply a compiler like [21] to the passive protocol. Another approach was taken by [10], who built an unconditionally UC-secure protocol for OLE based on tamper-proof hardware [23] as a setup assumption.

Currently, all of the above mentioned actively secure realizations of OPE or OLE require rather expensive computations or strong setup assumptions. In contrast, the most efficient passively secure constructions built from noisy encodings and OT require only simple field operations. However, to date a direct construction of a maliciously secure protocol in this setting has been elusive. While passive-to-active compilers such as [21] yield actively secure protocols with a constant overhead, such a transformation typically incurs a prohibitively large constant, resulting in efficiency only in an asymptotic sense.¹ Thus, the most efficient realizations currently follow from applying the techniques used for the precomputation of multiplied values in arithmetic MPC protocols such as SPDZ [8] or more recently MASCOT [25].

1.1 Our Contribution

We present a UC-secure protocol for oblivious linear function evaluation in the OT-hybrid model based on noisy encodings. The protocol is based on the semi-honest secure implementation of OLE by Ishai et al. [22], which is the most efficient protocol for passively secure OLE that we are aware of. Our actively secure protocol only has a constant overhead of 2 compared to the passively secure original construction. In numbers, this means:

- We need 16 OTs per OLE, compared to 8 for the semi-honest protocol [22].
- We communicate $16 \cdot (2 + c_{\text{OT}}) \cdot n + 8$ field elements for n multiplications, compared to $8 \cdot (2 + c_{\text{OT}}) \cdot n$ in [22], where c_{OT} is the cost for one OT.
- The computational overhead is twice that of the semi-honest case.

¹ We will compare in more detail to protocols obtained via [21] after comparing to known direct constructions.

One nice property of [22] and the main reason for its efficiency is that it directly allows to multiply a batch of values. This property is preserved in our construction, i.e., we can simultaneously evaluate several linear functions.

In order to achieve our result we solve the long standing open problem of finding an actively secure OLE/OPE protocol which can directly be reduced to the security of noisy encodings (and OT). This problem was not solved in [29] and has only been touched upon in follow-up work [22,26]. The key technical contribution of the paper is a reduction which shows that noisy encodings are robust against leakage in a strong sense, which allows their application in a malicious setting. As a matter of fact, our robustness results are more general and extend to *all* noisy encodings.

An immediate application of our UC-secure batch-OLE construction is a UC-secure OPE construction. The construction is very simple and has basically no overhead over the OLE construction. We follow the approach taken in [30], i.e., we use the fact that a polynomial of degree d can be decomposed into d linear functions. Such a decomposed polynomial is evaluated with the batch-OLE and then reconstructed. UC-security against the sender directly follows from the UC-security of the batch-OLE. In order to make the protocol secure against a cheating receiver, we only have to add one additional check that ensures that the receiver chooses the same input for each linear function. Table 1 compares the efficiency of our result with existing solutions in the literature.

	Assumption	OTs	Expon.	Security
[7]	OT	$O(d\kappa)$	0	passive
[30]	OT & Noisy Encodings	$O(d\kappa \log \kappa)$	0	passive
[22]	OT & Noisy Encodings	$O(d)$	0	passive
[19]	CRS & DCRP	0	$O(ds)$	UC
[18]	DDH	0	$O(d)^*$	active
This work	OT & Noisy Encodings	$O(d)$	0	UC

Table 1. Overview of OPE realizations, where d is the degree of the polynomial, k is a computational security parameter and s a statistical security parameter ([19] propose $s \approx 160$). We compare the number of OTs and exponentiations in the respective protocols.

We point out that [18] only realizes OPE in the exponent, which is still sufficient for many applications, but requires additional work to yield a full fledged OPE. In particular, this might entail additional expensive operations. Another important factor regarding the efficiency of OT-based protocols is the cheap extendability of OT [20,24] which shows that the asymptotic price of OT is only small constant number of applications of a symmetric primitive like for instance AES. Therefore, the concrete cost of the OTs is much less than the price of exponentiations if d is sufficiently large, or if several OPEs have to be carried out. In such a scenario, we get significant improvements over previous (actively

secure) solutions, which always require expensive operations in the degree of the polynomial.

Comparing to IPS. In [22], which provides the passive secure OLE protocol from which we depart, the authors propose to get active security using the IPS compiler [21]. Our protocol is inspired by this approach but we have spared a lot of the generic mechanism of the IPS compiler and obtained a protocol which both in terms of number of OTs and implementation is considerable simpler.

In the IPS-compiler one needs to specify an outer protocol for n servers which should be active secure against a constant fraction of corruptions. Since our goal here is to generate active secure OLEs, we will need an outer protocol computing OLEs and therefore an arithmetic protocol is natural. One should also specify an inner protocol, which is a passive secure two-party computation protocol allowing to emulate the individual parties in the outer protocol. This protocol should be based on passive secure OLEs, as this is our underlying primitive.

For the outer protocol it seems as a good choice to use a variant of [2]. The protocol has perfect active security and therefore fits the [22] framework and avoids complications with coinflipping into the well. Since the IPS compiler does not need error recovery, but only error detection, the entire dispute-control mechanism of [2] can be skipped, yielding a very simple protocol. By setting the number n of servers very high, the number of corrupted servers that the outer protocol should tolerate can become an arbitrarily small constant ϵ . In that case one could probably get a packed version of [2], where one computes the same circuit on $(n - 3\epsilon)/2$ different values in parallel. This means that each secure multiplication done by the outer protocol will involve just more than 2 local multiplications of the servers. For simplicity we ignore all other costs than these two multiplications.

Using the natural inner protocol based on the passive secure OLE from [22], each of the two emulated multiplications will consume 2 passive secure OLEs. Each passive secure OLE consumes 8 OTs, for a total cost of 32 OTs per active secure OLE generated. Our protocol uses 16 OTs per OLE.

Putting the inner and outer protocols together in the IPS framework involves considerably additional complications like setting up watch-lists and encrypting and sending all internal messages of emulated servers on these watch-list channels. In comparison our protocol has almost no additional mechanisms extra to the underlying passive secure protocol. Our overhead on local computation of the parties compared to the passive secure OLE protocol is 2. It seems unlikely that the IPS compiler would have an overhead even close to just 2.

In summary, our protocol saves a factor of 2 in consumptions of OTs and is much simpler and therefore easier to implement and has considerable lower computational cost.

1.2 Technical Overview

At the heart of our constructions are noisy encodings. These were introduced by Naor and Pinkas [29] in their paper on OPE and provide a very efficient means to

obliviously compute multiplications. A noisy encoding is basically an encoding of a message via a linear code that is mixed with random values in such a way that the resulting vector hides which elements belong to the codeword and which elements are random, thereby hiding the initial message. In a little more detail, the input $\mathbf{x} \in \mathbb{F}^t$ is used as t sampling points on locations α_i of an otherwise random polynomial P of some degree $d > t$. Then the polynomial is evaluated at e.g., $4d$ positions β_i , and half of these positions are replaced by uniformly random values, resulting in the encoding \mathbf{v} . It is assumed that this encoding is indistinguishable from a uniformly random vector.²

Robustness of noisy encodings The main problem of using noisy encodings in maliciously secure protocols is that the encoding is typically used in a non-black-box way. On one hand this allows for very efficient protocols, but on the other hand a malicious party obtains knowledge that renders the assumption that is made on the indistinguishability of noisy encodings useless. In a little more detail, consider a situation where the adversary obtains the encoding and manipulates it in a way that is not specified by the protocol. The honest party only obtains part of the encoding (this is usually necessary even in the passively secure case). In order to achieve active security, a check is performed which is supposed to catch a deviating adversary. But since the check is dependent on which part of the encoding the honest party learned, this check actually leaks some non-trivial information to the adversary, typically noisy positions of the codeword.

We show that noisy encodings as defined by [30,22] are very robust with respect to leakage. In particular, we show the following theorem that is basically a stronger version of a result previously obtained by Kiayias and Yung [26].

Theorem. (informal) *For appropriate choices of parameters, noisy encodings are resilient against non-adaptive leakage of $O(\log \kappa)$ noisy positions.*

In a little more detail, we show that a noisy encoding generated as described above remains indistinguishable from a random vector of field elements, even for an adversary that is allowed to *fix* the position of f noisy positions. Fixing f positions is of course stronger than being able to leak f positions. The security loss incurred by the fixing of f positions is 3^f .

We then show that an adversary which is given a noisy encoding cannot identify a super-logarithmic sized set consisting of only noisy positions.

Theorem. (informal) *For appropriate choices of parameters, an adversary cannot identify more than $O(\log \kappa)$ noisy positions in a noisy encoding.*

These theorems together show that we can tolerate the leakage of any number of noisy positions that might be guessed. This is the basis for the security of our protocol. Note that tolerance to leakage of a set of noisy positions that might

² The problem is related to efficient polynomial reconstruction, i.e., decoding Reed-Solomon codes, and as such well researched. The parameters have to be chosen in such a way that all known decoding algorithms fail.

be guessed is not trivial, as we are working with an indistinguishability notion. Hence leakage of a single bit might *a priori* break the assumption.

We describe the main idea behind our reduction proving the first theorem. Assume that there are a total of ρ noisy positions. Consider an adversary that is allowed to submit a set F . Then we generate a noisy encoding as above, except that all positions $i \in F$ are fixed to be noisy. The remaining $\rho - |F|$ noisy positions are picked at random. Denote the distribution by $\mathbf{v}^{\rho, F}$. Let $\mathbf{v}^\rho = \mathbf{v}^{\rho, \emptyset}$. Let \mathbf{v}^\S denote a vector with all positions being uniformly random. We start from the assumption that $\mathbf{v}^\rho \approx \mathbf{v}^\S$. Let n be the total number of positions. Then clearly $\mathbf{v}^{n, F} = \mathbf{v}^\S$ for all F .

We want to prove that $\mathbf{v}^{\rho, F} \approx \mathbf{v}^\S$ for small F . Let F be a set of size f . Assume that we managed to prove that $\mathbf{v}^{\rho, F'} \approx \mathbf{v}^\S$ for all sets F' of size $f - 1$. Assume also that we have managed to prove that $\mathbf{v}^{\rho+1, F} \approx \mathbf{v}^\S$.

For a set F let i be the smallest index in F and let $F = F' \cup \{i\}$. Consider the reduction which is given \mathbf{v} from $\mathbf{v}^{\rho, F'}$ or \mathbf{v}^\S and which adds noise to position i in \mathbf{v} and outputs the result \mathbf{v}' . If $\mathbf{v} \sim \mathbf{v}^\S$, then $\mathbf{v}' \sim \mathbf{v}^\S$. If $\mathbf{v} \sim \mathbf{v}^{\rho, F'}$, then $\mathbf{v}' \sim \alpha \mathbf{v}^{\rho+1, F} + (1 - \alpha) \mathbf{v}^{\rho, F}$, where α is the probability that i is not already a noisy position. Putting these together we get that $\mathbf{v}^{\rho, F'} \approx \mathbf{v}^\S$ implies that $\alpha \mathbf{v}^{\rho+1, F} + (1 - \alpha) \mathbf{v}^{\rho, F} \approx \mathbf{v}^\S$. We then use that $\mathbf{v}^{\rho+1, F} \approx \mathbf{v}^\S$ to get that $\alpha \mathbf{v}^\S + (1 - \alpha) \mathbf{v}^{\rho, F} \approx \mathbf{v}^\S$, which implies that $\mathbf{v}^{\rho, F} \approx \mathbf{v}^\S$, when α is not too large.

We are then left with proving that $\mathbf{v}^{\rho, F'} \approx \mathbf{v}^\S$ and $\mathbf{v}^{\rho+1, F} \approx \mathbf{v}^\S$. These are proven by induction. The basis for $\mathbf{v}^{\rho, F'} \approx \mathbf{v}^\S$ is $\mathbf{v}^{\rho, \emptyset} \approx \mathbf{v}^\S$. The basis for $\mathbf{v}^{\rho+1, F} \approx \mathbf{v}^\S$ is $\mathbf{v}^{n, F} = \mathbf{v}^\S$. Controlling the security loss in these polynomially deep and nested inductions is tricky. We give the full details later.

We now give the intuition behind the proof of the second theorem. Assume that some adversary can guess a set S of s noisy positions with polynomial probability p_1 given an encoding $\mathbf{v}^\rho \approx \mathbf{v}^\S$ and assume that s is super-logarithmic and that ρ/n is a non-zero constant. We prove the theorem for noise level ρ but have to start with the assumption that $\mathbf{v}^{\rho - c\kappa} \approx \mathbf{v}^\S$ for an appropriate constant $c \in (0, 1)$ and where κ is the security parameter.

Consider the reduction which is given a sample \mathbf{v} from $\mathbf{v}^{\rho - c\kappa}$ or \mathbf{v}^\S . It starts by adding κ random positions R to \mathbf{v} to get \mathbf{v}' . Then it feeds \mathbf{v}' to \mathcal{A} to get a set S . Then it uses its knowledge of R to sample the size of the intersection between S and the noisy positions in \mathbf{v}' . If it is "large" we guess that $\mathbf{v} \sim \mathbf{v}^{\rho - c\kappa}$. Otherwise we guess that $\mathbf{v} \sim \mathbf{v}^\S$. We pick c such that the total number of random positions in \mathbf{v}' is ρ with polynomial probability when $\mathbf{v} \sim \mathbf{v}^{\rho - c\kappa}$, in which case S is a subset of the noisy positions with probability p_1 , which will give a large intersection. If $\mathbf{v} \sim \mathbf{v}^\S$, then R is uniformly hidden to the adversary, and the expected size of the intersection will be smaller by a constant factor depending on ρ and c . The calibration of c and "small" is done as to allow a formal proof using a Chernoff bound. The details are given in the following.

Efficient OLE from noisy encodings. We build a UC-secure OLE protocol inspired by the passively secure multiplication protocol of Ishai et al. [22]. Let us briefly recall their construction on an intuitive level. One party, let us call it the sender, has as input t values $a_1, \dots, a_t \in \mathbb{F}$, while the receiver has an input $b_1, \dots, b_t \in \mathbb{F}$. A set of distinct points $\alpha_1, \dots, \alpha_{n/4}$ is fixed. The high-level idea is as follows: both sender and receiver interpolate a degree $n/4 - 1$ polynomial through the points (α_i, a_i) and (α_i, b_i) (picking a_i, b_i for $i > t$ randomly), to obtain $A(x)$ and $B(x)$, respectively. They also agree on n points β_1, \dots, β_n . Now the receiver replaces half of the points $B(\beta_i)$ with uniformly random values (he creates a noisy encoding) and sends these n values $\bar{B}(\beta_i)$ to the sender. He keeps track of the noiseless positions using an index set L . The sender draws an additional random polynomial R of degree $2(n/4 - 1)$ to mask the output. He then computes $Y(\beta_i) = A(\beta_i) \cdot \bar{B}(\beta_i) + R(\beta_i)$ and uses these points as input into a $n/2 - 1$ -out-of- n OT, from which the receiver chooses the $n/2 - 1$ values in L . He can then interpolate the obtained points of $Y(\beta_i)$ to reconstruct Y and learn $a_i \cdot b_i + r_i$ in the positions α_i . This also directly yields an OLE: the polynomial R is simply used as another input of the sender, since it can be generated identically to A .

The passive security for the sender follows from the fact that the receiver obtains only $n/2 - 1$ values and thus R completely masks the inputs a_1, \dots, a_t . Passive security of the receiver follows from the noisy encoding, i.e., the sender cannot learn B from the noisy encoding.

In order to achieve actively secure OLE from the above protocol, we have to ensure several things: first of all, we need to use an actively secure k -out-of- n OT. But instead of using a black-box realization, which incurs an overhead of $n \log n$ on the number of OTs, we use n 1-out-of-2 OTs and ensure that the right number of messages was picked via a secret sharing, which the receiver has to reconstruct. This protocol first appeared in [32]. It does not have active security against the sender, who can guess some choice bits. A less efficient but active secure version was later given in [9], using verifiable secret sharing. We can, however, use the more efficient but less secure original variant as we can tolerate leakage of a few choice bits in the overall protocol.

Secondly, we also need to make sure that the parties used the right inputs in the computation, i.e., valid polynomials A, B and R . In order to catch deviations, we add two checks—one in each direction. The check is fairly simple: one party selects a random point z and the other party sends a pair $A(z), R(z)$, or $B(z), Y(z)$ respectively. Each party can now locally verify that the values satisfy the equation $A(z) \cdot B(z) + R(z) = Y(z)$.

As it turns out, both of these additions to the protocol, while ensuring protocol compliance w.r.t. the inputs, are dependent on the encoding. But this also means that a malicious sender can do selective failure attacks, e.g., it inputs incorrect shares for the secret sharing, and gets some leakage on the “secret key” of the encoding. This problem does not occur when considering semi-honest security.

2 Preliminaries

We use the standard notions of probabilistic polynomial time (PPT) algorithms, negligible and overwhelming functions. Further, we denote by $\mathbf{x} \in \mathbb{F}^n$ a vector of length n , x_i as the i th element of \mathbf{x} and $\mathbf{x}_{|I}$ all x_i for $i \in I$. Unless noted otherwise, $P(x)$ denotes a polynomial in $\mathbb{F}[X]$, and X denotes a distribution.

We will typically denote a value \hat{x} chosen or extracted by the simulator, while x^* is chosen by the adversary \mathcal{A} .

2.1 Universal Composability Framework

We state and prove our results in the Universal Composability (UC) framework of Canetti [5]. Security is defined via the comparison of an *ideal model* and a *real model*. In the real model, a protocol Π between the protocol participants is executed, while in the ideal model the parties only communicate with an ideal functionality \mathcal{F} that is supposed to model the ideal security guarantees of the protocol. For an adversary \mathcal{A} in the real protocol who coordinates the behavior of all malicious parties, there has to exist a simulator \mathcal{S} for \mathcal{A} in the ideal protocol. An environment \mathcal{Z} , which is plugged to both the real and the ideal protocol, provides the inputs to the parties and can read the outputs. The simulator \mathcal{S} has to ensure that \mathcal{Z} is not able to distinguish these models. Thus, even with concurrently executed protocols (running in the environment) the security holds. Usually, we assume that \mathcal{A} is a dummy adversary controlled by \mathcal{Z} , which means that \mathcal{Z} can adaptively choose its inputs depending on protocol messages it received and send messages on behalf of a (corrupted) protocol party.

More formally, let $\text{Real}_{\Pi}^{\mathcal{A}}(\mathcal{Z})$ denote the random variable describing the output of \mathcal{Z} when interacting with the real model, and let $\text{Ideal}_{\mathcal{F}}^{\mathcal{S}}(\mathcal{Z})$ denote the random variable describing the output of \mathcal{Z} when interacting with the ideal model. A protocol Π is said to *UC-realize* a functionality \mathcal{F} if for any (PPT) adversary \mathcal{A} , there exists a PPT simulator \mathcal{S} such that for any (PPT) environment \mathcal{Z} , $\text{Real}_{\Pi}^{\mathcal{A}}(\mathcal{Z}) \approx \text{Ideal}_{\mathcal{F}}^{\mathcal{S}}(\mathcal{Z})$.

For our constructions we assume active adversaries and static corruption. We prove security in the hybrid model access to oblivious transfer (OT).

2.2 Commitment Scheme

A commitment scheme COM consists of two algorithms (COM.Commit, COM.Open). It is a two party protocol between a sender and a receiver. In the commit phase of the protocol, when the sender wants to commit to some secret value m , it runs COM.Commit(m) and gets back two values (com, unv). It sends com to the receiver. Later on in the unveil phase, the sender sends the unveil information unv to the receiver, who can use COM.Open to verify that the commitment com contains the actual secret m .

A commitment scheme must satisfy two security properties; 1) *Hiding*: The receiver cannot learn any information about the committed secret before the unveil phase, and 2) *Binding*: The sender must not be able to change the committed

secret after the commit phase. For our purpose we need efficient UC-secure commitment schemes that can be realized in \mathcal{F}_{OT} -hybrid model and in \mathcal{F}_{OLE} -hybrid model.

In [6] the authors proposed an UC-secure commitment scheme in the \mathcal{F}_{OT} -hybrid model. Their protocol gives the first UC commitment scheme with "optimal properties" of rate approaching 1 and linear time complexity, in a "amortized sense". In our UC-secure OLE protocol also we need to commit to many values at a time, so we can use their UC commitment scheme in our protocol.

3 Noisy Encodings

The security of our protocols is based on a noisy encoding assumption. Very briefly, a noisy encoding is an encoding of a message, e.g., via a linear code, that is mixed with random field elements. It is assumed that such a codeword, and in particular the encoded message, cannot be recovered. This assumption seems reasonable due to its close relationship to decoding random linear codes or the efficient decoding of Reed-Solomon codes with a large fraction of random noise.

Noisy encodings were first introduced by Naor and Pinkas [29], specifically for the purpose of realizing OPE. Their encoding algorithm basically generates a random polynomial P of degree $k-1$ with $P(0) = x$. The polynomial is evaluated at $n > k$ locations, and then $n - k$ positions are randomized. Generalizing the approach of [30], Ishai et al. [22] proposed a more efficient encoding procedure that allows to encode several field elements at once instead of a single element, using techniques of [12]. Basically, they use Reed-Solomon codes and then artificially blind the codeword with random errors in order to mask the location of the codeword elements in the resulting string.

The encoding procedure depicted in Figure 1 is nearly identical to the procedure given in [22], apart from the fact that we do not fix the signal-to-noise ratio (because this will be dependent on the protocol). We also allow to pass a set of points \mathcal{P} as an argument to `Encode` to simplify the description of our protocol later on. This change has no impact on the assumption, since these points are made public anyway via G .

[29] propose two different encodings and related assumptions, tailored to their protocols. One of these assumptions was later broken by [3] and [4], and a fixed version was presented in [30]. We are only interested in the unbroken assumption. The same assumption was used by [22] and we will adopt their notation in the following.

Assumption 1 *Let κ be a security parameter and $n, \rho \in \text{poly}(\kappa)$. Further let $x, y \in \mathbb{F}^{t(\kappa)}$. Then the ensembles $\{\text{Encode}_{n,\rho}(x)\}_{\kappa}$ and $\{\text{Encode}_{n,\rho}(y)\}_{\kappa}$ are computationally indistinguishable, for $t \leq \frac{\ell}{4}$.*

As it is, our security reductions do not hold with respect to Assumption 1, but rather a variant of Assumption 1 which was already discussed in [29]. Instead of requiring indistinguishability of two encodings, we require that an encoding is

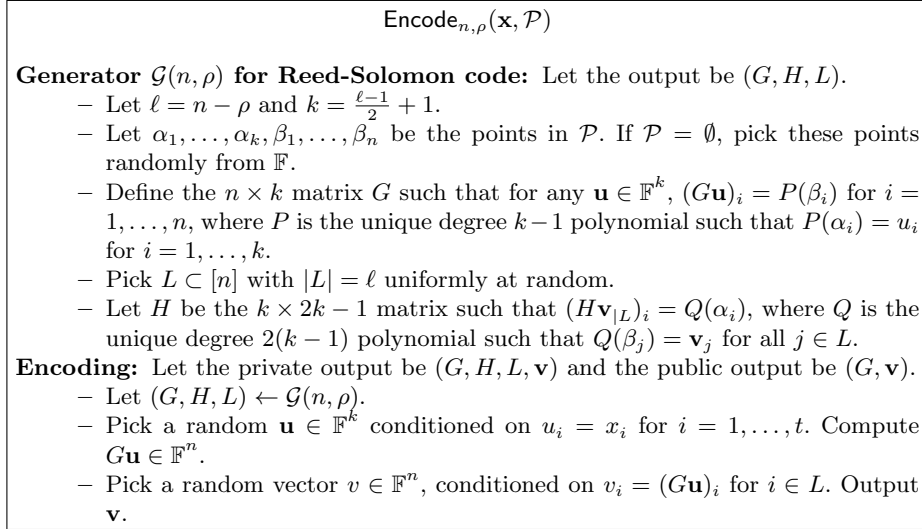


Fig. 1. Encoding procedure for noisy encodings.

pseudorandom. In order for these assumption to hold, [22] propose $n = 4\kappa, \rho = 2\kappa + 1$ as parameters, or $n = 8\kappa, \rho = 6\kappa + 1$ on the conservative side.

Assumption 2 *Let κ be a security parameter and $n, \rho \in \text{poly}(\kappa)$. Further let $x \in \mathbb{F}^{t(\kappa)}$. Then the ensembles $\{\text{Encode}_{n,\rho}(x)\}_\kappa$ and $\{G \leftarrow \mathcal{G}(n, \rho), v \leftarrow \mathbb{F}^n\}_\kappa$ are computationally indistinguishable, for $t \leq \frac{\ell}{4}$.*

Clearly, Assumption 2 implies Assumption 1, while the other direction is unclear. Apart from being a very natural assumption, Kiayias and Yung [26] provide additional evidence that Assumption 2 is valid. They show that if an adversary cannot decide for a random position i of the encoding whether it is part of the codeword or not, then the noisy codeword is indeed pseudorandom.

4 Noisy Encodings are Robust Against Leakage

In this section we show that a large class of computational assumptions can be extended to allow some leakage without loss of asymptotic security. This is one of the main technical contributions of the paper and we deem the reductions to be of independent interest. We first define the class of assumptions we consider.

Definition 1. *Let a finite field \mathbb{F} be given. For a positive integer n we use \mathbf{U}^n to denote the uniform distribution on \mathbb{F}^n . Let n be the length of a codeword, ρ the number of randomised positions, G a generator for the encoding and F some fixed random positions with $|F| \leq \rho$. The distribution $\mathbf{Y}^{n,\rho,G,F}$ is sampled as follows.*

1. Sample $(x_1, \dots, x_n) \leftarrow G(1^\kappa)$, where κ is the security parameter.

2. Sample uniform $R \subseteq [n]$ under the restriction $|R| = \rho$ and $F \subseteq R$.
3. Sample uniform $(e_1, \dots, e_n) \leftarrow \mathbf{U}^n$ under the restriction $e_i = 0$ for $i \notin R$.
4. Output $\mathbf{y} = \mathbf{x} + \mathbf{e}$.

Clearly it holds that the above defined `Encode` generalises Definition 1, i.e., all of the following results hold for noisy encodings as well.

We will mostly consider the case that $n = \Theta(\kappa)$ and $\rho = \Theta(\kappa)$. Typically n and G are fixed, in which case we denote the distribution of \mathbf{y} by $\mathbf{Y}^{\rho, F}$. Note that $\mathbf{Y}^{n, F} = \mathbf{U}^n$.

We are going to assume that for sufficiently large ρ it holds that $\mathbf{Y}^\rho \approx \mathbf{Y}^n$, where \approx means the distributions are computationally indistinguishable. For example, this is given by Assumption 2 for noisy encodings with appropriate parameters. We will use this to prove that the same result holds given limited leakage on R and that it is hard to compute a lot of elements of R given only \mathbf{y} .

When we prove the first result, we are going to do an argument with two nested recursive reductions. To make it easier to track the security loss, we are going to shift to a concrete security notation for a while and then switch back to asymptotic security when we have control over the security loss.

Given two distributions A_0 and A_1 and a distinguisher D let $\text{Adv}_D(A_0, A_1) = \Pr[A_1 = 1] - \Pr[A_0 = 1]$. We use $A_0 \approx_\epsilon^t A_1$ to denote that it holds for all distinguishers computable in time t that $\text{Adv}_D(A_0, A_1) \leq \epsilon$.

Given two distributions A_0 and A_1 and $0 \leq \alpha_0 \leq 1$ and $\alpha_1 = 1 - \alpha_0$ we use $B = \alpha_0 A_0 + \alpha_1 A_1$ to denote the distribution generated by the following procedure. Sample $c \in \{0, 1\}$ with $\Pr[c = i] = \alpha_i$. Then sample $b \leftarrow A_c$ and output b .

We will use the following simple facts in our proofs.

Lemma 1. *Let A_0, A_1 and Z be distributions.*

$$\begin{aligned} A_0 \approx_\epsilon^t A_1 &\Rightarrow \alpha_0 A_0 + \alpha_1 Z \approx_{\alpha_0 \epsilon}^t \alpha_0 A_1 + \alpha_1 Z . \\ \alpha_0 A_0 + \alpha_1 Z \approx_\epsilon^t \alpha_0 A_1 + \alpha_1 Z &\Rightarrow A_0 \approx_{\alpha_0^{-1} \epsilon}^t A_1 . \end{aligned}$$

Proof. Let $B_i = \alpha_0 A_0 + \alpha_1 Z$. We first prove the first implication. Consider a distinguisher D for B_0 and B_1 . Then

$$\begin{aligned} \text{Adv}_D(B_0, B_1) &= \Pr[D(B_1) = 1] - \Pr[D(B_0) = 1] \\ &= \sum_i \alpha_i \Pr[D(B_1) = 1 | c = i] - \sum_i \alpha_i \Pr[D(B_0) = 1 | c = i] \\ &= \alpha_0 \Pr[D(B_1) = 1 | c = 0] - \alpha_0 \Pr[D(B_0) = 1 | c = 0] \\ &= \alpha_0 \Pr[D(A_1) = 1] - \alpha_0 \Pr[D(A_0) = 1] \\ &= \alpha_0 \Pr[D(A_1) = 1] - \alpha_0 \Pr[D(A_0) = 1] , \end{aligned}$$

from which it follows that

$$\text{Adv}_D(B_0, B_1) = \alpha_0 \text{Adv}_D(A_0, A_1) . \quad (1)$$

From (1) it follows that $\text{Adv}_D(B_0, B_1) \leq \alpha_0 \epsilon$ for all D , which proves the claim in the lemma. Consider a distinguisher D for A_0 and A_1 . It can also act as distinguisher for B_0 and B_1 , so from (1) we have that

$$\text{Adv}_D(A_0, A_1) = \alpha_0^{-1} \text{Adv}_D(B_0, B_1) .$$

From this the second claim follows. \square

In the following we will show that if $Y^\rho \approx_\epsilon^{t+t'} Y^n$ and F is not too large, then $Y^{\rho, F} \approx_{\epsilon'}^t Y^{n, F}$ for ϵ' polynomially related to ϵ and t' a small fixed polynomial. Since the reduction will be recursive and will modify ϵ multiplicatively, we will keep explicit track of ϵ to ensure the security loss is not too large. As for t , each reduction will only *add* a small t' to t , namely the time to sample a distribution. The time will therefore clearly grow by at most a polynomial term. We therefore do not keep track of t , for notational convenience.

Lemma 2. *If $\rho - |F| \geq 2n/3$ and $Y^j \approx_\epsilon Y^n$ for all $j \geq \rho$, then $Y^{\rho, F} \approx_{\sigma_\rho} Y^{n, F}$ for $\sigma_\rho = 3^{|F|} \epsilon$.*

Proof. We prove the claim by induction in the size of F . If $|F| = 0$ it follows by assumption. Consider then the following randomised function f with inputs (y_1, \dots, y_n) and F . Let i be the largest element in F and let $F' = F \setminus \{i\}$. Sample uniformly random $y'_i \in \mathbb{F}$. For $j \neq i$, let $y'_j = y_j$. Output \mathbf{y}' . Consider the distribution $Y^{\rho, F'}$. Let R denote the randomised positions. If $i \in R$, then $f(Y^{\rho, F'}) = Y^{\rho, F}$. If $i \notin R$, then $f(Y^{\rho, F'}) = Y^{\rho+1, F}$, as we added one more noisy point. The point i is a fixed point not in F' . There are $n - |F| + 1$ points not in F' . There are ρ randomised points, i.e., $|R| = \rho$. Exactly $|F| - 1$ of these points are the points of F' . The other points are uniform outside F' . So there are $\rho - |F| + 1$ such points. Therefore the probability that $i \in R$ is $p = \frac{\rho - |F| + 1}{n - |F| + 1}$. It follows that

$$f(Y^{\rho, F'}) = Y^{n, F} , \quad f(Y^{\rho, F'}) = pY^{\rho, F} + (1-p)Y^{\rho+1, F} .$$

It then follows from $Y^{n, F'} \approx_{\epsilon'} Y^{\rho, F'}$ (where $\epsilon' = 3^{|F'|} \epsilon$) that

$$Y^{n, F} \approx_{\epsilon'} pY^{\rho, F} + (1-p)Y^{\rho+1, F} .$$

We now claim that $Y^{\rho, F} \approx_{3\epsilon'} Y^{n, F}$. The claim is trivially true for $\rho = n$, so we can assume that $\rho < n$ and assume the claim is true for all $\rho' > \rho$. Using Lemma 1 and the induction hypothesis we get that

$$pY^{\rho, F} + (1-p)Y^{\rho+1, F} \approx_{(1-p)3\epsilon'} pY^{\rho, F} + (1-p)Y^{n, F} .$$

Clearly

$$Y^{n, F} \approx_0 pY^{n, F} + (1-p)Y^{n, F} .$$

Putting these together we get that

$$pY^{n, F} + (1-p)Y^{n, F} \approx_{\epsilon' + (1-p)3\epsilon'} pY^{\rho, F} + (1-p)Y^{n, F} .$$

Using Lemma 1 we get that $Y^{n,F} \approx_{p^{-1}(\epsilon' + (1-p)3\epsilon')} Y^{\rho,F}$. We have that $p \geq 2/3$ so

$$p^{-1}(\epsilon' + (1-p)3\epsilon') \leq \frac{3}{2}(\epsilon' + \frac{1}{3}3\epsilon') = 3\epsilon' = 3^{|F|}\epsilon.$$

□

In the rest of the section, assume that n and ρ are functions of a security parameter κ and that $n, \rho = \Theta(\kappa)$. Also assume that $\rho \geq 2n/3$ and that $n - \rho = \Theta(\kappa)$. We say that $Y^\rho \approx Y^n$ if $Y^{\rho'} \approx_{1/p(\kappa)}^{q(\kappa)} Y^{n(\kappa)}$ for all polynomials p and q and all sufficiently large κ and all $\rho' \geq \rho$.

From $3^{O(\log \kappa)}$ being a polynomial in κ we get that

Corollary 1 *If $Y^\rho \approx Y^n$ and $F \subseteq [n]$ has size $O(\log \kappa)$, then $Y^{\rho,F} \approx Y^n$.*

Now assume that $Y^\rho \approx Y^n$. Let $Y^{\rho,F,\neg}$ be defined as $Y^{\rho,F}$ except that R is sampled according to the restriction that $F \not\subseteq R$, i.e., F has at least one element outside R . Let $p(F)$ be the probability that $F \subseteq R$ for a uniform R . Then by definition and the law of total probability $Y^\rho = pY^{\rho,F} + (1-p)Y^{\rho,F,\neg}$. We have that $Y^\rho \approx Y^n$ and that $Y^{\rho,F} \approx_{3^{|F|}} Y^{n,F} = Y^n$. Putting these together we have that $Y^n \approx_{p3^{|F|}} pY^n + (1-p)Y^{\rho,F,\neg}$. We then get that $Y^n \approx_{3^{|F|}} Y^{\rho,F,\neg}$.

Corollary 2 *If $Y^\rho \approx Y^n$ and $F \subseteq [n]$ has size $O(\log \kappa)$, then $Y^{\rho,F,\neg} \approx Y^n$.*

We now prove that given one small query on R does not break the security.

Definition 2. *Let \mathcal{A} be a PPT algorithm and Y as defined in Definition 1. The game $\mathcal{G}_{\text{leak}}$ is defined as follows.*

1. Run \mathcal{A} to get a subset $Q \subseteq [n]$.
2. Sample a uniformly random bit c .
 - If $c = 0$, then sample $\mathbf{y} \leftarrow Y^\rho$ and let R be the subset used in the sampling
 - If $c = 1$, then sample $\mathbf{y} \leftarrow Y^n$ and let $R \subseteq [n]$ be a uniformly random subset of size ρ .
3. Let $r \in \{0, 1\}$ be 1 iff $Q \subseteq R$ and input (r, \mathbf{y}) to \mathcal{A} .
4. Run \mathcal{A} to get a guess $g \in \{0, 1\}$.

The advantage of \mathcal{A} is $\text{Adv}_{\mathcal{A}} = \Pr[g = 1 | c = 1] - \Pr[g = 1 | c = 0]$.

Theorem 1. *Assume that $n, \rho = \Theta(\kappa)$ and $\rho \geq \frac{2}{3}n$. If $Y^\rho \approx Y^n$, then $\text{Adv}_{\mathcal{A}} \approx 0$.*

Proof. Let p be the probability that $Q \subseteq R$. If p is negligible, then in item 3 of the game we could send the constant $r = 0$ to \mathcal{A} and it would only change the advantage by a negligible amount. But in the thus modified game $\text{Adv}_{\mathcal{A}} \approx 0$ because $Y^n \approx Y^\rho$. So assume that p is a polynomial.³ Let Y_0 be the distribution

³ Formally we should consider the case where it is a polynomial for infinitely many κ , but the following argument generalises easily to this case.

of (r, y) when $c = 0$. Let Y_1 be the distribution of (r, \mathbf{y}) when $c = 1$. If $c = 0$, then (r, \mathbf{y}) is distributed as follows

$$Y_0 = p \cdot (1, Y^{\rho, Q}) + (1 - p) \cdot (0, Y^{\rho, Q, \neg}) .$$

When p is polynomial, then $|F| = O(\log \kappa)$ as $n - \rho = \Theta(\kappa)$. From this we get

$$Y_0 \approx p \cdot (1, Y^n) + (1 - p) \cdot (0, Y^n) = Y_1 ,$$

using the above asymptotic corollaries. \square

We will then prove that it is hard to compute a lot of elements of R .

Definition 3. Let \mathcal{A} be a PPT algorithm and Y as defined in Definition 1. The game $\mathcal{G}_{\text{ident}}$ is defined as follows.

1. Sample $\mathbf{y} \leftarrow Y^\rho$ and let R denote the randomized positions.
2. Input \mathbf{y} to \mathcal{A} .
3. Run \mathcal{A} and denote the output by $Q \subseteq [n]$. We require that $|Q| = s$.
4. Let $r \in \{0, 1\}$ be 1 iff $Q \subseteq R$.
5. Output r .

The advantage of \mathcal{A} is $\text{Adv}_{\mathcal{A}}^{\rho, s} = \Pr[r = 1]$.

Theorem 2. Let $n = \Theta(\kappa)$ and assume that $Y^\sigma \approx Y^n$. Then $\text{Adv}_{\mathcal{A}}^{\rho, s} \approx 0$ is true in both of these cases:

1. Let $\sigma = n \frac{\rho - \kappa}{n - \kappa}$ and $s = \kappa$.
2. Let $\sigma = \frac{n\kappa}{n - \rho - \kappa}$ and $s \in \omega(\log \kappa)$.

Proof. Let \mathcal{A} be an adversary such that when $Q \leftarrow \mathcal{A}(Y^\rho)$, then $Q \subseteq R$ with non-negligible probability p . The argumentation is similar for both cases. For the first part of the theorem, consider the following adversary $\mathcal{B}(\mathbf{y})$ receiving $\mathbf{y} \in \mathbb{F}^n$. It samples a uniform $X \subset [n]$ of size κ . For $i \in X$ let y'_i be uniformly random. For $i \notin X$ let $y'_i = y_i$. Compute $Q \leftarrow \mathcal{A}(\mathbf{y})$. If $|Q \cap X| \geq \frac{\kappa^2}{\rho}$, then output 1. Otherwise output 0.

We now prove that $\Pr[\mathcal{B}(Y^n) = 1] \approx 0$ and that $\Pr[\mathcal{B}(Y^\sigma) = 1]$ is non-negligible, which proves the first statement of the theorem.

Let R be the positions that were randomised in \mathbf{y} . Let $R' = R \cup X$. Note that if $\mathbf{y} \leftarrow Y^\sigma$, then

$$\mathbb{E}[|R'|] = \kappa + \sigma - \mathbb{E}[|X \cap R|] = \kappa + \sigma - \kappa \frac{\sigma}{n} = \rho .$$

It is straight forward to verify that $\Pr[|R'| = \rho] = 1/O(\kappa)$, which implies that $\Pr[Q \subseteq R] = p/O(\kappa)$, which is non-negligible when p is non-negligible. Let E denote the event that $Q \subseteq R$. By a simple application of linearity of expectation we have that

$$\mathbb{E}[|Q \cap X| | E] = \frac{\kappa^2}{\rho} ,$$

as X is a uniformly random subset $X \subseteq R$ given the view of \mathcal{A} . From this it follows that $\Pr[\mathcal{B}(Y^\sigma) = 1]$ is non-negligible.

Then consider $\mathcal{B}(Y^n)$. Note that now $R = [n]$ and again X is a uniformly random subset of R independent of the view of \mathcal{A} . Therefore

$$\mathbb{E}[|Q \cap X|] = \frac{\kappa^2}{n} =: \mu.$$

Then

$$\Pr[|Q \cap X| \geq \frac{s(\rho - \kappa)}{\rho}] = \Pr[|Q \cap X| \geq \frac{n}{\rho} \mu] = \Pr[|Q \cap X| \geq (1 + \delta)\mu]$$

for $\delta \in (0, 1)$. It follows that

$$\Pr[|Q \cap X| \geq \frac{\kappa^2}{\rho}] \leq e^{-\frac{\mu\delta^2}{3}} = e^{-\Theta(\mu)} = e^{-\Theta(\kappa)} = \text{negl}(\kappa).$$

To see this let $X = \{x_1, \dots, x_\kappa\}$ and let X_i be the indicator variable for the event that the i 'th element of X ends up in Q . Then $\Pr[X_i = 1] = \frac{\kappa}{n}$ and $|X \cap Q| = \sum_i X_i$. Consider then the modified experiment called *Independent Sampling*, where we sample the κ elements for X uniformly at random from $[n]$ and independently, i.e., it may happen that two of them are identical. In that case the inequality is a simple Chernoff bound. It is easy to see that when we go back to *Dependent Sampling*, where we sample x_i uniformly at random except that they must be different from x_1, \dots, x_{i-1} , then we only lower the variance of the sum $\sum_i X_i$ compared to Independent Sampling, so $\Pr[|Q \cap X| \geq (1 + \delta)\mu]$ will drop. To see this, consider the sequence $x, x_1 + x_2, \dots, \sum_i x_i$ as a random walk. In the Dependent Sampling case, when $\sum_i x_i$ is larger than the expectation, then x_{i+1} is less likely to be in Q compared to the Independent Sampling case, as an above expectation number of slots in Q is already taken. Similarly, when $\sum_i x_i$ is smaller than the expectation, then x_{i+1} is more likely to be in Q compared to the Independent Sampling case, as a below expectation number of slots in Q is already taken. Therefore the random walk in the Dependent Sampling case will always tend closer to average compared to the Independent Sampling random walk.

The second statement of Theorem 2 follows by setting X to be a uniform subset of size $\rho - \kappa$. As above, if \mathcal{A} outputs Q such that $|Q \cap X| \geq \frac{s(\rho - \kappa)}{\rho}$, then $\mathcal{B}(y)$ outputs 1. Otherwise it outputs 0. Let again R be the positions that were randomised in y . Let $R' = R \cup X$. If $y \leftarrow Y^\sigma$, then

$$\mathbb{E}[|R'|] = \rho - \kappa + \sigma - \mathbb{E}[|X \cap R|] = \rho - \kappa + \sigma - (\rho - \kappa) \frac{\sigma}{n} = \rho.$$

Let E denote the event that $Q \subseteq R$. Following the above argumentation,

$$\mathbb{E}[|Q \cap X| | E] = \frac{s(\rho - \kappa)}{\rho}.$$

From this it follows that $\Pr[\mathcal{B}(Y^\sigma) = 1]$ is non-negligible. Then consider $\mathcal{B}(Y^n)$. Note that now $R = [n]$ and again X is a uniformly random subset of R independent of the view of \mathcal{A} . Therefore

$$\mathbb{E}[|Q \cap X|] = \frac{s(\rho - \kappa)}{n} =: \mu.$$

It follows that

$$\Pr[|Q \cap X| \geq \frac{s(\rho - \kappa)}{\rho}] \leq e^{-\frac{\mu^2}{3}} = e^{-\Theta(\mu)} = e^{-\omega(\log \kappa)} = \text{negl}(\kappa).$$

□

5 Constant Overhead Oblivious Linear Function Evaluation

Oblivious linear function evaluation (OLE) is the task of computing a linear function $f(x) = ax + b$ in the following setting. One party, let's call it the sender S , provides the function, namely the values a and b . The other party, the receiver R , wants to evaluate this function on his input x . This task becomes non-trivial if the parties want to evaluate the function in such a way that the sender learns nothing about x , while the receiver learns only $f(x)$, but not a and b . OLE can be seen as a special case of oblivious polynomial evaluation (OPE) as proposed by Naor and Pinkas [29], where instead of a linear function f , the sender provides a polynomial p .

5.1 Ideal Functionality

The efficiency of our protocol follows in part from the fact that we can directly perform a batch of multiplications. This is reflected in the ideal UC-functionality for $\mathcal{F}_{\text{OLE}}^t$ (cf. Figure 2), which allows both sender and receiver to input vectors of size t .

Functionality $\mathcal{F}_{\text{OLE}}^t$
<ol style="list-style-type: none"> 1. Upon receiving a message (inputS, \mathbf{a}, \mathbf{b}) from S with $\mathbf{a}, \mathbf{b} \in \mathbb{F}^t$, verify that there is no stored tuple, else ignore that message. Store \mathbf{a} and \mathbf{b} and send a message (input) to \mathcal{A}. 2. Upon receiving a message (inputR, \mathbf{x}) from R with $\mathbf{x} \in \mathbb{F}^t$, verify that there is no stored tuple, else ignore that message. Store \mathbf{x} and send a message (input) to \mathcal{A}. 3. Upon receiving a message (deliver, S) from \mathcal{A}, check if both \mathbf{a}, \mathbf{b} and \mathbf{x} are stored, else ignore that message. Send (delivered) to S. 4. Upon receiving a message (deliver, R) from \mathcal{A}, check if both \mathbf{a}, \mathbf{b} and \mathbf{x} are stored, else ignore that message. Set $y_i = a_i \cdot x_i + b_i$ for $i \in [t]$ and send (output, \mathbf{y}) to R.

Fig. 2. Ideal functionality for an oblivious linear function evaluation.

5.2 Our Protocol

Our starting point is the protocol of Ishai et al. [22] for *passively* secure batch multiplication. Their protocol is based on noisy encodings, similar to our construction. We will now briefly sketch their construction (with minor modifications) and then present the high-level ideas that are necessary to make the construction actively secure.

In their protocol, the receiver first creates a noisy encoding $(G, H, L, \mathbf{v}) \leftarrow \text{Encode}(\mathbf{x})$ (as described in Section 3, Figure 1) and sends (G, \mathbf{v}) to the sender. At this point, the locations $i \in L$ of \mathbf{v} hide a degree $\frac{\ell-1}{2}$ polynomial over the points β_1, \dots, β_n which evaluates to the input $\mathbf{x} = x_1, \dots, x_t$ in the positions $\alpha_1, \dots, \alpha_t$. The sender picks two random polynomials A and B with the restriction that $A(\alpha_i) = a_i$ and $B(\alpha_i) = b_i$ for $i \in [t]$. The degree of A is $\frac{\ell-1}{2}$, and the degree of B is $\ell - 1$.⁴ This means that B completely hides A and therefore the inputs of the sender. Now the sender simply computes $w_i = A(\beta_i) \cdot v_i + B(\beta_i)$. Sender and receiver engage in an ℓ -out-of- n OTs, and the receiver picks the ℓ positions in L . He applies H to the obtained values and interpolates a polynomial Y which evaluates in position α_i to $a_i \cdot x_i + b_i$.

We keep the generic structure of the protocol of [22] in our protocol. In order to ensure correct and consistent inputs, we have to add additional checks. The complete description is given in Figure 3, and we give a high-level description of the ideas in the following paragraph.

First, we need to ensure that the receiver can only learn ℓ values, otherwise he could potentially reconstruct part of the input. Instead of using an expensive ℓ -out-of- n OT, we let the sender create a (ρ, n) -secret sharing (remember that $\rho + \ell = n$) of a random value e and the share s_i in the i 'th OT, letting the other message offered be a random value t_i . Depending on his set L , the receiver chooses t_i or the share s_i . Then he uses the shares to reconstruct e and sends it to the sender. This in turn might leak some information on L to the sender, if he can provide an inconsistent secret sharing. We thus force the sender to commit to e and later provide an unveil. Here the sender can learn some information on L , if he cheats but is not caught, but we can use our results from the previous section to show that this leakage is tolerable. The receiver can proceed and provide the encoding \mathbf{v} , which allows the sender to compute \mathbf{w} .

Second, we have to make sure that the sender computes the correct output. In order to catch a cheating sender, we add a check to the end of the protocol. Recall that the receiver knows the output Y . He can compute another polynomial X of his input and then pick a uniformly random challenge z_R . He sends it to the sender, who has to answer with $A(z_R), B(z_R)$. Now the receiver can verify that $Y(z_R) = A(z_R)X(z_R) + B(z_R)$, i.e., the sender did not cheat in the noiseless positions. Again this leaks some information to the sender, but with the correct choice of parameters this leakage is inconsequential.

⁴ The value ℓ is fixed by the encoding, but we require that ℓ is uneven due to the fact that we have to reconstruct a polynomial of even degree $\frac{\ell-1}{2} + \frac{\ell-1}{2} = \ell - 1$, which requires ℓ values.

Security against a malicious receiver basically follows from the passively secure protocol. We only have to make sure that the extraction of his input is correct and that no information about the sender's inputs is leaked if e is incorrect. We thus mask the w_i by a one-time-pad and add the following check. This time the sender chooses z_S and the receiver has to answer with $X(z_S), Y(z_S)$, which enforces correct extraction.

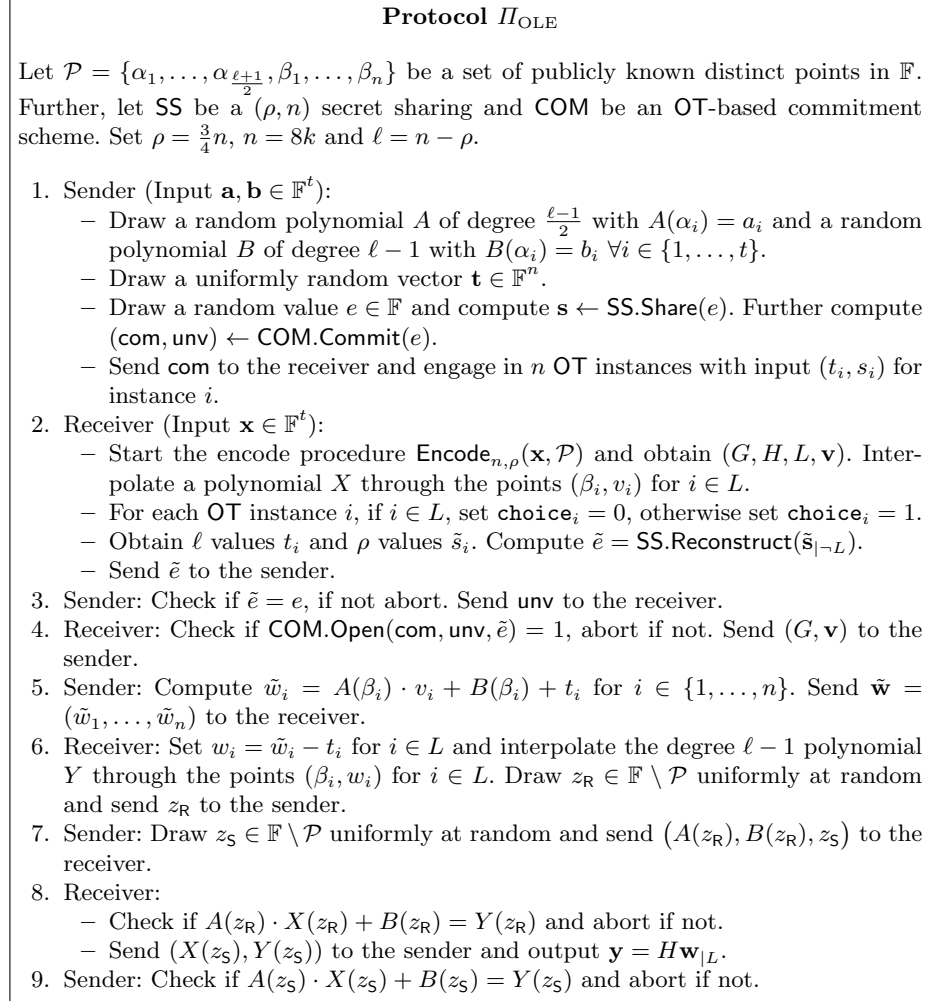


Fig. 3. Actively secure realization of \mathcal{F}_{OLE} in the OT-hybrid model.

Theorem 3. *The protocol Π_{OLE} UC-realizes \mathcal{F}_{OLE} in the OT-hybrid model with computational security.*

Proof. Corrupted sender. In the following we present a simulator \mathcal{S}_5 which provides a computationally indistinguishable simulation of Π_{OLE} to a malicious sender \mathcal{A}_5 (cf. Figure 4).

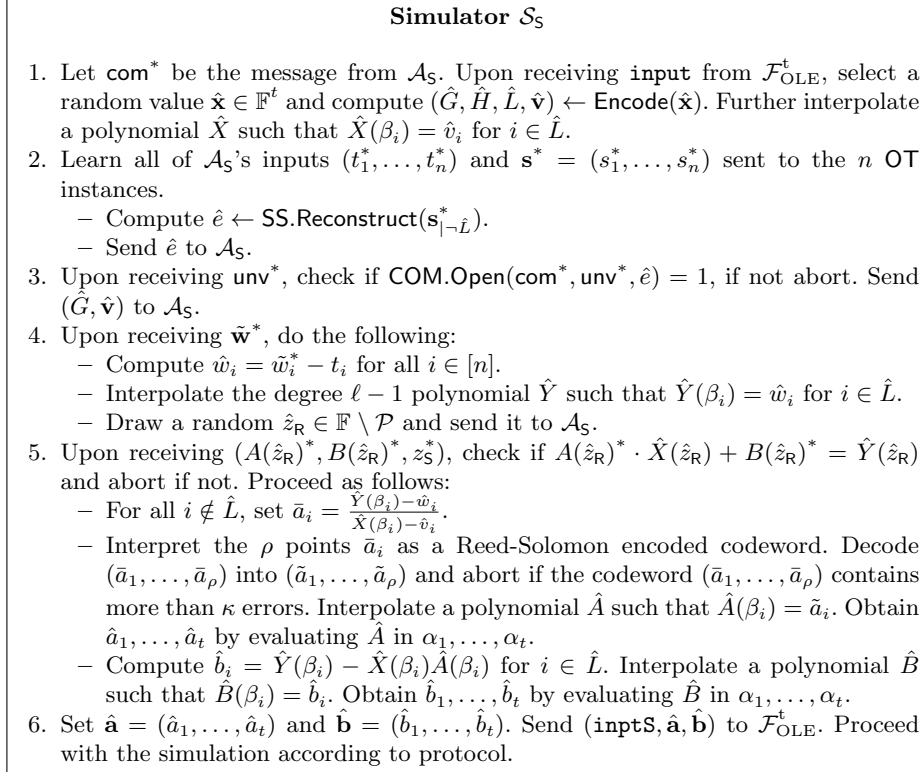


Fig. 4. Simulator against a corrupted sender in Π_{OLE} .

The main idea behind the extraction is the following. Since \mathcal{S}_5 learns all inputs into the OTs, it can use the now available noisy elements \hat{v}_i with $i \notin L$ to learn the input \mathbf{a} . The noiseless $\hat{v}_i, i \in L$ can be extrapolated to the noisy positions via a polynomial \hat{Y} (\hat{w}_i values imply a degree $\ell - 1$ polynomial for $i \in L$, and the receiver always learns ℓ values).

Ignoring for the moment that \mathcal{A}_5 might provide inconsistent inputs, the simulator now knows two values for each position $\beta_i, i \notin L$: $\hat{w}_i = a_i \cdot \hat{v}_i + b_i$ and $\hat{Y}(\beta_i)$. Therefore, assuming that \mathcal{A}_5 is honest, by computing $\hat{w}_i - \hat{Y}(\beta_i)$ the simulator gets $a_i \cdot \hat{v}_i + b_i - a_i \cdot \hat{x}_i + b_i = a_i(\hat{v}_i - \hat{x}_i)$, where \hat{x}_i is the value that his input $\hat{\mathbf{x}} \in \mathbb{F}^t$ would imply according to the encoding $\hat{\mathbf{v}}_{|L}$ on position β_i . Since the simulator knows \hat{v}_i and \hat{x}_i , it can simply compute a_i . From $\frac{\ell-1}{2} + 1$ of these points it can then compute the degree- $\frac{\ell-1}{2}$ polynomial A . From $Y = AZ + B$, it can then compute B and therefore the b_i s. For this to work we only need $\frac{\ell-1}{2} + 1$ points. Therefore, if the corrupted sender sends incorrect values in at most κ

positions $i \notin L$ and $\frac{\ell-1}{2} + 2\kappa < n$ there are still enough points to at least define a *correct* A and therefore also a *correct* $B = Y - AX$.

We now show that for every PPT environment \mathcal{Z} , the two distributions $\text{Real}_{\mathcal{H}_{\text{OLE}}}^{\mathcal{A}_5}(\mathcal{Z})$ and $\text{Ideal}_{\mathcal{F}_{\text{OLE}}}^{\mathcal{S}_5}(\mathcal{Z})$ are indistinguishable. Consider the following series of hybrid experiments.

Hybrid 0: This is the real protocol.

Hybrid 1: Identical to Hybrid 0, except that \mathcal{S}_1 extracts all inputs (s_i, t_i) input into the OT's by \mathcal{A}_5 .

Hybrid 2: Identical to Hybrid 1, except that \mathcal{S}_2 extracts the values \bar{a} as shown in Figure 4 and aborts if the check in Step 8 is passed, but $\bar{a}_1, \dots, \bar{a}_\rho$ has more than κ errors.

Hybrid 3: Identical to Hybrid 2, except that \mathcal{S}_3 encodes a random value $\hat{\mathbf{x}}$ as its input.

The indistinguishability of Hybrids 0 and 1 is immediate. We show that Hybrid 1 and Hybrid 2 are computationally indistinguishable in Lemma 1, and then we prove indistinguishability of Hybrid 2 and Hybrid 3 in Lemma 2.

Lemma 1 *Hybrids 1 and 2 are computationally indistinguishable from \mathcal{Z} 's view given that Assumption 2 holds.*

Proof. In order to prove the lemma, we have to show the following two statements.

1. \mathcal{S}_2 correctly extracts the input $\hat{\mathbf{a}}, \hat{\mathbf{b}}$, if there are less than κ errors in noisy positions.
2. The probability that \mathcal{S}_2 aborts due to more than κ errors in noisy positions is negligible in κ .

There are two ways in which \mathcal{A}_5 can cheat and prevent the correct extraction: (1) it uses an inconsistent input for a noiseless value $\hat{v}_i, i \in L$ which leads to a wrong polynomial \hat{Y} (and also an incorrect \bar{a}_i); (2) it uses an inconsistent input for a noisy value $\hat{v}_i, i \notin L$, which leads to incorrectly extracted values \bar{a}_i .

In case (1), the honest party will abort due to the check in Step 8 with overwhelming probability. It has to hold that $A(z)^* \cdot \hat{X}(z) + B(z)^* = \hat{Y}(z)$ for a uniformly chosen z . From Assumption 2 it follows that \hat{X} (and thus \hat{Y}) are unknown to \mathcal{Z} , as they would be unconditionally hidden by a completely random vector. By the fundamental theorem of algebra there are at most $\deg(\hat{Y}) = \ell - 1$ possible values z for which $A(z)^* \cdot \hat{X}(z) + B(z)^* = \hat{Y}(z)$ for incorrect $A(z)^*, B(z)^*$. Since $z_{\mathbb{R}}$ is chosen uniformly at random from \mathbb{F} , the probability that the check succeeds with incorrect $A(z)^*, B(z)^*$ is thus upper bounded by $\frac{\ell-1}{|\mathbb{F}|}$, which is negligible in the security parameter. This means that the check in Step 8 ensures that *all* the values \hat{w}_i for $i \in L$ are correct.

For case (2), we first argue that the extraction also succeeds if \mathcal{A}_5 adds less than κ errors in noisy positions (the simulator will abort if more than κ errors occur). By the choice of parameters it holds that $\rho > 3\kappa = 6\ell$, and

the simulator learns ρ values a_i that are supposed to represent a degree $\frac{\ell-1}{2}$ polynomial. Applying a standard Reed-Solomon decoder then yields the correct values a_i , i.e., if less than κ errors occur, \mathcal{S}_5 extracts the correct $\mathbf{a} \in \mathbb{F}^t$ (and thus also the correct $\mathbf{b} \in \mathbb{F}^t$).

This shows that as long as there are less than κ errors in noisy positions, the extracted values are correct.

We claim that a \mathcal{Z} that places more than κ errors in noisy positions breaks Assumption 2. The scenario of \mathcal{Z} in the simulation is identical to the game $\mathcal{G}_{\text{ident}}$: \mathcal{Z} gets an encoding $\mathbf{v} \leftarrow \text{Encode}_{n,\rho}(\mathbf{x})$ with ρ noisy positions and has to output a set of positions $Q \subseteq [n]$ such that $Q \subseteq R$ and $|Q| \geq \kappa$.

As discussed in Section 3, we can assume that $\text{Encode}_{n,n/2}$ yields encodings that are indistinguishable from $\text{Encode}_{n,n}$, i.e., truly random strings. In order to meet the requirements of Theorem 2, it therefore has to hold that $\sigma = n \frac{\rho - \kappa}{n - \kappa} \geq \frac{n}{2}$. Thus, we get that ρ has to be larger than $\frac{n + \kappa}{2}$, which by our choice of parameters is the case. Thus the claim directly follows from Theorem 2. \square

Lemma 2 *Hybrids 2 and 3 are computationally indistinguishable from \mathcal{Z} 's view given that Assumption 2 holds and COM is a UC commitment scheme.*

Proof. Assume that there exists a PPT \mathcal{Z} that distinguishes Hybrids 2 and 3 with non-negligible probability ϵ . We will show that \mathcal{Z} breaks Assumption 2 with non-negligible probability.

We have to consider all the messages that \mathcal{A}_5 receives during a protocol run. First note that \mathcal{S}_5 (resp. R) outputs either e or aborts in Step 4. Assume for the sake of contradiction that \mathcal{A}_5 manages to create two secret sharings $s_{1,1}, \dots, s_{1,n}$ and $s_{2,1}, \dots, s_{2,n}$ for values e, e' such that R outputs both of e or e' with non-negligible probability ϵ without aborting depending on the set L and L' , respectively. Then we create an adversary \mathcal{B} from \mathcal{Z} that breaks the binding property of COM. \mathcal{B} simulates the protocol and learns all values s_i^* , then draws two uniformly random sets L, L' . \mathcal{B} samples via L and L' two subsets of secret sharings that reconstruct to e and e' , respectively, with non-negligible probability. It must hold for both values that $\text{COM.Open}(\text{com}, \text{unv}, e) = \text{COM.Open}(\text{com}, \text{unv}, e') = 1$, otherwise \mathcal{B} aborts as the real R would. Since \mathcal{A}_5 achieves that R outputs e or e' with non-negligible probability, \mathcal{B} outputs com, e, e' with non-negligible probability to the binding experiment and thereby breaks the binding property of COM.

The next message he receives is the encoding. Recall that the choice bits into the OTs are derived from the set L of the encoding, i.e., a cheating \mathcal{A}_5 might try to use inconsistent inputs (e.g., incorrect s_i values in positions that are supposedly not in L) in the OT such that R aborts depending on the set L . However, \mathcal{A}_5 has to cheat before knowing the encoding \mathbf{v} and as shown above always learns the same e , thus he can obtain at most 1 bit of leakage, namely whether the cheating was detected or not. We will now show that the leakage in Step 4 does not help \mathcal{Z} to distinguish. The situation for a malicious \mathcal{Z} is identical to game $\mathcal{G}_{\text{leak}}$. First, \mathcal{A}_5 has to decide on a set of values which he believes are not in L . Then he is informed (via a successful check) that his guess was correct, and given the encoding. Now he has to decide whether he is given a random

encoding or not. We can directly apply Theorem 1, given that $\rho \geq \frac{2}{3}n$, and get that \mathcal{Z} 's distinguishing advantage is negligible.

After learning \mathbf{v} , \mathcal{A}_S has to compute the values \mathbf{w} , which are checked in Step 8. By cheating in noisy positions, Step 8 will succeed, but \mathcal{A}_S learns some noisy positions by learning the bit whether the check succeeded. This case is more involved than the above step, since now \mathcal{A}_S can decide on the set Q after seeing the encoding \mathbf{v} . We argue that the distinguishing advantage of \mathcal{Z} remains negligible. It is obvious that \mathcal{A}_S can always find $O(\log \kappa)$ noisy positions with polynomial probability simply by guessing. However, Theorem 2 guarantees that in this scenario \mathcal{A}_S cannot find more than $O(\log \kappa)$ noisy positions, if $Y^\sigma \approx Y^n$ for $\sigma = \frac{n\kappa}{n-\rho-\kappa}$. From Theorem 1 we know that if $Q = O(\log \kappa)$ and $\sigma > \frac{2}{3}n$, then $Y^\sigma \approx Y^n$. Combined, we have that for $\rho = n - \frac{\kappa}{2}$, \mathcal{A}_S cannot find more than $O(\log n)$ noisy positions and the distinguishing advantage of \mathcal{Z} is negligible. This concludes the proof. \square

Corrupted receiver. In the following we present a simulator \mathcal{S}_R which provides a statistically indistinguishable simulation of Π_{OLE} to a malicious receiver \mathcal{A}_R (cf. Figure 5). Conceptually the simulation is straight forward. The simulator learns all choice bits and thus can reconstruct the set L , which is sufficient to decode the codeword \mathbf{v} . Knowing X , \mathcal{S}_R can easily derive consistent inputs A, B . Care has to be taken since \mathcal{A}_R obtains one additional pair of values related to the polynomials A and B , thus he can tamper with the extraction. In a little more detail, he obtains one more value than necessary to reconstruct Y and can therefore play both with the degree of his input as well as with the correctness of L and \mathbf{v} . We describe and analyze a subtle attack in Lemma 4, which makes the analysis a bit more complex.

We now show the indistinguishability of the simulation in a series of hybrid experiments. For every PPT environment \mathcal{Z} , the two distributions $\text{Real}_{\Pi_{\text{OLE}}}^{\mathcal{A}_S}(\mathcal{Z})$ and $\text{Ideal}_{\mathcal{F}_{\text{OLE}}}^{\mathcal{S}_S}(\mathcal{Z})$ are indistinguishable.

Hybrid 0: This is the real protocol.

Hybrid 1: Identical to Hybrid 0, except that \mathcal{S}_1 extracts all inputs choice_i input into OT by \mathcal{A}_R .

Hybrid 2: Identical to Hybrid 1, except that \mathcal{S}_2 aborts if \mathcal{A}_R passes the check in Step 3, although he selects less than ρ values s_i .

Hybrid 3: Identical to Hybrid 2, except that \mathcal{S}_3 reconstructs \hat{X} as shown in Figure 5 and aborts if $\hat{Y}(\hat{z}_S) \neq Y^*(\hat{z}_S)$, $\hat{X}(\hat{z}_S) \neq X^*(\hat{z}_S)$ or $\hat{Y} = R$.

Indistinguishability of Hybrids 0 and 1 is trivial. We show the indistinguishability of Hybrids 1 and 2 in Lemma 3, based on the privacy of the secret sharing and the hiding property of the commitment. In Lemma 4 we show that we can always extract the correct input of \mathcal{A}_R and thus Hybrid 2 and Hybrid 3 are statistically indistinguishable.

Lemma 3 *Hybrids 1 and 2 are statistically indistinguishable from \mathcal{Z} 's view given that SS is a perfectly private secret sharing and COM is a statistically hiding commitment scheme.*

Simulator \mathcal{S}_R

1. Upon receiving a message **input** from \mathcal{F}_{OLE}^t , simulate the first part of Π_{OLE} with random inputs.
 - Draw a uniformly random vector $\hat{\mathbf{t}} \in \mathbb{F}^n$ and a random value $\hat{e} \in \mathbb{F}$.
 - Compute $\hat{\mathbf{s}} \leftarrow \text{SS.Share}(\hat{e})$ and $(\widehat{\text{com}}, \widehat{\text{unv}}) \leftarrow \text{COM.Commit}(\hat{e})$.
 - Send $\widehat{\text{com}}$ to \mathcal{A}_R and engage in n OT instances with input (\hat{t}_i, \hat{s}_i) in OT i .
2. Learn all choice bits $(\text{choice}_1^*, \dots, \text{choice}_n^*)$ of \mathcal{A}_R from the n OT instances. Reconstruct \hat{L} as follows: for each $i \in [n]$, if $\text{choice}_i^* = 0$ then $i \in \hat{L}$.
3. Upon receiving e^* , check if $\hat{e} = e^*$, otherwise abort. Send $\widehat{\text{unv}}$ to \mathcal{A}_R .
4. Upon receiving (G^*, \mathbf{v}^*) from \mathcal{A}_R , proceed as follows.
 - (a) Let $\deg(P_{\hat{L}})$ denote the degree of the polynomial defined by $\mathbf{v}_{|\hat{L}}$.
 - If $|\hat{L}| = \ell - 1$, interpolate the polynomial $P_{\hat{L}}$ defined over $\mathbf{v}_{|\hat{L}}$. If $\deg(P_{\hat{L}}) \leq \frac{\ell-1}{2}$, set $\hat{X} = P_{\hat{L}}$.
 - If $|\hat{L}| = \ell$, interpolate the polynomial $P_{\hat{L}}$ defined over $\mathbf{v}_{|\hat{L}}$. If $\deg(P_{\hat{L}}) \leq \frac{\ell-1}{2}$, set $\hat{X} = P_{\hat{L}}$. If $\deg(P_{\hat{L}}) > \frac{\ell+1}{2}$, try for all $\hat{i} \in \hat{L}$ if it holds that for $\hat{L}' = \hat{L} \setminus \hat{i}$, $\deg(P_{\hat{L}'}) \leq \frac{\ell-1}{2}$. If such an \hat{i} exists, set $\hat{X} = P_{\hat{L}'}$ and $\hat{L} = \hat{L}'$.
 - (b) Compute $\hat{x}_i = \hat{X}(\alpha_i)$, $i \in [t]$ and send $(\text{inputR}, \hat{\mathbf{x}})$ to \mathcal{F}_{OLE}^t . Let $(\text{output}, \hat{\mathbf{y}})$ be the result. Pick a random polynomial \hat{Y} such that $\deg(\hat{Y}) = \deg(\hat{X}) + \frac{\ell-1}{2}$ and $\hat{Y}(\alpha_i) = \hat{y}_i$, $i \in [t]$. If no \hat{X} was extracted in Step 4a, set \hat{Y} to be a random degree $\ell - 1$ polynomial R .
 - (c) For $i \in \hat{L}$, set $\hat{w}_i = \hat{Y}(\beta_i) + t_i$, otherwise pick a uniform \hat{w}_i and send $\hat{\mathbf{w}}$ to \mathcal{A}_R .
5. Upon receiving z_R^* , draw $\hat{z}_S \in \mathbb{F}$ and proceed as follows:
 - If $\hat{Y} \neq R$, compute $\hat{X}(z_R^*)$, $\hat{Y}(z_R^*)$ and sample a random $\hat{b} \in \mathbb{F}$. Set $\hat{a} = \frac{\hat{Y}(z_R^*) - \hat{b}}{\hat{X}(z_R^*)}$ and send $(\hat{a}, \hat{b}, \hat{z}_S)$ to \mathcal{A}_R .
 - If $\hat{Y} = R$, pick random $\hat{a}, \hat{b} \in \mathbb{F}$ and send $(\hat{a}, \hat{b}, \hat{z}_S)$ to \mathcal{A}_R .
6. Upon receiving $(X^*(\hat{z}_S), Y^*(\hat{z}_S))$, proceed as follows:
 - If $\hat{Y} \neq R$, check if $Y^*(\hat{z}_S) = \hat{Y}(\hat{z}_S)$ and $X^*(\hat{z}_S) = \hat{X}(\hat{z}_S)$ and abort if not.
 - If $\hat{Y} = R$, abort.

Fig. 5. Simulator against a corrupted receiver in Π_{OLE} .

Proof. Assume for the sake of contradiction that there exists an environment \mathcal{Z} that distinguishes the hybrids, i.e., \mathcal{Z} has to make \mathcal{S}_2 abort with non-negligible probability ε . We will construct from \mathcal{Z} an adversary \mathcal{B} that breaks the hiding property of COM with non-negligible probability. \mathcal{B} simulates the protocol exactly like \mathcal{S}_2 , but creates a secret sharing of a random r and picks two random e, e' , which he sends to the hiding experiment. The hiding experiment returns a commitment com on one of these values. Then \mathcal{B} integrates the commitment and secret sharing into the simulation and checks whether \mathcal{Z} inputs less than ρ values $\text{choice}_i = 1$ into OT, otherwise \mathcal{B} aborts. Since SS is a perfectly private secret sharing and \mathcal{Z} obtains less than ρ values s_i , these values leak nothing about r and the simulation of \mathcal{B} is indistinguishable from \mathcal{S}_2 's simulation. Let now e^* be \mathcal{Z} 's answer in the simulated protocol. \mathcal{B} simply forwards e^* to the hiding experiment. Since it has to hold that $e^* = e$ or $e^* = e'$ with non-negligible probability ε (otherwise the check in Step 3 fails), \mathcal{B} breaks the hiding property

of COM with the same probability. From this contradiction it follows that \mathcal{A}_R learns at most ℓ values t_i through OT. \square

Lemma 4 *Hybrids 2 and 3 are statistically indistinguishable from \mathcal{Z} 's view.*

Proof. In order to distinguish Hybrids 2 and 3, \mathcal{Z} must pass the check in Step 9, even though it holds that \mathcal{S}_3 picked a random polynomial R (allowing to distinguish the simulation from the real protocol). First note that the result \mathbf{w} always defines a polynomial of degree $\ell - 1$ if \mathcal{A}_R 's input polynomial has degree less than $\frac{\ell-1}{2}$. As we know from Lemma 3, \mathcal{A}_R learns at most ℓ values through the OTs and then one additional pair (a, b) via the check in Step 9.

Before we look at the details of the extraction, let us first describe an generic adversarial strategy that we have to cover. The adversary gets 1 free query and might try to use this query to prevent extraction. Say he picks a polynomial of degree $\frac{\ell-1}{2}$, but only uses $\ell - 1$ values of L . In the choice phase, he selects a random index $i^* \notin L$ and sets $\text{choice}_{i^*} = 0$, i.e., \mathcal{S}_3 will assume this index is also in L . Towards the same goal, \mathcal{A}_R can simply set the value v_i for a random index i to a random value. \mathcal{S} will then extract a wrong polynomial (with degree greater than $\frac{\ell+1}{2}$), while \mathcal{A}_R can still reconstruct Y via the additional values. However, since \mathcal{A}_R can only add exactly 1 random element, \mathcal{S}_3 can identify the index by trying for each $i \in L$ whether the set $L' = L \setminus i$ defines a polynomial of degree $\frac{\ell-1}{2}$ over the v_i . Here it is essential that there are no two sets L_1, L_2 with $|L_1| = \ell - 1, |L_2| = \ell$ such that $L_1 \subset L_2$ and $\deg(P_{L_1}) = \frac{\ell-1}{2}, \deg(P_{L_2}) = \frac{\ell+1}{2}$, i.e., there is only one possible index i that can be removed. This follows from the fact that the polynomial $P = P_{L_2} - P_{L_1}$ has only $\frac{\ell+1}{2}$ roots, but L_1 and L_2 have to agree on $\ell - 1$ positions. If that scenario were possible, \mathcal{S}_3 would not be able to distinguish these cases.

Let in the following $\deg(P_{\hat{L}})$ denote the degree of the polynomial that is defined by the points v_i for $i \in \hat{L}$.

- $|\hat{L}| \leq \ell - 2$: \mathcal{A}_R obtains at most $\ell - 2 + 1$ points, but Y is of degree $\ell - 1$ and thus underspecified. Clearly \mathcal{A}_R 's probability of answering the check in Step 9 with a correct $X^*(z_S), Y^*(z_S)$ is negligible in \mathbb{F} . Since \mathcal{S}_3 aborts as well, Hybrids 2 and 3 are indistinguishable in this case.
- $|\hat{L}| = \ell - 1$: In this case it holds that $\hat{Y} = R$ only if $\deg(P_{\hat{L}}) \geq \frac{\ell+1}{2}$.
 - $\deg(P_{\hat{L}}) = \frac{\ell-1}{2}$: In this case \mathcal{A}_R can reconstruct Y and pass the check in Step 9, but \mathcal{S}_3 extracts the correct \hat{X} . From the argument above, there cannot exist another polynomial X' that fits with the set \hat{L} and thus Hybrids 2 and 3 are indistinguishable.
 - $\deg(P_{\hat{L}}) = \frac{\ell+1}{2}$: In this case \mathcal{A}_R obtains $\ell - 1 + 1$ points, but the resulting Y is of degree $\frac{\ell-1}{2} + \frac{\ell+1}{2} = \ell$, i.e., \mathcal{A}_R needs $\ell + 1$ points to reconstruct Y . By the same argument as above, Hybrids 2 and 3 are indistinguishable.
 - $\deg(P_{\hat{L}}) > \frac{\ell+1}{2}$: In this case \mathcal{A}_R can behave as described above, i.e., add a random i to the set \hat{L} and thereby artificially increase $\deg(P_{\hat{L}})$. But since $|\hat{L}| = \ell - 1$, removing an additional value from \hat{L} leads to the case $|\hat{L}| \leq \ell - 2$ and thus indistinguishability of Hybrids 2 and 3.

- $|\hat{L}| = \ell$: In this case it holds that $\hat{Y} = R$ only if $\deg(P_{\hat{L}}) > \frac{\ell+1}{2}$ and no index i can be identified to reduce $\deg(P_{\hat{L}})$ to $\frac{\ell-1}{2}$.
 - $\deg(P_{\hat{L}}) = \frac{\ell-1}{2}$: In this case \mathcal{A}_R can reconstruct Y and pass the check in Step 9, but \mathcal{S}_3 extracts the correct \hat{X} .
 - $\deg(P_{\hat{L}}) = \frac{\ell+1}{2}$: In this case \mathcal{A}_R obtains $\ell + 1$ points, and the resulting Y is of degree $\frac{\ell-1}{2} + \frac{\ell+1}{2} = \ell$. Thus \mathcal{A}_R can reconstruct Y and pass the check, but \mathcal{S}_3 extracts the correct \hat{X} .
 - $\deg(P_{\hat{L}}) > \frac{\ell+1}{2}$: In this case \mathcal{A}_R can behave as described above, i.e., add a random i to the set \hat{L} and thereby artificially increase $\deg(P_{\hat{L}})$. Removing an additional value from \hat{L} leads to the case $|\hat{L}| = \ell - 1$, i.e., \mathcal{S}_3 will simulate correctly. Otherwise, \mathcal{S}_3 will abort, but \mathcal{A}_R cannot reconstruct Y and thus fails the check in Step 9.
- $|\hat{L}| > \ell$: \mathcal{S}_3 aborts, and from Lemma 3 it follows that Hybrids 2 and 3 are indistinguishable.

The correctness of the simulation follows from the fact that either \mathcal{S}_3 extracts the correct input \hat{X} , or the check in Step 9 fails with overwhelming probability, in which case $\hat{X} = R$. Thus, the event that \mathcal{Z} can provoke an abort is negligible, i.e., Hybrids 2 and 3 are indistinguishable. \square

This concludes the proof. \square

6 Efficient Oblivious Polynomial Evaluation

The ideal functionality \mathcal{F}_{OPE} for OPE is depicted in Figure 6. It allows the sender S to input a polynomial P and the receiver R to input $\alpha \in \mathbb{F}$. In the

Functionality \mathcal{F}_{OPE}
<ol style="list-style-type: none"> 1. Upon receiving a message (inputS, P) from S where $P \in \mathbb{F}[X]$, verify that there is no stored tuple, else ignore that message. Store P and send a message (input) to \mathcal{A}. 2. Upon receiving a message (inputR, α) from R with $\alpha \in \mathbb{F}$, verify that there is no stored tuple, else ignore that message. Store α and send a message (input) to \mathcal{A}. 3. Upon receiving a message (deliver, S) from \mathcal{A}, check if both P and α are stored, else ignore that message. Send (delivered) to S. 4. Upon receiving a message (deliver, R) from \mathcal{A}, check if both P and α are stored, else ignore that message. Send (output, $P(\alpha)$) to R.

Fig. 6. Ideal functionality for an oblivious polynomial evaluation.

remainder of this section we will establish the following theorem.

Theorem 4. *There exists a (constant-round) protocol Π_{OPE} that UC-realizes \mathcal{F}_{OPE} with unconditional security in the $\mathcal{F}_{\text{OLE}}^t$ -hybrid model. In particular, for a polynomial P of degree d , $t = d + 2$.*

Our roadmap is as follows. We first show how to reduce \mathcal{F}_{OPE} to an intermediate OLE-based functionality $\mathcal{F}_{\text{OLE}}^{t,1}$. After establishing this we present an efficient reduction of $\mathcal{F}_{\text{OLE}}^{t,1}$ to $\mathcal{F}_{\text{OLE}}^t$ (or \mathcal{F}_{OLE}).

We follow the generic idea of Naor and Pinkas [30] of using the linearization technique from [17] to construct an oblivious polynomial evaluation protocol. They decompose a polynomial P of degree d into d linear functions. These functions can then be evaluated using our OLE with input α for each of the functions, and the receiver can reconstruct the value $P(\alpha)$. We state the lemma here, a proof can be found in [30] and the full version of this paper [15].

Lemma 3 ([17]). *For every polynomial P of degree d , there exist d linear polynomials P_1, \dots, P_d , such that an OPE of P can be reduced to a parallel execution of an OLE of each of P_1, \dots, P_d , where all the linear polynomials are evaluated at the same point.*

In the semi-honest case, this approach directly works with the $\mathcal{F}_{\text{OLE}}^t$ for $t = d$. But unlike the construction of [30], our batch-OLE does not enforce the receiver to use the same input α in all of the OLEs. Therefore we cannot use the reduction of [30] that shows malicious security against a receiver. In particular, a malicious receiver might learn some non-trivial linear combinations of the coefficients of P .

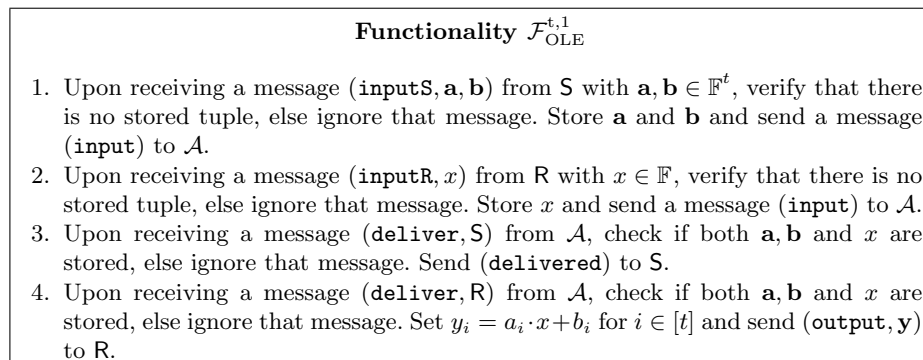


Fig. 7. Ideal functionality for a $(t, 1)$ -oblivious linear function evaluation.

Reducing \mathcal{F}_{OPE} to $\mathcal{F}_{\text{OLE}}^{t,1}$. As a first step we reduce OPE to a variant of OLE where the receiver has only one input $x \in \mathbb{F}$, while the sender inputs two vectors \mathbf{a}, \mathbf{b} . This is depicted in Figure 7.

The reduction of \mathcal{F}_{OPE} to $\mathcal{F}_{\text{OLE}}^{t,1}$ is straightforward, given Lemma 3. The sender decomposes his polynomial P into d linear functions f_1, \dots, f_d with coefficients (a_i, b_i) and inputs these into $\mathcal{F}_{\text{OLE}}^{d,1}$. The receiver chooses his input α and obtains d linear evaluations, from which he can reconstruct $P(\alpha)$. The number of OLEs required is only dependent on the realization of $\mathcal{F}_{\text{OLE}}^{d,1}$.

Lemma 4. *The protocol Π_{OPE} UC-realizes \mathcal{F}_{OPE} in the $\mathcal{F}_{\text{OLE}}^{d,1}$ -hybrid model with unconditional security.*

Proof. The security of Π_{OPE} is immediate: the simulator simulates $\mathcal{F}_{\text{OLE}}^{d,1}$ and learns all inputs, which it simply forwards to \mathcal{F}_{OPE} (and reconstructs if necessary). The correctness of the decomposition of P follows from Lemma 3. \square

Note that by taking our approach, we also remove the need for the stronger assumption of [30], while having a comparable efficiency in the resulting protocol.

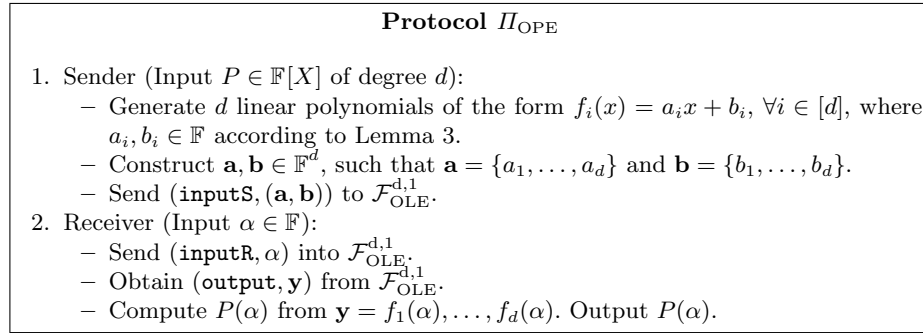


Fig. 8. Reduction of \mathcal{F}_{OPE} to $\mathcal{F}_{\text{OLE}}^{d,1}$.

Reducing $\mathcal{F}_{\text{OLE}}^{t,1}$ to $\mathcal{F}_{\text{OLE}}^{t+2}$. As a second step, we need to realize $\mathcal{F}_{\text{OLE}}^{t,1}$ from $\mathcal{F}_{\text{OLE}}^t$. Döttling et al. [11] describe a black-box protocol that realizes $\mathcal{F}_{\text{OLE}}^{t,1}$ from $\mathcal{F}_{\text{OLE}}^t$ (or our batch variant) with unconditional UC-security. The protocol has a constant multiplicative overhead of $2 + \varepsilon$ in the number of OLEs, and works for any field \mathbb{F} . While this protocol basically solves our problem, we propose a more efficient variant that makes essential use of the fact that we only consider a large field \mathbb{F} . Our new approach requires only two additional OLEs and thus has overhead $1 + \epsilon$.

Our solution for $\mathcal{F}_{\text{OLE}}^{t,1}$ is as follows. Let $\mathbf{a}, \mathbf{b} \in \mathbb{F}^t$ be given as input to the sender. It now needs to choose one additional pair of inputs (a_{t+1}, b_{t+1}) such that $\sum_{i=1}^{t+1} a_i = 0$ and b_{t+1} is uniformly random in \mathbb{F} . The sender inputs $\mathbf{a}', \mathbf{b}' \in \mathbb{F}^{t+1}$ into $\mathcal{F}_{\text{OLE}}^{t+1}$, while the receiver inputs $\mathbf{x}' = (x, \dots, x) \in \mathbb{F}^{t+1}$. Now the receiver locally computes $c = \sum_{i=1}^{t+1} y_i = \sum_{i=1}^{t+1} a_i x + \sum_{i=1}^{t+1} b_i = \sum_{i=1}^{t+1} b_i$ and sends a commitment to c to the sender. This commitment can also be based on OLE, even in such a way that we can use $\mathcal{F}_{\text{OLE}}^{t+2}$ by precomputing the commitment (a detailed description is given in the full version [15]). The sender answers with $c' = \sum_{i=1}^{t+1} b_i$, which the receiver can verify. This makes sure that the sender chose \mathbf{a}' correctly, while c' itself does not give the receiver any new information. Now the receiver unveils, which shows the sender whether the receiver used the same x in each invocation. There is one small problem left: if the receiver cheated, he will be caught, but he might still learn some information about the sender's

inputs that cannot be simulated. In order to solve this issue, we let \mathbf{a}' and \mathbf{b}' be uniformly random and then replace these with the inputs after the check succeeded. A detailed description of the protocol is given in Figure 9.

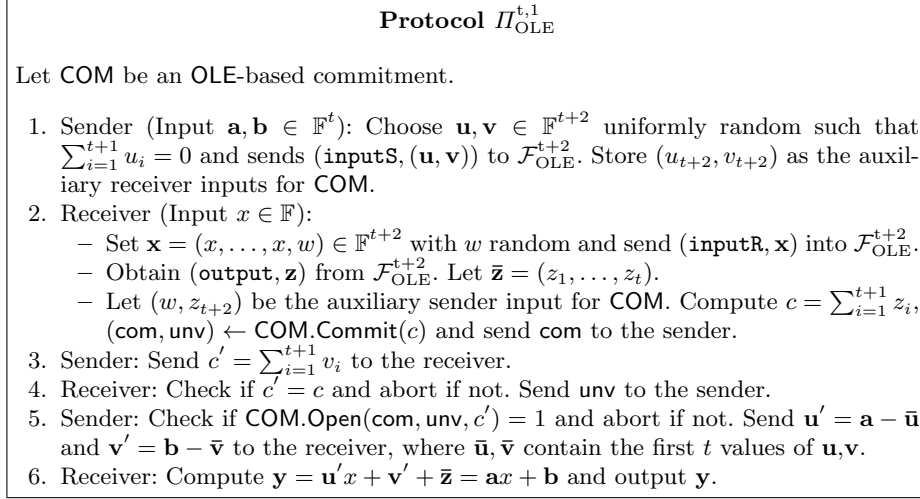


Fig. 9. Reduction of $\mathcal{F}_{\text{OLE}}^{t,1}$ to $\mathcal{F}_{\text{OLE}}^{t+2}$.

Lemma 5. *The protocol $\Pi_{\text{OLE}}^{t,1}$ UC-realizes $\mathcal{F}_{\text{OLE}}^{t,1}$ in the $\mathcal{F}_{\text{OLE}}^{t+2}$ -hybrid model with unconditional security.*

Proof. Corrupted sender. The simulator \mathcal{S}_S simulates $\mathcal{F}_{\text{OLE}}^{t+2}$ for the corrupted sender \mathcal{A}_S . It extracts all the inputs, namely $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$. We do not need to extract the commitment, which also uses $\mathcal{F}_{\text{OLE}}^{t+2}$. \mathcal{S}_S sends a commitment to $\hat{c} = \sum_{i=1}^{t+1} \hat{v}_i$ to the receiver. If it holds that $\sum_{i=1}^{t+1} u_i \neq 0$, but the check in Step 4 succeeds, \mathcal{S}_S aborts. Otherwise, it computes $\hat{\mathbf{a}} = \hat{\mathbf{u}}^{t*} + \hat{\mathbf{u}}$ and $\hat{\mathbf{b}} = \hat{\mathbf{v}}^{t*} + \hat{\mathbf{v}}$ and inputs the first t elements of each into $\mathcal{F}_{\text{OLE}}^{t,1}$.

First note that if $\sum_{i=1}^{t+1} u_i \neq 0$, the commitment com contains an incorrect value. As long as the receiver always aborts in this case, the hiding property of COM guarantees indistinguishability of the simulation. So the only way that a malicious environment \mathcal{Z} can distinguish the simulation from the real protocol is by forcing an abort. Note that if $\sum_{i=1}^{t+1} u_i = e \neq 0$, then c depends on x and is thus uniformly distributed, since

$$c = \sum_{i=1}^{t+1} z_i = \sum_{i=1}^{t+1} u_i x + \sum_{i=1}^{t+1} v_i = ex + c'.$$

Thus, the probability that $c' = c$ is negligible.

Corrupted receiver. The simulator \mathcal{S}_R against the corrupted receiver \mathcal{A}_R simulates $\mathcal{F}_{\text{OLE}}^{t+2}$ and learns $\hat{\mathbf{x}}$. It chooses $\hat{\mathbf{u}}, \hat{\mathbf{v}} \in \mathbb{F}^{t+2}$ according to $\Pi_{\text{OLE}}^{t,1}$, and computes

$\hat{\mathbf{z}} \in \mathbb{F}^{t+2}$, where $\hat{z}_i = \hat{u}_i \hat{x}_i + \hat{v}_i \forall i \in [1, t+2]$. \mathcal{S}_R sends $\hat{\mathbf{z}}$ to \mathcal{A}_R . After receiving the commitment, \mathcal{S}_R sends $\hat{c}' = \sum_{i=1}^{t+1} \hat{z}_i$. It aborts if the commitment unveils correctly, even though $x_i \neq x_j$ for some $i, j \in [t+1]$. If that is not the case, it inputs \hat{x} into $\mathcal{F}_{OLE}^{t,1}$ and obtains \mathbf{y} . \mathcal{S}_R picks $\hat{\mathbf{v}}' \in \mathbb{F}^t$ uniformly at random, sets $\hat{u}'_i = \frac{y_i - \hat{z}_i - \hat{v}'_i}{x} \forall i \in [t]$. It sends $\hat{\mathbf{u}}', \hat{\mathbf{v}}'$ to \mathcal{A}_R .

For an honest receiver, the check in Step 5 always succeeds. A malicious \mathcal{Z} can only distinguish between the simulation and the real protocol by producing a correct commitment on c , even though $x_i \neq x_j$ for some $i, j \in [t+1]$. Since the commitment is binding, \mathcal{A}_R must commit to some value c before seeing \hat{c}' . Let w.l.o.g. $x_j = (x + e) \neq x$ for some j . Then we have

$$c = \sum_{i=1}^{t+1} z_i = \sum_{\substack{i=1 \\ i \neq j}}^{t+1} u_i x + u_j (x + e) + \sum_{i=1}^{t+1} v_i = \sum_{i=1}^{t+1} u_i x + u_j e + \sum_{i=1}^{t+1} v_i = \hat{c}' + u_j e.$$

But this means that \hat{c}' is uniformly distributed from \mathcal{A}_R 's point of view, because u_j is chosen uniformly and unknown to \mathcal{A}_R . As a consequence, the probability that \mathcal{Z} can distinguish the simulation from the real protocol is negligible. \square

Combining the results from this section we get that \mathcal{F}_{OPE} for a polynomial P of degree d requires $\mathcal{F}_{OLE}^{d,1}$, which in turn can be based on \mathcal{F}_{OLE}^{d+2} . This establishes Theorem 4.

Remark 1. It is possible to evaluate several polynomials in parallel with the batch-OLE functionality, given that t is chosen of appropriate size. Then, for each polynomial the above described protocol is carried out (including making sure that the receiver uses the same α in all OLEs relevant to the respective polynomial).

References

1. Applebaum, B., Ishai, Y., Kushilevitz, E.: How to garble arithmetic circuits. In: Ostrovsky, R. (ed.) 52nd FOCS. pp. 120–129. IEEE Computer Society Press (Oct 2011)
2. Beerliová-Trubíniová, Z., Hirt, M.: Perfectly-secure MPC with linear communication complexity. In: TCC 2008. LNCS, vol. 4948, pp. 213–230. Springer (2008)
3. Bleichenbacher, D., Nguyen, P.Q.: Noisy polynomial interpolation and noisy Chinese remaindering. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 53–69. Springer, Heidelberg (May 2000)
4. Boneh, D.: Finding smooth integers in short intervals using CRT decoding. In: 32nd ACM STOC. pp. 265–272. ACM Press (May 2000)
5. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: 42nd FOCS. pp. 136–145. IEEE Computer Society Press (Oct 2001)
6. Cascudo, I., Damgård, I., David, B., Döttling, N., Nielsen, J.B.: Rate-1, linear time and additively homomorphic UC commitments. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part III. LNCS, vol. 9816, pp. 179–207. Springer, Heidelberg (Aug 2016)

7. Chang, Y.C., Lu, C.J.: Oblivious polynomial evaluation and oblivious neural learning. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 369–384. Springer, Heidelberg (Dec 2001)
8. Damgård, I., Pastro, V., Smart, N.P., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 643–662. Springer, Heidelberg (Aug 2012)
9. David, B.M., Nishimaki, R., Ranellucci, S., Tapp, A.: Generalizing efficient multiparty computation. In: Lehmann, A., Wolf, S. (eds.) ICITS 15. LNCS, vol. 9063, pp. 15–32. Springer, Heidelberg (May 2015)
10. Döttling, N., Kraschewski, D., Müller-Quade, J.: David & Goliath oblivious affine function evaluation - asymptotically optimal building blocks for universally composable two-party computation from a single untrusted stateful tamper-proof hardware token. Cryptology ePrint Archive, Report 2012/135 (2012), <http://eprint.iacr.org/2012/135>
11. Döttling, N., Kraschewski, D., Müller-Quade, J.: Statistically secure linear-rate dimension extension for oblivious affine function evaluation. In: Smith, A. (ed.) ICITS 12. LNCS, vol. 7412, pp. 111–128. Springer, Heidelberg (Aug 2012)
12. Franklin, M.K., Yung, M.: Communication complexity of secure computation (extended abstract). In: 24th ACM STOC. pp. 699–710. ACM Press (May 1992)
13. Freedman, M.J., Ishai, Y., Pinkas, B., Reingold, O.: Keyword search and oblivious pseudorandom functions. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 303–324. Springer, Heidelberg (Feb 2005)
14. Freedman, M.J., Nissim, K., Pinkas, B.: Efficient private matching and set intersection. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 1–19. Springer, Heidelberg (May 2004)
15. Ghosh, S., Nielsen, J.B., Nilges, T.: Maliciously secure oblivious linear function evaluation with constant overhead. IACR Cryptology ePrint Archive 2017, 409 (2017), <http://eprint.iacr.org/2017/409>
16. Gilboa, N.: Two party RSA key generation. In: Wiener, M.J. (ed.) CRYPTO’99. LNCS, vol. 1666, pp. 116–129. Springer, Heidelberg (Aug 1999)
17. Gilboa, N.: Topics in private information retrieval. Ph.D. thesis, Thesis (Doctoral)–Technion - Israel Institute of Technology, Faculty of Computer Science, 2001, Haifa (2001)
18. Hazay, C.: Oblivious polynomial evaluation and secure set-intersection from algebraic PRFs. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 90–120. Springer, Heidelberg (Mar 2015)
19. Hazay, C., Lindell, Y.: Efficient oblivious polynomial evaluation with simulation-based security. Cryptology ePrint Archive, Report 2009/459 (2009), <http://eprint.iacr.org/2009/459>
20. Ishai, Y., Kilian, J., Nissim, K., Petrank, E.: Extending oblivious transfers efficiently. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 145–161. Springer, Heidelberg (Aug 2003)
21. Ishai, Y., Prabhakaran, M., Sahai, A.: Founding cryptography on oblivious transfer - efficiently. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 572–591. Springer, Heidelberg (Aug 2008)
22. 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 (Mar 2009)
23. Katz, J.: Universally composable multi-party computation using tamper-proof hardware. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 115–128. Springer, Heidelberg (May 2007)

24. Keller, M., Orsini, E., Scholl, P.: Actively secure OT extension with optimal overhead. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 724–741. Springer, Heidelberg (Aug 2015)
25. Keller, M., Orsini, E., Scholl, P.: MASCOT: Faster malicious arithmetic secure computation with oblivious transfer. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 16. pp. 830–842. ACM Press (Oct 2016)
26. Kiayias, A., Yung, M.: Cryptographic hardness based on the decoding of reed-solomon codes. *IEEE Trans. Information Theory* 54(6), 2752–2769 (2008)
27. Kilian, J.: Founding cryptography on oblivious transfer. In: 20th ACM STOC. pp. 20–31. ACM Press (May 1988)
28. Lindell, Y., Pinkas, B.: Privacy preserving data mining. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 36–54. Springer, Heidelberg (Aug 2000)
29. Naor, M., Pinkas, B.: Oblivious transfer and polynomial evaluation. In: 31st ACM STOC. pp. 245–254. ACM Press (May 1999)
30. Naor, M., Pinkas, B.: Oblivious polynomial evaluation. *SIAM J. Comput.* 35(5), 1254–1281 (2006)
31. Rabin, M.O.: How to exchange secrets with oblivious transfer. Technical Report TR-81, Aiken Computation Lab, Harvard University (1981)
32. Shankar, B., Srinathan, K., Rangan, C.P.: Alternative protocols for generalized oblivious transfer. In: ICDCN 2008. Lecture Notes in Computer Science, vol. 4904, pp. 304–309 (2008)
33. Tonicelli, R., Nascimento, A.C.A., Dowsley, R., Müller-Quade, J., Imai, H., Hanaoka, G., Otsuka, A.: Information-theoretically secure oblivious polynomial evaluation in the commodity-based model. *International Journal of Information Security* 14(1), 73–84 (2015), <http://dx.doi.org/10.1007/s10207-014-0247-8>
34. Zhu, H., Bao, F.: Augmented oblivious polynomial evaluation protocol and its applications. In: di Vimercati, S.D.C., Syverson, P.F., Gollmann, D. (eds.) ESORICS 2005. LNCS, vol. 3679, pp. 222–230. Springer, Heidelberg (Sep 2005)