

There Is Always an Exception: Controlling Partial Information Leakage in Secure Computation^{*}

Máté Horváth, Levente Buttyán, Gábor Székely, and Dóra Neubrandt

Laboratory of Cryptography and Systems Security (CrySyS)
Department of Networked Systems and Services
Budapest University of Technology and Economics
{mhorvath, buttyan, gszekely, dneubrandt}@crysys.hu

Abstract. Private Function Evaluation (PFE) enables two parties to jointly execute a computation such that one of them provides the input while the other chooses the function to compute. According to the traditional security requirements, a PFE protocol should leak no more information, neither about the function nor the input, than what is revealed by the output of the computation. Existing PFE protocols inherently restrict the scope of computable functions to a certain function class with given output size, thus ruling out the direct evaluation of such problematic functions as the identity map, which would entirely undermine the input privacy requirement. We observe that when not only the input x is confidential but certain partial information $g(x)$ of it as well, standard PFE fails to provide meaningful input privacy if g and the function f to be computed fall into the same function class.

Our work investigates the question whether it is possible to achieve a reasonable level of input and function privacy simultaneously even in the above cases. We propose the notion of Controlled PFE (CPFE) with different flavours of security and answer the question affirmatively by showing simple, generic realizations of the new notions. Our main construction, based on functional encryption (FE), also enjoys strong reusability properties enabling, e.g. fast computation of the same function on different inputs. To demonstrate the applicability of our approach, we show a concrete instantiation of the FE-based protocol for inner product computation that enables secure statistical analysis (and more) under the standard Decisional Diffie–Hellman assumption.

Keywords: Cryptographic Protocols · Private Function Evaluation · Functional Encryption · Oblivious Transfer · Secure Data Markets

^{*} This work was partially performed in the frames of the projects no. FIEK_16-1-2016-0007 and no. 2017-1.3.1-VKE-2017-00042, implemented with the support provided from the National Research, Development and Innovation Fund of Hungary, financed under the FIEK_16 and the 2017-1.3. funding schemes, and it was also partially supported by the National Research, Development and Innovation Office NKFIH, under grant contract no. 116675 (K). The first author was also supported by the Sándor Csibi Grant.

1 Introduction

Secure two-party computation (2PC) a.k.a. secure function evaluation (SFE) protocols enable two parties, Alice and Bob, to compute a function of their choice on their private inputs without disclosing their secrets to each other or anyone else (see Fig. 1a). In real life, however, the participants not necessarily have interchangeable roles. We call private function evaluation (PFE) a protocol if one party can alone choose the function to evaluate, while the other provides the input to it (see Fig. 1b) while both of them intends to hide their contribution. PFE can be realized by invoking 2PC after the function was turned into data. A universal function [25] is a “programmable function” that can implement *any* computation up to a given complexity. It takes two inputs, the description of the function to be computed and the input to it. By evaluating a public universal function using 2PC, all feasibility results extend from 2PC to PFE. Improving efficiency turns out to be more challenging. Indeed, universal functions cause significant – for complex computations even prohibitive – overhead, and the elimination of this limitation was the primary focus of PFE research [20,18].

In this work, we initiate the study of a security issue that – to the best of our knowledge – received no attention earlier. More concretely, we focus on the opportunities of the input provider to control the information leakage of her input. As PFE guarantees Bob that his function is hidden from Alice, he can learn some information about the input of Alice such that it remains hidden what was exactly revealed. Disclosing the entire input by evaluating the identity function is typically ruled out by the restriction that the computable function class has shorter output length than input length. At the same time, the following question arises: is it really possible to determine the computable function class so that no function is included which could reveal sensitive information about the input? We argue that most often exceptions occur in every function class, so measures are required to also protect such partial information besides the protection of the input as a whole. As intentional partial information recovery does not cause anomalies when only the function provider, Bob receives the function’s output, later on we consider this scenario.

For a simple and illustrative example, let us recall one of the most popular motivating applications for PFE. In privacy-preserving credit checking [22, §7], Alice feeds her private data to a Boolean function of her bank (or another service provider) that decides whether she is eligible for credit or not. Using PFE for such computation allows Alice to keep her data secret and the bank to hide its crediting policy. Notice that the function provider can extract *any binary information* about the input and use it, e.g. to discriminate clients. The leaked partial information can be, e.g. gender or the actual value of any indicator variable about the data that should not be necessary to reveal for credit checking. Our goal is to enable Alice to rule out the leakage of specific sensitive information in PFE without exposing what partial information she wants to hide.




Fig. 1: Comparison of the *ideal functionality* of different concepts for secure function evaluation, realized with the help of a trusted third party (TTP). The key difference lies in which information Alice and Bob can or cannot have access to.

1.1 Our Contributions

Our contributions can be summarized as follows.

- We initiate the study of partial information protection in the context of private function evaluation.
- To take the first step, we put forward the notion of Controlled PFE (CPFE) and formally define its security (see Fig. 1c for its ideal functionality). We also devise a relaxed definition, called rCPFE (see Fig. 1d) that guarantees weaker (but still reasonable) k -anonymity style function privacy leading to a trade-off between security and efficiency.
- Then we show conceptually simple, generic realizations of both CPFE and rCPFE. In the latter case, we utilize the modified function privacy guarantee (through using functional encryption) to enable the reusability of the protocol messages in case of multiple function evaluations. As a result, in our rCPFE when evaluating the same function(s) on multiple, say d inputs, the communication and online computation overhead only increases with an additive factor proportional to d instead of a multiplicative factor as in ordinary PFE.
- To demonstrate the practicality of the rCPFE approach, we instantiate our generic protocol for the inner product functionality enabling secure statistical analysis in a controlled manner under the standard Decisional Diffie–Hellman (DDH) assumption. Our proof of concept implementation shows that the reusability property indeed results in a significant performance improvement over the state of the art secure inner product evaluation method [9].

1.2 Applications

We believe that in most PFE applications, the evaluated function class also permits the leakage of potentially sensitive partial information about the input as our above example demonstrates this even for very restricted Boolean functions. To motivate our inner product rCPFE, we mention two of its possible application scenarios.

Logistic Regression Evaluation. The linear part of logistic regression computation is an inner product of the input and weight vectors. Our inner product rCPFE can help to rule out weight vectors that are unlikely to belong to a model but are base vectors that could reveal a sensitive input vector element.

Location Privacy. Let us assume that a “data broker” (DB) periodically collects location-based information in vector form, where vector elements correspond to information related to specific positions. Such data can be important for service providers (SP), offering location-based services, without the proper infrastructure to collect the necessary data. During their interaction that can be an inner product computation,¹ the SP should hide the location of its users, while the DB may want to protect the exact information in specific locations or to adjust higher price if specific measurements are used. These can be achieved by having control over the possible queries of SP.

1.3 Related Work

Some PFE variants share ideas with our concepts. Semi-private function evaluation (semi-PFE) [22,17] for instance, also relaxes the function privacy requirement of PFE by revealing the topology of the function being evaluated. While this relaxation also leads to a useful trade-off between function privacy and efficiency, unfortunately, the available extra information about the function does not necessarily allow Alice to rule out the evaluation of functions that are against her interest.

Selective private function evaluation (SPFE) [7] deals with a problem that is orthogonal to the one considered in this paper. Namely, SPFE also aims to conceal information that is leaked in PFE. However, instead of protecting Alice (the data owner), it intends to increase the security of Bob by hiding from Alice the location of the function’s input in her database via using private information retrieval (PIR).

Leaving the field of PFE and comparing our work to related problems in secure computation, we find that hiding the computed function raises similar issues in other contexts. [4] put forth the notion of verifiable obfuscation that is motivated by the natural fear for executing unknown programs. The goal here is similar than in our setting: some assurance is required that the hidden functionality cannot be arbitrary. However, the fundamental difference between our CPFE and the verifiable obfuscation and verifiable FE of [4] is that while the latter ones enforce correctness when an obfuscator or authority may be dishonest, CPFE tries to disable semi-honest parties to evaluate specific functions (i.e. to handle exceptions in PFE).

Our rCPFE is built upon functional encryption (FE) in a black-box manner. This generalization of traditional encryption was first formalized by [6]. While

¹ E.g. multiplying the data vector with a position vector (that is non-zero in all positions representing locations close to the user – possibly containing weights depending on the distance – and zero otherwise) can give useful information.

The functionality is parametrized by two integers $k < n$, and two parties: a sender \mathcal{S} and a receiver \mathcal{R} .

FUNCTIONALITY:

On input m_1, \dots, m_n messages from \mathcal{S} and an index set $\{i_1, \dots, i_k\} \subset [n]$ from \mathcal{R}

- \mathcal{S} obtains no output,
- \mathcal{R} receives m_{i_1}, \dots, m_{i_k} but nothing else.

Fig. 2: Ideal functionality $\mathcal{F}_{OT_k^n}$ of k out of n OT.

general-purpose FE candidates [12,13] currently rely on untested assumptions like the existence of indistinguishability obfuscation or multilinear maps, our application does not require such heavy hammers of cryptography (see details in §2.2). In the context of FE, [21] raised the question of controllability of function evaluation. The essential difference, compared to our goals, is that they want to limit repeated evaluations of the *same* function² that they solve with the involvement of a third party.

Finally, we sum up the state of the art of private inner product evaluation. The provably secure solutions are built either on partially homomorphic encryption schemes [14,10] or 2PC protocols [9] but public-key inner product FE [1] is also capable of the same task. At the same time, several ad-hoc protocols achieve better performance in exchange for some information leakage (see, e.g. [26] and the references therein), but these constructions lack any formal security argument.

2 Preliminaries

In this section, we briefly summarize the relevant background for the rest of the paper. We will always assume that the participants of the considered protocols are semi-honest, i.e. while following the protocol honestly, they try to recover as much information from the interactions as they can. We also use the OT-hybrid model that assumes that the parties have access to an ideal process that securely realizes oblivious transfer, which we discuss in more detail in §2.1.

2.1 Oblivious Transfer

Oblivious transfer (OT) is one of the most fundamental primitives in cryptography and a cornerstone of secure computation. It enables transferring data between two parties, the sender (\mathcal{S}) and the receiver (\mathcal{R} , a.k.a. chooser), in a way that protects both of them. \mathcal{S} can be sure that \mathcal{R} only obtains a subset of the sent messages, while \mathcal{R} is assured that \mathcal{S} does not know which messages

² In FE schemes, the control over the computable functions is in the hand of the master secret key holder, so this is not an issue unlike in PFE.

he selected to reveal. In Fig. 2 the ideal functionality of k out of n OT [8] is represented that we are also going to rely on.

While being a public-key primitive, so-called OT-extension protocols enable rather efficient OT evaluation. To do so, the participants first pre-compute a limited number of “base-OTs” with certain inputs that are independent of their real inputs. Then using the obtained values, they can evaluate a much larger number of OTs by executing more efficient symmetric-key operations only. This kind of efficiency improvement automatically applies to our protocols after substituting plain OT, with OT-extension with the same functionality [19,23].

2.2 Functional Encryption

As we already introduced, FE is a generalized encryption scheme that enables certain computations on hidden data for authorized parties. Both public- and secret-key variants are known, but here we limit ourselves to the secret-key setting that suffices for our purposes. An sk-FE scheme consists of the following four algorithms.

- FE.Setup(λ) \rightarrow ($\text{msk}_{\text{FE}}, \text{pp}_{\text{FE}}$) Upon receiving a security parameter λ it produces the public system parameters pp_{FE} and the master secret key msk_{FE} .
- FE.Enc($\text{msk}_{\text{FE}}, x$) \rightarrow ct The encryption algorithm takes the master secret key msk_{FE} and a message x and outputs a ciphertext ct.
- FE.KeyGen($\text{msk}_{\text{FE}}, f$) \rightarrow fsk_f The key generation algorithm can be used to generate a functional secret key fsk_f for a function f with the help of the msk_{FE} .
- FE.Dec(ct, fsk_f) \rightarrow y Having a functional secret key fsk_f (for function f) and a ciphertext ct (corresponding to x), the decryption outputs the value y .

The correctness of FE requires that if fsk_f and ct were indeed generated with the corresponding algorithms using inputs f and x respectively, then $y = f(x)$ must hold. Regarding security, in this work we are going to use the non-adaptive simulation-based security definition of FE [15], which we recall in Appendix A. We note that while the SIM security of FE is impossible to realize in general [6], for several restricted – yet important – cases it is still achievable, e.g. when the number of functional keys are a priori bounded [15], or when the computable function class is restricted [2]. As our applications also use these restrictions, known FE impossibility results do not affect the way we use FE.

3 General Approaches for Securing Partial Input Information in PFE

In this part, we introduce the notion of controlled PFE and in §3.1 formally define its security in different flavours. Next, in §3.2–3.3, we propose two general protocols satisfying these security requirements.

PARAMETERS: participants P_1, P_2 , a class $\mathcal{F} = \{f : \mathcal{X} \rightarrow \mathcal{Y}\}$ of deterministic functions [and an integer $\kappa > k$]

FUNCTIONALITY:

- On inputs $x_1, \dots, x_d \in \mathcal{X}$ and $\mathcal{F}_A \subset \mathcal{F}$ from P_1 ; and $\mathcal{F}_B = \{f_1, \dots, f_k\} \subset \mathcal{F}$ from P_2
- P_1 receives no output, [or P_1 receives \mathcal{F}_R s.t. $\mathcal{F}_B \subset \mathcal{F}_R \subset \mathcal{F}$ and $|\mathcal{F}_R| = \kappa$]
 - P_2 obtains $\{y'_{i,j} = f'_j(x_i)\}_{i \in [d], j \in [k]} \subset \mathcal{Y} \cup \{\perp\}$ for

$$f'_j(x_i) = \begin{cases} f_j(x_i) & \text{if } f_j \notin \mathcal{F}_A \\ \perp & \text{otherwise.} \end{cases}$$

Fig. 3: Ideal functionalities for $\mathcal{F}_{\text{CPFE}}$ and $\mathcal{F}_{\text{ICPFE}}$ (see the extensions in brackets) formulated generally for multiple inputs and multiple functions.

3.1 Definitional Framework

Our first security definition for controlled PFE captures the intuitive goal of extending the PFE functionality with a blind function verification step by P_1 to prevent unwanted information leakage. See the corresponding ideal functionality $\mathcal{F}_{\text{CPFE}}$ in Fig. 3 that we call controlled PFE, and the security definition below. For the ease of exposition, later on we denote the inputs of the participants as $\text{inp} = (\{x_i\}_{i \in [d]}, \mathcal{F}_A, \{f_j\}_{j \in [k]})$ with the corresponding parameters.

Definition 1 (SIM security of CPFE wrt. semi-honest adversaries). Let Π denote a Controlled PFE (CPFE) protocol for a function class \mathcal{F} with functionality $\mathcal{F}_{\text{CPFE}}$ (according to Fig. 3). We say that Π achieves SIM security against semi-honest adversaries, if the following criteria hold.

- Correctness: the output computed by Π is the required output, i.e.

$$\Pr[\text{output}^\Pi(1^\lambda, \text{inp}) \neq \mathcal{F}_{\text{CPFE}}(\text{inp})] \leq \text{negl}(\lambda).$$

- Function Privacy: there exist a PPT simulator \mathcal{S}_{P_1} , s.t.

$$\{\mathcal{S}_{P_1}(1^\lambda, \{x_i\}_{i \in [d]}, \mathcal{F}_A)\}_{\lambda, x_i, \mathcal{F}_A} \stackrel{c}{\equiv} \{\text{view}_{P_1}^\Pi(1^\lambda, \text{inp})\}_{\lambda, x_i, f_j, \mathcal{F}_A}.$$

- Data Privacy: there exist a PPT simulator \mathcal{S}_{P_2} , s.t.

$$\{\mathcal{S}_{P_2}(1^\lambda, \{f_j\}_{j \in [k]}, \{y'_{i,j}\}_{i \in [d], j \in [k]})\}_{\lambda, f_j} \stackrel{c}{\equiv} \{\text{view}_{P_2}^\Pi(1^\lambda, \text{inp})\}_{\lambda, x_i, f_j, \mathcal{F}_A}$$

where $\text{inp} = (\{x_i\}_{i \in [d]}, \mathcal{F}_A, \{f_j\}_{j \in [k]})$, $f_j \in \mathcal{F}$, $\mathcal{F}_A \subset \mathcal{F}$, $x_i \in \mathcal{X}$, $y'_{i,j} \in \mathcal{Y} \cup \{\perp\}$, and $\lambda \in \mathbb{N}$.

We also propose a relaxation of Def. 1, which on the one hand gives up perfect function privacy but on the other, allows us to construct efficient protocols while still maintaining a k -anonymity style guarantee for function privacy. As SIM security alone cannot measure how much information is leaked by a set of functions, we formulate an additional requirement to precisely characterise function privacy.

Definition 2 (SIM security of relaxed CPFE wrt. semi-honest adversaries). Let Π denote a relaxed CPFE (rCPFE) protocol for a function class \mathcal{F} with functionality $\mathcal{F}_{\text{rCPFE}}$ (according to Fig. 3). We say that Π achieves SIM security against semi-honest adversaries, if the following criteria hold.

- Correctness: the output computed by Π is the required output, i.e.

$$\Pr[\text{output}^\Pi(\lambda, \kappa, \text{inp}) \neq \mathcal{F}_{\text{rCPFE}}(\kappa, \text{inp})] \leq \text{negl}(\lambda).$$

- Function Privacy: is defined in two flavours:

- κ -relaxed function privacy holds, if $\exists \mathcal{S}_{P_1}$, a PPT simulator, s.t.

$$\{\mathcal{S}_{P_1}(1^\lambda, \kappa, \{x_i\}_{i \in [d]}, \mathcal{F}_A)\}_{\lambda, x_i, \mathcal{F}_A} \stackrel{c}{\equiv} \{\text{view}_{P_1}^\Pi(1^\lambda, \kappa, \text{inp})\}_{\lambda, x_i, f_j, \mathcal{F}_A}.$$

- Strong κ -relaxed function privacy holds if besides the existence of the above \mathcal{S}_{P_1} , it also holds that for any PPT \mathcal{A} :

$$\left| \Pr[\mathcal{A}(\text{aux}, \mathcal{F}_R) = f \in \mathcal{F}_B] - \frac{k}{\kappa} \right| \leq \text{negl}(\lambda)$$

where $\text{aux} \in \{0, 1\}^*$ denotes some a priori known auxiliary information about \mathcal{F}_B .

- Data Privacy: there exist a PPT simulator \mathcal{S}_{P_2} , s.t.

$$\{\mathcal{S}_{P_2}(\lambda, \kappa, \{f_j\}_{j \in [k]}, y'_{i,j})_{\lambda, f_j} \stackrel{c}{\equiv} \{\text{view}_{P_2}^\Pi(\lambda, \kappa, \text{inp})\}_{\lambda, x_i, f_j, \mathcal{F}_A}$$

where $\text{inp} = (\{x_i\}_{i \in [d]}, \mathcal{F}_A, \{f_j\}_{j \in [k]})$, $f_j \in \mathcal{F}$, $\mathcal{F}_A \subset \mathcal{F}$, $x_i \in \mathcal{X}$, $y'_{i,j} \in \mathcal{Y} \cup \{\perp\}$, and $\lambda, \kappa \in \mathbb{N}$.

3.2 Universal Circuit-based CPFE

The natural approach for realizing CPFE comes from the traditional way of combining universal circuits and SFE to obtain PFE. Fig. 4 shows how the same idea with conditional evaluation leads to CPFE in the single input, single function setting. The following theorem is a straightforward consequence of the security of SFE.

Theorem 1. *The CPFE protocol of Fig. 4 is secure according to Def. 1, if the used SFE protocol is SIM secure in the semi-honest model.*

The main drawback of this approach is that when extending the protocol to handle multiple inputs or functions, its complexity will multiplicatively depend on the number of inputs or functions because of the single-use nature of 2PC.

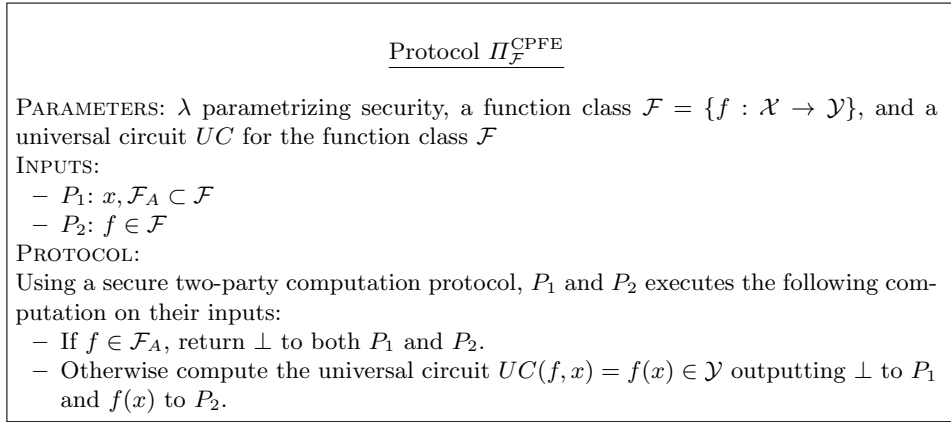


Fig. 4: General 2PC-based CPFE

3.3 Reusable Relaxed CPFE from FE

We observe that the notion of rCPFE not only allows the input provider to verify the functions to be evaluated but also opens the door for making parts of the protocol messages reusable multiple times, thus leading to significant efficiency improvements.

A naive first attempt to realize rCPFE is to execute the computation on the side of P_1 . Upon receiving a κ function descriptions (including both the intended and dummy functions) P_1 can easily verify the request and evaluate the allowed ones on her input. The results then can be shared with P_2 , using an OT scheme achieving both the required data and function privacy level. Unfortunately, the κ function evaluations lead to scalability issues. The subsequent natural idea is to shift the task of function evaluation to P_2 , to eliminate the unnecessary computations and to hide the output from P_1 entirely. Since at this point P_1 has both the inputs and the functions to evaluate, the task resembles secure outsourcing of computation where function evaluation must be under the strict control of P_1 . These observations lead us to the usage of FE and the protocol in Fig. 5 in which both ciphertext and functional keys can be reused in multiple computations. When instantiated with the FE scheme of [15], $\Pi_{\mathcal{F}}^{\text{rCPFE}}$ can be used for all polynomial sized functions in theory (in practice verifying the circuits would be a bottleneck).

Theorem 2. *The protocol of Fig. 5 is SIM secure according to Def. 2 achieving κ -relaxed function privacy for k function queries by P_2 , if the underlying FE scheme is k -query non-adaptive SIM secure (k -NA-SIM) for a single message and the used OT protocol is SIM secure against semi-honest adversaries.*

The proof of the theorem is postponed to Appendix B.

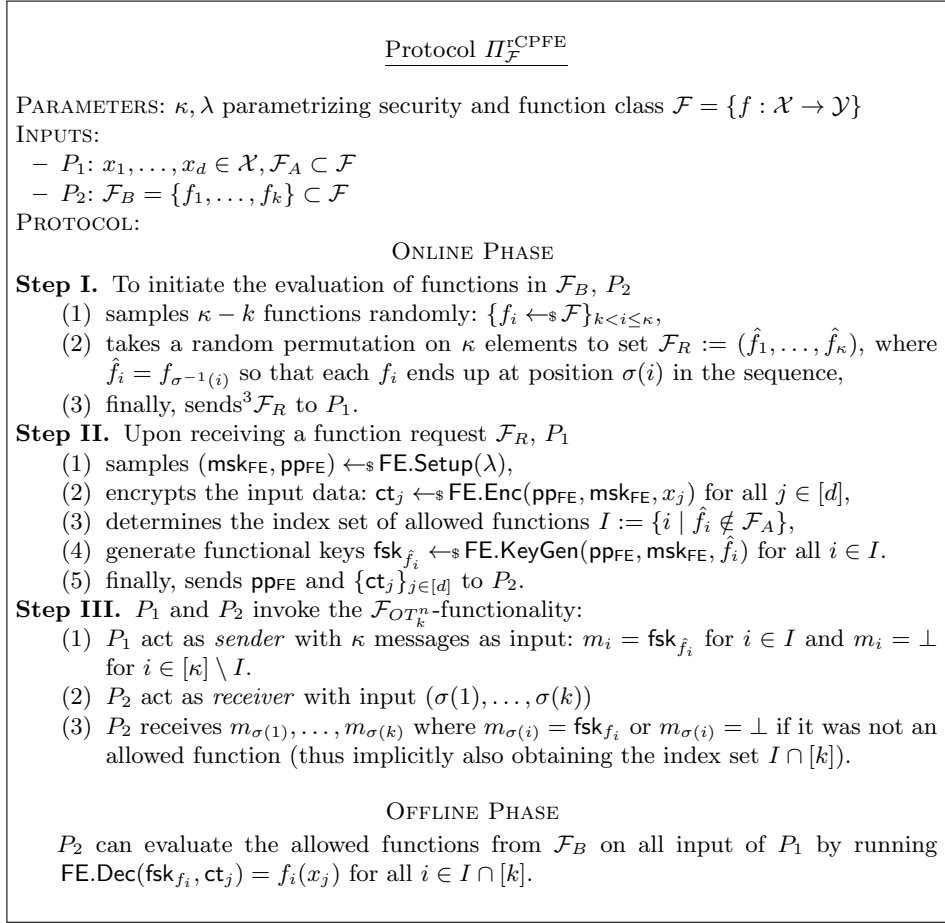


Fig. 5: General rCPFE construction.

Corollary 1. *The protocol of Fig. 5 also achieves strong κ -relaxed function privacy if in (1) of Step I., all f_i are sampled from the same distribution as the elements of \mathcal{F}_B and $\text{aux} = \perp$.*

4 Concrete Instantiation for Inner Products

To demonstrate the practicality of our approach, we instantiate our generic rCPFE protocol (Fig. 5) using the k -NA-SIM secure FE scheme of [2] for the

³ Depending on \mathcal{F} and the sampling of the dummy functions, communication cost of transferring the function descriptions can be reduced. In [16] we describe such optimizations for the inner product function class.

inner product functionality and the semi-honest 1 out of κ OT protocol of [24]. Theorem 2 and the assumptions of [2,24] directly imply the following theorem.

Theorem 3. *There is a SIM secure rCPFE protocol (according to Def. 2) for inner product computation, achieving κ -relaxed function privacy, if the DDH assumption holds.*

Corollary 2. *The inner product rCPFE protocol derived from $\Pi_{\mathcal{F}}^{\text{rCPFE}}$ (on Fig. 5) also achieves strong κ -relaxed function privacy (as defined in Def. 2) if $\mathbf{aux} = \perp$ and the dummy function vectors are chosen from the same distribution as the real ones.*

For the detailed description of the inner product rCPFE (or IP-rCPFE for short) we refer to the full version of this paper [16].

4.1 Performance and Possible Optimizations

For our IP-CPFE protocol, we prepared a proof of concept implementation using the Charm framework [3]. To evaluate its performance in two scenarios, we compared its running times and communication costs with that of the state of the art secure arithmetic inner product computation method of the ABY framework [9]. For our experiments we used a commodity laptop with a 2.60GHz Intel[®] Core[™] i7-6700HQ CPU and 4GB of RAM.

Simulating regression model evaluation. In the first use-case, we do not assume that the vectors have a special structure. The vectors to be multiplied can correspond to data and weight vectors of a binary regression model, in which case it is likely that the same model (weight vector) is evaluated over multiple inputs. Fig. 6a and 6d depict running times and overall communication costs respectively depending on the number of inputs to the same model. Fig. 6c and 6f show the cost of the dummy queries. In the same setting, our experiments show that without optimizations⁴ IP-rCPFE reaches the running time of ABY for $\kappa \approx 6200$. For this scenario, we also propose a method (denoted as rCPFE opt) to pre-compute the dummy function queries of Step I. thus reducing both the online communication and computation costs. The key insight of this is that sending a value together with dummy values is essentially the same as hiding the value with a one time pad (OTP) and attaching the OTP key together with dummy keys. The gain comes from the fact that the OTP keys can be computed and sent beforehand, moreover it is enough to transmit the used seeds for a pseudo-random generator instead of the entire keys (see details in [16]). Security is not affected as long as $\mathbf{aux} = \perp$.

⁴ We note that while our implementation is only a proof of concept without any code level optimization, ABY has a very efficient and parallelizable implementation.

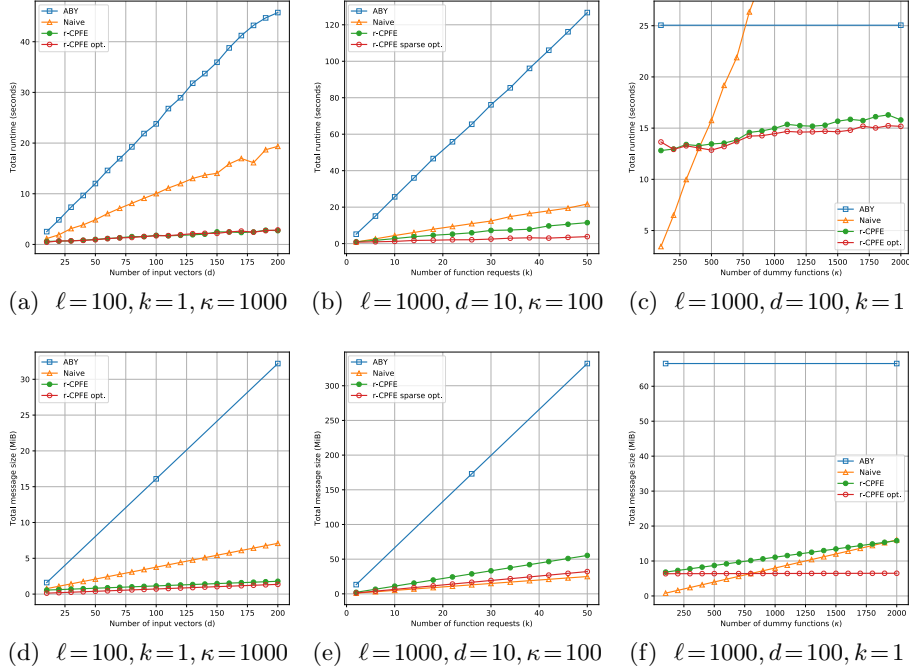


Fig. 6: Comparisons of the overall running times (6a–6c) and communication costs (6d–6f) of our rCPFE protocols with the ABY framework [9] and the naive OT-based approach for inner product computation (ℓ denotes vector dimension, d and k are the number of input and “function” vectors, while κ is the number of dummy vectors).

Sparse vector products for location privacy. The location privacy scenario of §1.2 implies the usage of sparse query vectors. Fig. 6b and 6e show how the number of queries (k) affects running time and message sizes respectively, when roughly 5% of the vector elements are non-zero. We note that as queries are related to real-time user requests, batching these requests, as done in Step I. of the protocol, can be unrealistic when data vectors are not changing in real time but, e.g. periodically. Because of this, in our implementation, we allowed P_2 to repeat Step I. for a single function and P_1 to answer the queries independently of encrypting the data.⁵ While sparsity disables the above optimization, after masking the places of non-zero elements, the above idea can be extended for sparse vectors as long as other structural properties are not known about the vector in form of auxiliary information. For more details on the optimized variants, we refer to [16].

⁵ It means that (3)–(4) of Step II., and Step III. are repeated until the input data changes at the end of the period.

5 Conclusion and Open Directions

In this work, we attempted to draw attention to the problem of possibly sensitive partial information leakage in the context of private function evaluation. We proposed a definitional framework for protocols that aim to prevent such leakage and showed both generic and concrete protocols to solve the problem. The main advantage of our FE-based protocol is that it turns the privacy sacrifice required by controllability into performance improvement whenever more function evaluations are necessary.

Our work also leaves open several problems for future work. For instance, it would be important to investigate the effects of having different types of auxiliary information about the evaluated functions. Transmission and verification of dummy functions can be serious bottlenecks in our rCPFE in case of complex functions, making further efficiency improvements desirable. A first step towards this could be to find a way for restricting the set of forbidden functions – as most often very simple functions are the only undesired ones. Finally, looking for different trade-offs between function privacy and efficiency can also be interesting direction for future work.

References

1. Abdalla, M., Bourse, F., Caro, A.D., Pointcheval, D.: Simple functional encryption schemes for inner products. In: Katz, J. (ed.) *Proceedings of Public-Key Cryptography - PKC 2015*. LNCS, vol. 9020, pp. 733–751. Springer (2015)
2. Agrawal, S., Libert, B., Stehlé, D.: Fully secure functional encryption for inner products, from standard assumptions. In: Robshaw, M., Katz, J. (eds.) *Advances in Cryptology - CRYPTO 2016, Proceedings, Part III*. Lecture Notes in Computer Science, vol. 9816, pp. 333–362. Springer (2016). https://doi.org/10.1007/978-3-662-53015-3_12
3. Akinyele, J.A., Garman, C., Miers, I., Pagano, M.W., Rushanan, M., Green, M., Rubin, A.D.: Charm: a framework for rapidly prototyping cryptosystems. *J. Cryptographic Engineering* **3**(2), 111–128 (2013). <https://doi.org/10.1007/s13389-013-0057-3>
4. Badrinarayanan, S., Goyal, V., Jain, A., Sahai, A.: Verifiable functional encryption. In: Cheon, J.H., Takagi, T. (eds.) *Advances in Cryptology - ASIACRYPT 2016, Proceedings, Part II*. LNCS, vol. 10032, pp. 557–587 (2016). https://doi.org/10.1007/978-3-662-53890-6_19
5. Bishop, A., Jain, A., Kowalczyk, L.: Function-hiding inner product encryption. In: *Advances in Cryptology - ASIACRYPT 2015, Proceedings, Part I*. pp. 470–491 (2015), http://dx.doi.org/10.1007/978-3-662-48797-6_20
6. Boneh, D., Sahai, A., Waters, B.: Functional encryption: Definitions and challenges. In: Ishai, Y. (ed.) *Theory of Cryptography - TCC 2011*. Proceedings. LNCS, vol. 6597, pp. 253–273. Springer (2011)
7. Canetti, R., Ishai, Y., Kumar, R., Reiter, M.K., Rubinfeld, R., Wright, R.N.: Selective private function evaluation with applications to private statistics. In: Kshemkalyani, A.D., Shavit, N. (eds.) *Proceedings of the Twentieth Annual ACM Symposium on Principles of Distributed Computing, PODC 2001*,. pp. 293–304. ACM (2001). <https://doi.org/10.1145/383962.384047>

8. Chu, C., Tzeng, W.: Efficient k -out-of- n oblivious transfer schemes with adaptive and non-adaptive queries. In: Vaudenay, S. (ed.) Public Key Cryptography - PKC 2005, Proceedings. LNCS, vol. 3386, pp. 172–183. Springer (2005). https://doi.org/10.1007/978-3-540-30580-4_12
9. Demmler, D., Schneider, T., Zohner, M.: ABY - A framework for efficient mixed-protocol secure two-party computation. In: 22nd Annual Network and Distributed System Security Symposium, NDSS 2015. The Internet Society (2015), <https://www.ndss-symposium.org/ndss2015/aby---framework-efficient-mixed-protocol-secure-two-party-computation>
10. Dong, C., Chen, L.: A fast secure dot product protocol with application to privacy preserving association rule mining. In: Tseng, V.S. et al. (eds.) Advances in Knowledge Discovery and Data Mining - 18th Pacific-Asia Conference, PAKDD 2014, Proceedings, Part I. LNCS, vol. 8443, pp. 606–617. Springer (2014). https://doi.org/10.1007/978-3-319-06608-0_50
11. Fouque, P., Joux, A., Tibouchi, M.: Injective encodings to elliptic curves. In: Boyd, C., Simpson, L. (eds.) Information Security and Privacy - 18th Australasian Conference, ACISP 2013. LNCS, vol. 7959, pp. 203–218. Springer (2013). https://doi.org/10.1007/978-3-642-39059-3_14
12. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, pp. 40–49. IEEE Computer Society (2013). <https://doi.org/10.1109/FOCS.2013.13>
13. Garg, S., Gentry, C., Halevi, S., Zhandry, M.: Functional encryption without obfuscation. In: Kushilevitz, E., Malkin, T. (eds.) Theory of Cryptography - 13th International Conference, TCC 2016-A, Proceedings, Part II. LNCS, vol. 9563, pp. 480–511. Springer (2016). https://doi.org/10.1007/978-3-662-49099-0_18
14. Goethals, B., Laur, S., Lipmaa, H., Mielikäinen, T.: On private scalar product computation for privacy-preserving data mining. In: Park, C., Chee, S. (eds.) Information Security and Cryptology - ICISC 2004, Revised Selected Papers. LNCS, vol. 3506, pp. 104–120. Springer (2004). https://doi.org/10.1007/11496618_9
15. Gorbunov, S., Vaikuntanathan, V., Wee, H.: Functional encryption with bounded collusions via multi-party computation. In: Safavi-Naini, R., Canetti, R. (eds.) Advances in Cryptology - CRYPTO 2012, Proceedings, LNCS, vol. 7417, pp. 162–179. Springer (2012). https://doi.org/10.1007/978-3-642-32009-5_11
16. Horváth, M., Buttyán, L., Székely, G., Neubrandt, D.: There is always an exception: Controlling partial information leakage in secure computation (full version). Cryptology ePrint Archive, Report 2019/ (2019), <https://eprint.iacr.org/2019/>
17. Kennedy, W.S., Kolesnikov, V., Wilfong, G.T.: Overlaying conditional circuit clauses for secure computation. In: Takagi, T., Peyrin, T. (eds.) Advances in Cryptology - ASIACRYPT 2017, Proceedings, Part II. LNCS, vol. 10625, pp. 499–528. Springer (2017). https://doi.org/10.1007/978-3-319-70697-9_18
18. Kiss, Á., Schneider, T.: Valiant’s universal circuit is practical. In: Fischlin, M., Coron, J. (eds.) Advances in Cryptology - EUROCRYPT 2016, Proceedings, Part I. LNCS, vol. 9665, pp. 699–728. Springer (2016). https://doi.org/10.1007/978-3-662-49890-3_27
19. Kolesnikov, V., Kumaresan, R., Rosulek, M., Trieu, N.: Efficient batched oblivious PRF with applications to private set intersection. In: Weippl, E.R. et al. (eds.) Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 818–829. ACM (2016). <https://doi.org/10.1145/2976749.2978381>

20. Kolesnikov, V., Schneider, T.: A practical universal circuit construction and secure evaluation of private functions. In: Tsudik, G. (ed.) *Financial Cryptography and Data Security*, 12th International Conference, FC 2008, Revised Selected Papers. LNCS, vol. 5143, pp. 83–97. Springer (2008). https://doi.org/10.1007/978-3-540-85230-8_7
21. Naveed, M., Agrawal, S., Prabhakaran, M., Wang, X., Ayday, E., Hubaux, J., Gunter, C.A.: Controlled functional encryption. In: Ahn, G., Yung, M., Li, N. (eds.) *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014. pp. 1280–1291. ACM (2014). <https://doi.org/10.1145/2660267.2660291>
22. Paus, A., Sadeghi, A., Schneider, T.: Practical secure evaluation of semi-private functions. In: Abdalla, M., et al. (eds.) *Applied Cryptography and Network Security*, 7th International Conference, ACNS 2009, Proceedings. LNCS, vol. 5536, pp. 89–106 (2009). https://doi.org/10.1007/978-3-642-01957-9_6
23. Rindal, P., Rosulek, M.: Improved private set intersection against malicious adversaries. In: Coron, J., Nielsen, J.B. (eds.) *Advances in Cryptology - EUROCRYPT 2017, Proceedings, Part I*. LNCS, vol. 10210, pp. 235–259 (2017). https://doi.org/10.1007/978-3-319-56620-7_9
24. Tzeng, W.: Efficient 1-out-of-n oblivious transfer schemes with universally usable parameters. *IEEE Trans. Computers* **53**(2), 232–240 (2004). <https://doi.org/10.1109/TC.2004.1261831>
25. Valiant, L.G.: Universal circuits (preliminary report). In: Chandra, A.K., Wotschke, D., Friedman, E.P., Harrison, M.A. (eds.) *Proceedings of the 8th Annual ACM Symposium on Theory of Computing*, pp. 196–203. ACM (1976). <https://doi.org/10.1145/800113.803649>
26. Zhu, Y., Wang, Z., Hassan, B., Zhang, Y., Wang, J., Qian, C.: Fast secure scalar product protocol with (almost) optimal efficiency. In: Guo, S., et al. (eds.) *Collaborative Computing: Networking, Applications, and Worksharing, Proceedings*. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, vol. 163, pp. 234–242. Springer (2015). https://doi.org/10.1007/978-3-319-28910-6_21

Appendix

A Simulation Security of Functional Encryption

For completeness we recall the simulation security of FE as defined in [15].

Definition 3 (q -NA-SIM and q -AD-SIM Security of FE). *Let FE be a functional encryption scheme for a circuit family $\mathcal{C} = \{\mathcal{C}_\kappa : \mathcal{X}_\kappa \rightarrow \mathcal{Y}_\kappa\}_{\kappa \in \mathbb{N}}$. For every PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and a PPT simulator $S = (S_1, S_2)$ consider the following two experiments:*

$\text{Exp}_{FE,A}^{\text{real}}(\lambda)$	$\text{Exp}_{FE,S}^{\text{ideal}}(\lambda)$
1 : $(\text{pp}_{FE}, \text{msk}_{FE}) \leftarrow_{\$} \text{FE.Setup}(\lambda)$	1 : $(\text{pp}_{FE}, \text{msk}_{FE}) \leftarrow_{\$} \text{FE.Setup}(\lambda)$
2 : $(x, \text{st}) \leftarrow_{\$} \mathcal{A}_1^{\text{FE.KeyGen}(\text{msk}_{FE}, \cdot)}(\text{pp}_{FE})$	2 : $(x, \text{st}) \leftarrow_{\$} \mathcal{A}_1^{\text{FE.KeyGen}(\text{msk}_{FE}, \cdot)}(pp)$ – Let (C_1, \dots, C_q) be \mathcal{A}_1 's oracle queries – Let fsk_{f_i} be the oracle reply to C_i – Let $\mathcal{V} := \{y_i = C_i(x), C_i, \text{fsk}_{f_i}\}$.
3 : $\text{ct} \leftarrow_{\$} \text{FE.Enc}(\text{pp}_{FE}, x)$	3 : $(\text{ct}, \text{st}') \leftarrow_{\$} S_1(\text{pp}_{FE}, \mathcal{V}, \lambda)$
4 : $\beta \leftarrow_{\$} \mathcal{A}_2^{O(\text{msk}_{FE}, \cdot)}(\text{pp}_{FE}, \text{ct}, \text{st})$	4 : $\beta \leftarrow_{\$} \mathcal{A}_2^{O'(\text{msk}_{FE}, \text{st}', \cdot)}(\text{pp}_{FE}, \text{ct}, \text{st})$
5 : $\text{Output}(\beta, x)$	5 : $\text{Output}(\beta, x)$

We distinguish between two case of the above experiment:

1. The adaptive case, where:
 - the oracle $O(\text{msk}_{FE}, \cdot) = \text{FE.KeyGen}(\text{msk}_{FE}, \cdot)$ and
 - the oracle $O'(\text{msk}_{FE}, \text{st}', \cdot)$ is the second stage of the simulator, namely $S_2^{U_x(\cdot)}(\text{msk}_{FE}, \text{st}', \cdot)$ where $U_x(C) = C(c)$ for any $C \in \mathcal{C}_{\mathcal{K}}$.

The simulator algorithm S_2 is stateful in that after each invocation, it updates the state st' which is carried over to its next invocation. We call a simulator algorithm $S = (S_1, S_2)$ admissible if, on each input C , S_2 just makes a single query to its oracle $U_x(\cdot)$ on C itself,

The functional encryption scheme FE is then said to be q -query-simulation-secure for one message against adaptive adversaries (q -AD-SIM secure for short) if there is an admissible PPT simulator $S = (S_1, S_2)$ such that for every ppt adversary $A = (A_1, A_2)$ that makes at most q queries, the following two distributions are computationally indistinguishable:

$$\left\{ \text{Exp}_{FE,A}^{\text{real}}(\lambda) \right\}_{\mathcal{K} \in \mathbb{N}} \stackrel{c}{\approx} \left\{ \text{Exp}_{FE,S}^{\text{ideal}}(\lambda) \right\}_{\mathcal{K} \in \mathbb{N}}$$

2. The non-adaptive case, where the oracles $O(\text{msk}_{FE}, \text{st}, \cdot)$ are both the "empty oracles" that return nothing: the functional encryption scheme FE is then said to be q -query-simulation-secure for one message against non-adaptive adversaries (q -NA-SIM secure, for short) if there is a PPT simulator $S = (S_1, \perp)$ such that for every PPT adversary $A = (A_1, A_2)$ that makes at most q queries, the two distributions above are computationally indistinguishable.

As shown by [15, Theorem A.1.], in the non-adaptive setting (that we also use), q -NA-SIM security for one message is equivalent to q -NA-SIM security for many messages.

B Proof of Theorem 2

We prove Theorem 2, by showing that the protocol of Fig. 5 fulfils the requirements of Definition 2 with the assumption that the underlying FE and OT are SIM secure against semi-honest adversaries. As correctness directly follows from

the correctness of the underlying FE and OT, we turn our attention towards the security requirements. We argue input and weak relaxed function privacy by showing that the view of both parties can be simulated (without having access to the inputs of the other party) using the simulators guaranteed by the SIM security of FE and OT.

Corrupted P_1 : Weak Relaxed Function Privacy. Besides its input and output, the view of P_1 consists of the received OT messages and the function query \mathcal{F}_R . Simulation becomes trivial because of the fact that the output of P_1 also contains \mathcal{F}_R . Thus $\mathcal{S}_{P_1}((x_1, \dots, x_d), \mathcal{F}_R)$ can return \mathcal{F}_R and the output of the sender's simulator \mathcal{S}_{OT}^S guaranteed by the SIM security of OT. The simulated view is clearly indistinguishable from the real one.

Corrupted P_2 : Input privacy. The following simulator \mathcal{S}_{P_2} simulates the view of a corrupt P_2 , that consists of its input (f_1, \dots, f_k) , output $\{y'_{i,j} = f'_i(x_j)\}_{i \in [k], j \in [d]}$, the used randomness and the incoming messages. \mathcal{S}_{P_2} first determines the index set $I^* = \{i \mid \exists j : y'_{i,j} \neq \perp\} \subseteq [k]$. Next, it sets up the parameters of the ideal experiment according to Def. 3. To do so, it samples $(\text{msk}_{\text{FE}}^*, \text{pp}_{\text{FE}}^*) \leftarrow \text{FE.Setup}(\lambda)$ and then for all $i \in I^*$ generates keys $\text{fsk}_{f_i}^* \leftarrow \text{FE.KeyGen}(\text{pp}_{\text{FE}}^*, \text{msk}_{\text{FE}}^*, f_i)$. For the simulation of the FE ciphertexts (corresponding to unknown messages), we can use the FE simulator \mathcal{S}_{FE} for many messages (implied by one-message q -NA-SIM security [15]). Thus $\mathcal{S}_{FE}(\text{pp}_{\text{FE}}^*, \{y_{i,j} = f_i(x_j), f_i, \text{fsk}_{f_i}^*\}_{i \in I^*, j \in [d]}, \lambda) = (\text{ct}_1^*, \dots, \text{ct}_d^*)$ can be appended to the simulated view together with pp_{FE}^* . The incoming messages of Step III. are simulated using the OT simulator \mathcal{S}_{OT}^R for the receiver. Finally the output of $\mathcal{S}_{OT}^R(\lambda, \{\text{fsk}_{f_i}^*\}_{i \in I^*} \cup \{\perp_i\}_{i \in [k] \setminus I^*})$ is appended to the simulated view.

Now we show the indistinguishability of the real and simulated views. As the inputs and outputs are the same in both cases, we have to compare the randomness and the incoming messages. First notice that pp_{FE} and pp_{FE}^* are generated with different random choices. At the same time, these cannot be told apart as otherwise the choices were not random. The rest of the incoming messages depend on these parameters. Observe that $I^* = I \cap [k]$. The security of the used FE scheme guarantees that $(\text{ct}_1^*, \dots, \text{ct}_d^*)$ even together with functional keys $\{\text{fsk}_{f_i}^*\}_{i \in I^*}$ are indistinguishable from $(\text{ct}_1, \dots, \text{ct}_d)$ with $\{\text{fsk}_{f_i}\}_{i \in I \cap [k]}$. Finally, the security of the OT simulation guarantees that $(\text{msg}_1^{\text{OT}}, \dots, \text{msg}_\kappa^{\text{OT}})$ and $(\text{msg}_1^{\text{OT}*}, \dots, \text{msg}_\kappa^{\text{OT}*})$ are indistinguishable. This also implies that functional keys for the same functions (with respect to either pp_{FE} or pp_{FE}^*) can be obtained both from the real and simulated OT messages. In other words, FE ciphertexts and functional keys are consistent in both cases (i.e. they allow one to obtain the same decryption outputs) due to the correctness of the FE simulation, which concludes our proof. \square