

Distributed (Correlation) Samplers: How to Remove a Trusted Dealer in One Round*

Damiano Abram, Peter Scholl, and Sophia Yakoubov

Aarhus University, Aarhus, Denmark

Abstract. Structured random strings (SRSs) and correlated randomness are important for many cryptographic protocols. In settings where interaction is expensive, it is desirable to obtain such randomness in as few rounds of communication as possible; ideally, simply by exchanging one reusable round of messages which can be considered public keys.

In this paper, we describe how to generate any SRS or correlated randomness in such a single round of communication, using, among other things, indistinguishability obfuscation. We introduce what we call a *distributed sampler*, which enables n parties to sample a single public value (SRS) from any distribution. We construct a semi-malicious distributed sampler in the plain model, and use it to build a semi-malicious *public-key PCF* (Boyle *et al.*, FOCS 2020) in the plain model. A public-key PCF can be thought of as a distributed *correlation* sampler; instead of producing a public SRS, it gives each party a private random value (where the values satisfy some correlation).

We introduce a general technique called an *anti-rusher* which compiles any one-round protocol with semi-malicious security without inputs to a similar one-round protocol with active security by making use of a programmable random oracle. This gets us actively secure distributed samplers and public-key PCFs in the random oracle model.

Finally, we explore some tradeoffs. Our first PCF construction is limited to *reverse-sampleable* correlations (where the random outputs of honest parties must be simulatable given the random outputs of corrupt parties); we additionally show a different construction without this limitation, but which does not allow parties to hold secret parameters of the correlation. We also describe how to avoid the use of a random oracle at the cost of relying on sub-exponentially secure indistinguishability obfuscation.

* Supported by the Independent Research Fund Denmark (DFF) under project number 0165-00107B (C3PO), the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 803096 (SPEC), and a starting grant from Aarhus University Research Foundation.

Table of Contents

Distributed (Correlation) Samplers: How to Remove a Trusted Dealer in One Round	1
<i>Damiano Abram, Peter Scholl, and Sophia Yakoubov</i>	
1 Introduction	3
1.1 Related Work	3
1.2 Our Contributions	5
1.3 Technical Overview	7
2 Preliminaries	9
2.1 Indistinguishability Obfuscation	9
2.2 Puncturable PRFs	10
2.3 Simulation-Extractable NIZKs	10
2.4 MHE with Private Evaluation	12
2.5 Somewhere Statistically Binding Hash Functions	13
2.6 Universal Samplers	14
3 Defining Distributed Samplers	16
3.1 Security	17
4 A Construction with Semi-Malicious Security	19
5 Upgrading to Active Security	28
5.1 Defeating Rushing	28
6 Public-Key PCFs for Reverse-Samplable Correlations	32
6.1 Correlation Functions and their Properties	33
6.2 Defining Public Key PCFs	35
6.3 Public-Key PCF with Trusted Setup	37
6.4 Our Public-Key PCFs	41
7 Ideal Public Key PCFs and Distributed Universal Samplers	43
7.1 Distributed Universal Samplers	44
7.2 Building Ideal Public Key PCFs upon Distributed Universal Samplers	48
A Proof of Theorem 5.2	51
B Proof of Theorem 6.7	62
C Proof of Theorem 6.10	69

1 Introduction

Randomness is crucial for many cryptographic protocols. Participants can generate some randomness locally (e.g. by flipping coins), but the generation of other forms of randomness is more involved. For instance, a *uniform reference string* (URS) must be produced in such a way that a coalition of corrupt protocol participants — controlled by the adversary — cannot bias it too much. Even more complex is the generation of a *structured* reference string (SRS, such as an RSA modulus), which can depend on secrets that should not be known to anyone. For instance, constructions such as cryptographic accumulators [Bd94] use an RSA modulus whose factorization is known to nobody, while many succinct zero-knowledge proofs such as SNARKs [BCCT12] require a more complex form of SRS.

In contrast to common reference strings, which are public, some protocols demand *correlated randomness*, where each participant holds a secret random value, but because the values must satisfy some relationship, they cannot be generated locally by the participants. An example of correlated randomness is random oblivious transfer, where one participant has a list of random strings, and another has one of those strings as well as its index in the list. Such correlated randomness often allows cryptographic protocols to run with a more efficient online phase.

Typically, in order to set up an SRS or correlated randomness without making additional trust assumptions, the parties must run a secure multi-party computation protocol, which takes several rounds of interaction. This is the case, for instance, in “setup ceremonies” [BGG19, BGM17] that have been designed to generate trusted SNARK parameters for applications. In this paper, we explore techniques that let parties sample *any* common reference string or correlation in just *one round* of interaction.

1.1 Related Work

There are a number of lines of work that can be used to generate randomness in different ways.

Universal samplers. A universal sampler [HJK⁺16] is a kind of SRS which can be used to obviously sample from any distribution that has an efficient sampling algorithm. That is, after a one-time trusted setup to generate the universal sampler, it can be used to generate arbitrary other SRSs. Hofheinz *et al.* [HJK⁺16] show how to build universal samplers from indistinguishability obfuscation and a random oracle, while allowing an unbounded number of adaptive queries. They also show how to build weaker forms of universal sampler in the standard model, from single-key functional encryption [LZ17]. A universal sampler is a very powerful tool, but in many cases impractical, due to the need for a trusted setup.

Non-interactive multiparty computation (NIMPC). Non-interactive multiparty computation (NIMPC, [BGI⁺14a]) is a kind of one-round protocol that allows

n parties to compute any function of their secret inputs in just one round of communication. However, NIMPC requires that the parties know one another’s public keys before that one round, so there is another implicit round of communication.¹ NIMPC for general functions can be constructed based on subexponentially-secure indistinguishability obfuscation [HLJ⁺17].

Spooky encryption. Spooky encryption [DHRW16] is a kind of encryption which enables parties to learn joint functions of ciphertexts encrypted under independent public keys (given one of the corresponding secret keys). In order for semantic security to hold, what party i learns using her secret key should reveal nothing about the value encrypted to party j ’s public key; so, spooky encryption only supports the evaluation of *non-signaling* functions. An example of a non-signaling function is any function where the parties’ outputs are an additive secret sharing. Dodis *et al.* [DHRW16] show how to build spooky encryption for any such additive function from the LWE assumption with a URS (this also implies multi-party homomorphic secret sharing for general functions). In the two-party setting, they also show how to build spooky encryption for a larger class of non-signaling functions from (among other things) sub-exponentially hard indistinguishability obfuscation.

Pseudorandom Correlation Generators and Functions (PCGs and PCFs). Pseudorandom correlation generators [BCG⁺19a, BCG⁺19b, BCG⁺20b] and functions [BCG⁺20a, OSY21] let parties take a small amount of specially correlated randomness (called the *seed* randomness) and expand it non-interactively, obtaining a large sample from a target correlation. Pseudorandom correlation generators (PCGs) support only a fixed, polynomial expansion; pseudorandom correlation functions (PCFs) allow the parties to produce exponentially many instances of the correlation (via evaluation of the function on any of exponentially many inputs).

PCGs and PCFs can be built for any *additively secret shared* correlation (where the parties obtain additive shares of a sample from some distribution) using LWE-based spooky encryption mentioned above. Similarly, with two parties, we can build PCGs and PCFs for more general *reverse-samplable* correlations by relying on spooky encryption from subexponentially secure iO. PCGs and PCFs with better concrete efficiency can be obtained under different flavours of the LPN assumption, for simpler correlations such as vector oblivious linear evaluation [BCGI18], oblivious transfer [BCG⁺19b] and others [BCG⁺20b, BCG⁺20a].

Of course, in order to use PCGs or PCFs, the parties must somehow get the correlated seed randomness. *Public-key* PCGs and PCFs allow the parties to instead derive outputs using their independently generated public keys, which can be published in a single round of communication. The above, spooky

¹ This requirement is inherent; otherwise, an adversary would be able to take the message an honest party sent, and recompute the function with that party’s input while varying the other inputs. NIMPC does allow similar recomputation attacks, but only with *all* honest party inputs fixed, which a PKI can be used to enforce.

encryption-based PCGs and PCFs are public-key, while the LPN-based ones are not. Public-key PCFs for OT and vector-OLE were recently built based on DCR and QR [OSY21]; however, these require a structured reference string consisting of a public RSA modulus with hidden factorization.

1.2 Our Contributions

In this paper, we leverage indistinguishability obfuscation to build public-key PCFs for *any* correlation. On the way to realizing this, we define several other primitives, described in Fig. 1. One of these primitives is a *distributed sampler*, which is a weaker form of public-key PCF which only allows the sampling of public randomness. (A public-key PCF can be thought of as a distributed *correlation* sampler.) Our constructions, and the assumptions they use, are mapped out in Fig. 2. We pay particular attention to avoiding the use of sub-exponentially secure primitives where possible (which rules out strong tools such as probabilistic iO [CLTV15]).

Primitive	Distribution	Output
Distributed Sampler (DS, Def. 3.1)	fixed	public
Reusable Distributed Universal Sampler (Def. 7.6)	any	public
Public-key PCF (pk-PCF, [OSY21])	fixed, reverse-samplable	private
Ideal pk-PCF (Def. 7.2)	any	private

Fig. 1. In this table we describe one-round n -party primitives that can be used for sampling randomness. They differ in terms of whether a given execution enables sampling from *any* distribution (or just a fixed one), and in terms of whether they only output public randomness (in the form of a URS or SRS) or also return private correlated randomness to the parties.

We begin by exploring constructions secure against semi-malicious adversaries, where corrupt parties are assumed to follow the protocol other than in their choice of random coins. We build a semi-malicious distributed sampler, and use it to build a semi-malicious public-key PCF. We then compile those protocols to be secure against active adversaries. This leads to a public-key PCF that requires a random oracle, and supports the broad class of *reverse-samplable* correlations (where, given only corrupt parties’ values in a given sample, honest parties’ values can be simulated in such a way that they are indistinguishable from the ones in the original sample).

We also show two other routes to public-key PCFs with active security. One of these avoids the use of a random oracle, but requires sub-exponentially secure building blocks. The other requires a random oracle, but can support general correlations, not just reverse-samplable ones. (The downside is that it does not support correlations *with master secrets*, which allow parties to have secret input parameters to the correlation.)

These are valuable trade-offs, as described below.

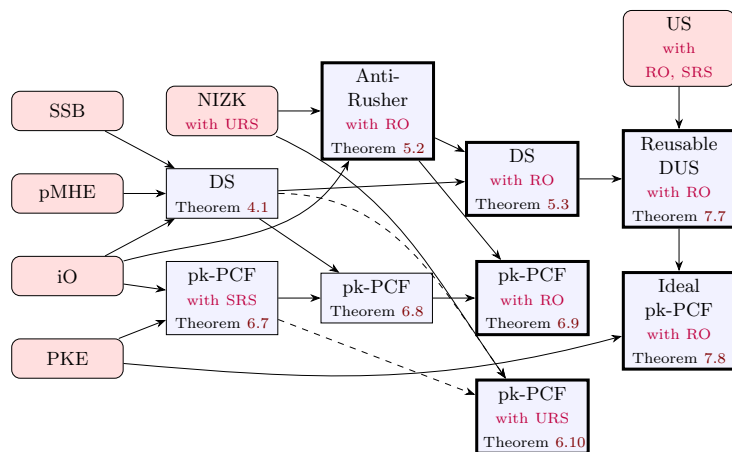


Fig. 2. In this table we describe the constructions in this paper. In pink are assumptions: they include somewhere statistically binding hash functions (SSB), multiparty homomorphic encryption with private evaluation (pMHE [AJJM20], a weaker form of multi-key FHE), indistinguishability obfuscation (iO), non-interactive zero knowledge proofs (NIZK), and universal samplers (US). In blue are constructions of distributed samplers (DS, Def. 3.1), reusable distributed universal samplers (reusable DUS, Def. 7.6) and public-key pseudorandom correlation functions (pk-PCFs, [OSY21]). Constructions with bold outlines are secure against active adversaries; the rest are secure against semi-malicious adversaries. In magenta are necessary setup assumptions. (Note that the availability of a random oracle (RO) immediately implies the additional availability of a URS.) Dashed lines denote the use of sub-exponentially secure tools.

- On the importance of realizing general correlations:** There are many valuable correlations that are *not* reverse-sampleable. An example of such a correlation gives a garbled circuit to one party, and all the wire labels to another. Since the labels cannot be reverse-sampled from the circuit (without violating the security properties of the garbling scheme), such a correlation is not reverse-sampleable; however, it is a valuable form of pre-processing for secure two-party computation.
- On the importance of avoiding a random oracle:** Random oracles are an idealized assumption with no known realization; there is a large gap between a random oracle and a hash function, which is often substituted for a random oracle in practice.
- On the importance of avoiding sub-exponential assumptions:** It may seem strange to want to avoid sub-exponentially secure primitives,² when many candidates for indistinguishability obfuscation itself are based on sub-exponential assumptions [JLS21]. However, despite informal arguments [LZ17], this is not known to be inherent: earlier iO candidates are based on polynomial hardness [GGH⁺13] (albeit for an exponential family of assumptions), and in future we may obtain iO from a single, polynomial hardness assumption. In general, it is always preferable to require a weaker form of security from a primitive, and this also leads to better parameters in practice. The problem of removing sub-exponential assumptions from iO, or applications of iO, has been studied previously in various settings [GPSZ17, LZ17].

1.3 Technical Overview

Distributed Samplers We start by introducing a new tool called a *distributed sampler* (DS, Section 3). A distributed sampler allows n parties to sample a single, public output from an efficiently sampleable distribution \mathcal{D} with just one round of communication (which is modelled by the exchange of public keys).

Semi-malicious distributed samplers. We use multiparty homomorphic encryption with private evaluation (pMHE [AJJM20], a weaker, setup-free version of multi-key FHE) and indistinguishability obfuscation to build semi-malicious distributed samplers in the plain model (Section 4). In our distributed sampler construction, all parties can compute an encryption of the sample from everyone's public keys (using, among other things, the homomorphic properties of the encryption scheme), and then use an obfuscated program in party i 's public key to get party i 's partial decryption of the sample. The partial decryptions can then be combined to recover the sample itself. The tricky thing is that, in the proof, we must ensure that we can replace the real sample with an ideal sample. To do this, we must remove all information about the real sample from the public keys. However, pMHE secret keys are not *puncturable*; that is, there is no way to ensure that they do not reveal any information about the contents of one cipher-

² By sub-exponential security, we mean that no PPT adversary cannot break the security of that primitive with probability better than $2^{-\lambda^c}$ for a constant c .

text, while correctly decrypting all others. We could, in different hybrids, hardcode the correct partial decryption for each of the exponentially many possible ciphertexts, but this would blow up the size of the obfuscated program. Therefore, instead of directly including a pMHE ciphertext in each party’s DS public key, we have each party obfuscate an additional program which produces a new pMHE ciphertext each time it is used. This way, when we need to remove all information about a given sample, we can remove the entire corresponding secret key (via the appropriate use of puncturable PRFs and hardcoded values). This technique may be useful for other primitives, such as NIMPC [BGI⁺14a] and probabilistic iO [CLTV15], to avoid the use of an exponential number of hybrids.

Achieving active security with a random oracle. Upgrading to active security is challenging because we need to protect against two types of attacks: malformed messages, and rushing adversaries, who wait for honest parties’ messages before sending their own. We protect against the former using non-interactive zero knowledge proofs. (This requires a URS which, though it is a form of setup, is much weaker than an SRS.) We protect against the latter via a generic transformation that we call an *anti-rusher* (Section 5.1). To use our anti-rusher, each party includes in her public key an obfuscated program which takes as input a hash (i.e. a random oracle output) of all parties’ public keys. It then samples *new* (DS) public keys, using this hash as a PRF nonce. This ensures that even an adversary who selects her public keys after seeing the honest party public keys cannot influence the selected sample other than by re-sampling polynomially many times.

Public-key PCFs We start by building a public-key PCF that requires an SRS (Section 6.3). The SRS consists of an obfuscated program that, given a nonce and n parties’ public encryption keys, uses a PRF to generate correlated randomness, and encrypts each party’s random output to its public key. We can then eliminate the need for a pre-distributed SRS by instead using a distributed sampler to sample it (Section 6.4).

Public-key PCFs without random oracles. The proofs of security for the constructions sketched above only require polynomially many hybrids, roughly speaking because the random oracle allows the simulator to predict and control the inputs to the obfuscated programs. We can avoid the use of the random oracle, at the cost of going through exponentially many hybrids in the proof of security, and thus requiring sub-exponentially secure primitives.

Public-key PCFs for any correlation with a random oracle. Boyle *et al.* [BCG⁺19b] prove that a public-key PCF in the plain model that can handle *any* correlation (not just reverse-sampleable ones) must have keys at least as large as all the correlated randomness it yields. We observe that we can use a random oracle to sidestep this lower bound by deriving additional randomness from the oracle.

As a stepping stone, we introduce a different flavour of the distributed sampler, which we call the *reusable distributed universal sampler* (reusable DUS).

It is *reusable* because it can be queried multiple times (without the need for additional communication), and it is *universal* because each query can produce a sample from a different distribution (specified by the querier). We build a reusable distributed universal sampler from a universal sampler, a random oracle and a distributed sampler (by using the distributed sampler to produce the universal sampler). Our last public-key PCF (Section 7) then uses the reusable distributed universal sampler to sample from a distribution that first picks the correlated randomness and then encrypts each party’s share under her public key.

2 Preliminaries

Notation. We denote the security parameter by λ and the set $\{1, 2, \dots, m\}$ by $[m]$. Our constructions are designed for an ordered group of n parties P_1, P_2, \dots, P_n . We will denote the set of (indexes of) corrupted parties by C , whereas its complementary, the set of honest players, is H .

We indicate the probability of an event E by $\mathbb{P}[E]$. We use the term *noticeable* to refer to a non-negligible quantity. A probability p is instead *overwhelming* if $1 - p$ is negligible. We say that a cryptographic primitive is sub-exponentially secure, if the advantage of the adversary is bounded by $2^{-\lambda^c}$ for some constant $c > 0$. When the advantage is negligible, we say that it is polynomially secure.

We use the simple arrow \leftarrow to assign the output of a deterministic algorithm $\text{Alg}(x)$ or a specific value a to a variable y , i.e. $y \leftarrow \text{Alg}(x)$ or $y \leftarrow a$. If Alg is instead probabilistic, we write $y \stackrel{\$}{\leftarrow} \text{Alg}(x)$ and we assume that the random tape is sampled uniformly. If the latter is set to a particular value r , we write however $y \leftarrow \text{Alg}(x; r)$. We use $\stackrel{\$}{\leftarrow}$ also if we sample the value of y uniformly over a set X , i.e. $y \stackrel{\$}{\leftarrow} X$. Finally, we refer to algorithms having no input as distributions. The latter are in most cases parametrised by λ . The terms *circuit* and *program* are used interchangeably.

2.1 Indistinguishability Obfuscation

We recall the formal definition of indistinguishability obfuscation (iO) [BGI⁺01, GGH⁺13]. Informally speaking an obfuscator is an efficient algorithm that “scrambles” any given circuit Cr until it is impossible to extract any information about Cr except its original input-output behaviour. Furthermore, the result of the operation is another program computing exactly the same function as Cr . We provide a formal definition.

Definition 2.1 (Indistinguishability Obfuscator). *Let $(\mathcal{L}_\lambda)_{\lambda \in \mathbb{N}}$ be a class of circuits such that every $\text{Cr} \in \mathcal{L}_\lambda$ maps a $\text{inp}(\lambda)$ -bit input into a $\text{out}(\lambda)$ -bit output.*

An indistinguishability obfuscator for $(\mathcal{L}_\lambda)_{\lambda \in \mathbb{N}}$ is a PPT algorithm iO with the following properties.

- **Correctness.** *For every $\lambda \in \mathbb{N}$, circuit $\text{Cr} \in \mathcal{L}_\lambda$ and input $x \in \{0, 1\}^{\text{inp}(\lambda)}$*

$$\mathbb{P}[\text{Cr}'(x) = \text{Cr}(x) \mid \text{Cr}' \stackrel{\$}{\leftarrow} \text{iO}(\mathbb{1}^\lambda, \text{Cr})] = 1$$

- **Security.** For every pair of circuits $\text{Cr}_0, \text{Cr}_1 \in \mathcal{L}_\lambda$ such that $\text{Cr}_0(x) = \text{Cr}_1(x)$ for each $x \in \{0, 1\}^{\text{inp}(\lambda)}$, the distributions given by $\text{iO}(\mathbb{1}^\lambda, \text{Cr}_0)$ and $\text{iO}(\mathbb{1}^\lambda, \text{Cr}_1)$ are computationally indistinguishable.

Observe that the obfuscator is tailored to a specific class of circuits. The latter often affects the size of the obfuscated programs, increasing it as the class becomes larger.

The first candidate indistinguishability obfuscator was presented in 2013 by Garg *et al.* [GGH⁺13]. The construction relies on subexponentially secure primitives. Subsequent work has more and more weakened the requirements under which it is possible to build obfuscation. However, all known constructions still rely on sub-exponentially secure primitives or exponentially-many assumptions. Indistinguishability obfuscators not suffering from these problems are still purpose of research.

2.2 Puncturable PRFs

We recall the formal definition of puncturable PRF [KPTZ13, BW13, BGI14b]. A puncturable PRF is a PRF construction F in which the keys K have an additional property: we can puncture them in any position x , obtaining a possibly longer key containing no information about $F_K(x)$ but still allowing computing $F_K(y)$ for every $y \neq x$.

Definition 2.2 (Puncturable PRF). A puncturable PRF with output space $(\mathcal{X}_\lambda)_{\lambda \in \mathbb{N}}$ is a pair of PPT algorithms (F, Punct) with the following properties.

- **Correctness.** For every $\lambda \in \mathbb{N}$, key $K \in \{0, 1\}^\lambda$ and values $x, y \in \{0, 1\}^*$ with $x \neq y$

$$\mathbb{P}[F_K(y) = F_{\hat{K}}(y) \mid \hat{K} \leftarrow \text{Punct}(K, x)] = 1$$

- **Security.** For every value $x \in \{0, 1\}^\lambda$, the following two distributions are indistinguishable.

$$\left\{ (\hat{K}, r) \left| \begin{array}{l} K \xleftarrow{\$} \{0, 1\}^\lambda \\ \hat{K} \leftarrow \text{Punct}(K, x) \\ r \leftarrow F_K(x) \end{array} \right. \right\} \quad \left\{ (\hat{K}, r) \left| \begin{array}{l} K \xleftarrow{\$} \{0, 1\}^\lambda \\ \hat{K} \leftarrow \text{Punct}(K, x) \\ r \xleftarrow{\$} \mathcal{X}_\lambda \end{array} \right. \right\}$$

As noticed in [BW13], it is easy to build puncturable PRFs by relying on the GGM construction [GGM86].

2.3 Simulation-Extractable NIZKs

We recall the formal definition of simulation-extractable NIZK [GO07]. Let \mathcal{R} be an efficiently computable relation consisting of pairs (x, w) , where x is called statement and w witness. As widely known, a NIZK for \mathcal{R} is a construction that allows a party to prove the knowledge of a witness w for a statement x with

only one message and without revealing any information about w itself (zero-knowledge). The procedure relies on a CRS, which, depending on the construction, can be structured or unstructured. Zero-knowledge is formulated by requiring the existence of two PPT simulators: the first one generates a fake CRS embedding a trapdoor τ into it, the second one leverages the knowledge of τ to produce valid proofs for various statements without needing the corresponding witnesses.

A simulation-extractable NIZK has also an additional property: there exists a PPT algorithm that, in conjunction with the two simulators, is able to extract the witness from any valid proof generated by the adversary. Such algorithm is called extractor and exploits the knowledge of the trapdoor in the CRS. We now formalise the syntax and the security properties of what we have just described.

Definition 2.3 (Non-Interactive Proof). *A non-interactive proof for a relation \mathcal{R} is a triple of PPT algorithms $(\text{Gen}, \text{Prove}, \text{Verify})$ with the following syntax.*

- *Gen takes as input the security parameter $\mathbb{1}^\lambda$ and outputs a CRS.*
- *Prove takes as input the security parameter $\mathbb{1}^\lambda$, the CRS, a statement x and a witness w . The output is a proof π for x .*
- *Verify takes as input a CRS, a proof π and a statement x . The output is a bit indicating whether the proof has been accepted or not.*

Definition 2.4 (Simulation-Extractable NIZK). *A non-interactive proof $(\text{Gen}, \text{Prove}, \text{Verify})$ for the relation \mathcal{R} is a simulation extractable NIZK if it satisfies the following properties.*

- **Completeness.** *For every $(x, w) \in \mathcal{R}$,*

$$\mathbb{P} \left[\text{Verify}(\text{crs}, \pi, x) = 1 \mid \begin{array}{l} \text{crs} \xleftarrow{\$} \text{Gen}(\mathbb{1}^\lambda) \\ \pi \xleftarrow{\$} \text{Prove}(\mathbb{1}^\lambda, \text{crs}, x, w) \end{array} \right] = 1$$

- **Multi-Theorem Zero Knowledge.** *There exist PPT simulators Sim_1 and Sim_2 such that, for every set of pairs $\{(x_i, w_i)\}_{i \in [m]}$ in \mathcal{R} , no PPT adversary is able to distinguish between the following distributions.*

$$\left\{ (\text{crs}, (x_i, \pi_i)_{i \in [m]}) \mid \begin{array}{l} \text{crs} \xleftarrow{\$} \text{Gen}(\mathbb{1}^\lambda) \\ \forall i \in [m] : \pi_i \xleftarrow{\$} \text{Prove}(\mathbb{1}^\lambda, \text{crs}, x_i, w_i) \end{array} \right\}$$

$$\left\{ (\text{crs}, (x_i, \pi_i)_{i \in [m]}) \mid \begin{array}{l} (\text{crs}, \tau) \xleftarrow{\$} \text{Sim}_1(\mathbb{1}^\lambda) \\ \forall i \in [m] : \pi_i \xleftarrow{\$} \text{Sim}_2(\text{crs}, \tau, x_i) \end{array} \right\}$$

- **Simulation-Extractability.** *There exists a PPT extractor Extract such that, for every set of statements $\{x_i\}_{i \in [m]}$ and PPT adversary \mathcal{A} , we have*

that

$$\mathbb{P} \left[\begin{array}{l} \forall i \in [m] : (\pi', x') \neq (\pi_i, x_i) \\ \text{Verify}(\text{crs}, \pi', x') = 1 \\ (x', w') \notin \mathcal{R} \end{array} \middle| \begin{array}{l} (\text{crs}, \tau) \xleftarrow{\$} \text{Sim}_1(\mathbb{1}^\lambda) \\ \forall i \in [m] : \pi_i \xleftarrow{\$} \text{Sim}_2(\text{crs}, \tau, x_i) \\ (x', \pi') \xleftarrow{\$} \mathcal{A}(\mathbb{1}^\lambda, \text{crs}, (x_i, \pi_i)_{i \in [m]}) \\ w' \leftarrow \text{Extract}(\text{crs}, \tau, x', \pi') \end{array} \right] = \text{negl}(\lambda)$$

Observe that simulation-extractability and zero-knowledge imply soundness.

NIZKs can be classified based on the type of CRS. When the latter is a uniform string of bits, we talk about NIZK with URS (uniform reference string). When the CRS is instead structured and possibly depends on secrets that must not be revealed to the parties, we talk about NIZKs with SRS (structured reference string).

2.4 MHE with Private Evaluation

We recall the definition of MHE with private evaluation [AJJM20]. This is a construction that permits any party P_i to encrypt a value x_i obtaining a ciphertext c_i and the corresponding partial decryption key sk_i . Once given the ciphertexts of the other parties, P_i can apply a circuit Cr on $(c_j)_{j \in [n]}$, obtaining the i -th partial decryption d_i by means of pk_i . By pooling the partial plaintexts $(d_j)_{j \in [n]}$, it is then possible to obtain $\text{Cr}(x_1, \dots, x_n)$ without learning any additional information on the inputs. The construction differs from multi-key FHE as there actually exists no public key but only a private one that changes from ciphertext to ciphertext. Furthermore, the encryption algorithm needs to be provided with the parameters (input size, output size and depth) of the circuits we are going to apply on the ciphertexts. The final decryption needs instead to know the exact circuit that was used to produce the partial plaintext it is provided with. It is possible to build MHE schemes with private evaluation from polynomially secure LWE [AJJM20].

Definition 2.5 (Multiparty Homomorphic Encryption with Private Evaluation). *A multiparty homomorphic encryption scheme with private evaluation (pMHE) is a triple of PPT algorithms $(\text{Enc}, \text{PrivEval}, \text{FinDec})$ with the following syntax.*

- Enc is a randomised algorithm that takes as input the security parameter $\mathbb{1}^\lambda$, the parameters of a circuit (input size, output size and depth), an index i and a message x_i . The output is a ciphertext c_i and a partial decryption key sk_i .
- PrivEval is a randomised algorithm. It takes as input a partial decryption key sk_i , the description of a circuit Cr mapping n inputs into a single output and n ciphertexts c_1, c_2, \dots, c_n . The output is a partial plaintext d_i .
- FinDec is a deterministic algorithm taking as input a circuit Cr and the corresponding n partial decryptions d_1, d_2, \dots, d_n . The output is a plaintext d .

Definition 2.6 (Reusable Semi-Malicious Security). We say that a pMHE scheme satisfies reusable semi-malicious security if it satisfies the following properties.

- **Correctness.** For every security parameter $\lambda \in \mathbb{N}$, circuit Cr mapping n inputs into one output and inputs x_1, x_2, \dots, x_n , we have

$$\mathbb{P} \left[d = \text{res} \begin{array}{l} \forall i \in [n] : (c_i, \text{sk}_i) \stackrel{\$}{\leftarrow} \text{Enc}(\mathbb{1}^\lambda, \text{Cr.params}, i, x_i) \\ \forall i \in [n] : d_i \stackrel{\$}{\leftarrow} \text{Eval}(\text{sk}_i, \text{Cr}, c_1, c_2, \dots, c_n) \\ d \leftarrow \text{FinDec}(\text{Cr}, d_1, d_2, \dots, d_n) \\ \text{res} \leftarrow \text{Cr}(x_1, x_2, \dots, x_n) \end{array} \right] = 1$$

- **Reusable Semi-Malicious Security.** There exist PPT simulators Sim_1 and Sim_2 such that, for every $n \in \mathbb{N}$, set of corrupted parties C , circuit Cr , inputs of the honest parties $(x_h)_{h \in H}$, no PPT adversary is able to win the game $\mathcal{G}_{\text{pMHE}}^{C, \text{Cr}, (x_h)_{h \in H}}(\lambda)$ (see Fig. 3) with noticeable advantage.

$\mathcal{G}_{\text{pMHE}}^{C, \text{Cr}, (x_h)_{h \in H}}(\lambda)$

Initialisation.

1. $b \stackrel{\$}{\leftarrow} \{0, 1\}$
2. $\forall h \in H : (c_h^0, \text{sk}_h^0) \stackrel{\$}{\leftarrow} \text{Enc}(\mathbb{1}^\lambda, \text{Cr.params}, h, x_h)$
3. $(\tau, (c_h^1)_{h \in H}) \stackrel{\$}{\leftarrow} \text{Sim}_1(\mathbb{1}^\lambda, H, \text{Cr.params})$
4. Activate the adversary with input $(\mathbb{1}^\lambda, (c_h^b)_{h \in H})$.
5. Receive the inputs $(x_i)_{i \in C}$ and the randomness $(r_i)_{i \in C}$ of the corrupted parties from the adversary.
6. $\forall i \in C : (c_i^0, \text{sk}_i^0) \leftarrow \text{Enc}(\mathbb{1}^\lambda, \text{Cr.params}, i, x_i; r_i)$

Decryption Queries. On input $(\text{Decrypt}, \text{Cr}')$ where $\text{Cr}'.\text{params} = \text{Cr.params}$.

1. $\forall h \in H : d_h^0 \stackrel{\$}{\leftarrow} \text{PrivEval}(\text{sk}_h, \text{Cr}', c_1^0, c_2^0, \dots, c_n^0)$
2. $(\tau, (d_h^1)_{h \in H}) \stackrel{\$}{\leftarrow} \text{Sim}_2(\tau, \text{Cr}', \text{Cr}'(x_1, x_2, \dots, x_n), (x_i, r_i)_{i \in C})$
3. Reply with $(d_h^b)_{h \in H}$.

Output. The adversary wins if the final output is b .

Fig. 3. The pMHE Security Game

Observe that correctness states that if we evaluate a circuit Cr over encrypted values x_1, x_2, \dots, x_n and we pool the partial decryptions, we obtain $C(x_1, x_2, \dots, x_n)$. Security instead declares that if we publish the encryption of a value x_i and we later on disclose the partial plaintext corresponding to a certain circuit Cr , we reveal nothing more than the output of Cr .

2.5 Somewhere Statistically Binding Hash Functions

We recall the formal definition of somewhere statistically binding hashing (SSB) [HW15], which can be regarded as a more powerful, obfuscation-friendly version

of hash functions. The notion was first introduced by Hubáček and Wichs in [HW15] and can be constructed from FHE [HW15].

The primitive is composed of two algorithms `Gen` and `Hash`. The first one is used to generate a key hiding a special index i . The second one instead produces a digest after receiving the key and a message as inputs. The interesting property of SSB hashing is that there exists no pair of messages with different i -th block but same digest. We provide now formal definitions.

Definition 2.7 (Somewhere Statistically Binding Hash Function). *A somewhere statistically binding hash function (SSB) with block alphabet Σ is a pair of PPT algorithms $(\text{Gen}, \text{Hash})$ with the following syntax.*

- `Gen` is a randomised procedure taking as input the security parameter $\mathbb{1}^\lambda$, the maximum number of blocks B and an index $i \in [B]$. The output is a hash key hk .
- `Hash` is a deterministic procedure taking as input a hash key hk and a message $x \in \Sigma^B$ where B is the maximum number of blocks supported by hk . The output is a digest y .

Definition 2.8 (Security of SSB Hashing). *A somewhere statistically binding hash function is secure if it satisfies the following properties.*

- **Index Hiding.** *For every $B \in \mathbb{N}$ and $i, j \in [B]$, no PPT adversary can distinguish between $\text{Gen}(\mathbb{1}^\lambda, B, i)$ and $\text{Gen}(\mathbb{1}^\lambda, B, j)$.*
- **Somewhere Statistically Binding.** *For every $B \in \mathbb{N}$ and $i \in [B]$,*

$$\mathbb{P} \left[\begin{array}{c} \exists x, x' \in \Sigma^B \\ x_i \neq x'_i \\ \text{Hash}(\text{hk}, x) = \text{Hash}(\text{hk}, x') \end{array} \middle| \text{hk} \stackrel{\$}{\leftarrow} \text{Gen}(\mathbb{1}^\lambda, B, i) \right] = \text{negl}(\lambda).$$

Observe that the SSB property, in conjunction with index hiding, implies the collision resistance of the hash function.

The definition described above is actually weaker than the original one. In [HW15], an SSB hash function featured two additional algorithms `Prove` and `Verify`. The first one was used to generate short openings of the j -th preimage block of a digest, for any $j \in [B]$. The second one was used to verify them. These two algorithms are not used in this work.

2.6 Universal Samplers

We recall the definitions of universal sampler (US) [HJK⁺16], describing the syntax and the various security notions. Informally speaking, a US is a trusted-dealer-based construction that permits to deterministically derive samples from any distribution, learning no information about the corresponding randomness. This primitive is fundamental for the construction of our distributed universal samplers (see Section 7.1 and Section 7.1).

Definition 2.9 (Universal Sampler). Let $\ell(\lambda)$, $r(\lambda)$ and $t(\lambda)$ be polynomials. A universal sampler for (ℓ, r, t) -distributions is a pair of PPT algorithms (**Setup**, **Sample**) with the following syntax.

1. **Setup** is a randomised algorithm taking as input the security parameter $\mathbb{1}^\lambda$ and outputting a sampler U .
2. **Sample** is a deterministic procedure possibly interacting with a random oracle \mathcal{H} . It takes as input a sampler U and the description \mathcal{D} of an (ℓ, r, t) -distribution, outputting a sample R .

Observe that in the above definition, we ask that **Setup** does not interact with any random oracle. If that was not the case, it would be impossible to generate a universal sampler inside an obfuscated program. The original definition of US [HJK⁺16] was not as restrictive. However, all the constructions presented in [HJK⁺16] satisfy the property described above, including those on which we rely in this work.

One-time universal samplers. We now formalise the weaker security notion of the primitive. In particular, we require that the samples look random only for one specific distribution selected ahead of time. In [HJK⁺16], the authors present a construction satisfying this definition without random oracle.

Definition 2.10 (One-Time Universal Sampler). A universal sampler (**Setup**, **Sample**) for (ℓ, r, t) -distributions satisfies one-time security if there exists a PPT simulator **Sim** such that, for every (ℓ, r, t) -distribution \mathcal{D} , the following two distributions are computationally indistinguishable

$$\left\{ U, R \left| \begin{array}{l} U \stackrel{s}{\leftarrow} \text{Setup}(\mathbb{1}^\lambda) \\ R \leftarrow \text{Sample}(U, \mathcal{D}) \end{array} \right. \right\} \quad \left\{ U, R \left| \begin{array}{l} R \stackrel{s}{\leftarrow} \mathcal{D} \\ U \stackrel{s}{\leftarrow} \text{Sim}(\mathbb{1}^\lambda, \mathcal{D}, R) \end{array} \right. \right\}$$

Notice that security states that one-time USs can be programmed to output ideal samples from the distribution \mathcal{D} , without the adversary noticing it.

Adaptively secure universal samplers. We present the stronger security notion of universal samplers. We now ask that the samples look random for every distribution adaptively chosen by the adversary. In particular, we do not care only about one distribution, but multiple ones. This definition can only be satisfied in the random oracle model. Hofheinz *et al.* presented an example of such construction in [HJK⁺16].

Definition 2.11 (Adaptively Secure Universal Sampler). A universal sampler (**Setup**, **Sample**) for (ℓ, r, t) -distributions satisfies adaptive security if there exist PPT simulators **SimGen** and **SimRO** such that no PPT adversary can win the game $\mathcal{G}_{US}(\lambda)$ (see Fig. 4) with noticeable advantage.

As for the one-time case, security states that the sampler U and the random oracle \mathcal{H} could be programmed to output ideal samples, without the adversary noticing it. The only difference is that now the adversary can adaptively choose the distribution multiple times.

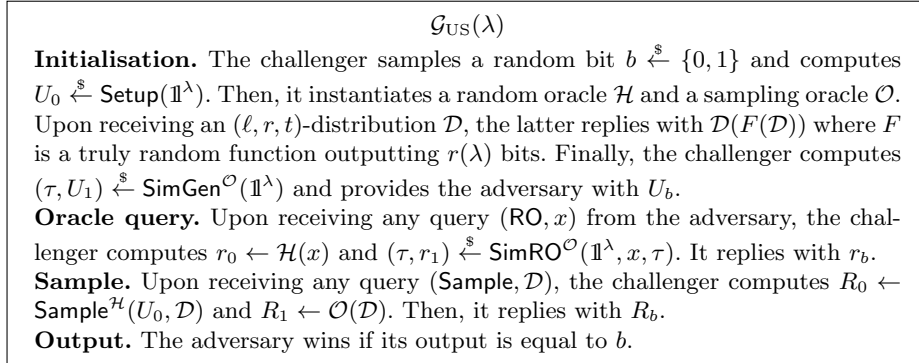


Fig. 4. The Universal Sampler Game

3 Defining Distributed Samplers

Informally speaking, a distributed sampler (DS) for the distribution \mathcal{D} is a construction that allows n parties to obtain a random sample R from \mathcal{D} with just one round of communication and without revealing any additional information about the randomness used for the generation of R . The output of the procedure can be derived given only the public transcript, so we do not aim to protect the privacy of the result against passive adversaries eavesdropping the communications between the parties.

If we assume an arbitrary trusted setup, building a DS becomes straightforward; we can consider the trivial setup that directly provides the parties with a random sample from \mathcal{D} . Obtaining solutions with a weaker (or no) trusted setup is much more challenging.

The structure and syntax of distributed samplers is formalised as follows. We then analyse different flavours of security definitions.

Definition 3.1 (n -party Distributed Sampler for the Distribution \mathcal{D}).

An n -party distributed sampler for the distribution \mathcal{D} is a pair of PPT algorithms $(\text{Gen}, \text{Sample})$ with the following syntax:

1. **Gen** is a probabilistic algorithm taking as input the security parameter $\mathbb{1}^\lambda$ and a party index $i \in [n]$ and outputting a sampler share U_i for party i . Let $\{0, 1\}^{L(\lambda)}$ be the space from which the randomness of the algorithm is sampled.
2. **Sample** is a deterministic algorithm taking as input n shares of the sampler U_1, U_2, \dots, U_n and outputting a sample R .

In some of our security definitions, we will refer to the *one-round protocol* Π_{DS} that is induced by the distributed sampler $\text{DS} = (\text{Gen}, \text{Sample})$. This is the natural protocol obtained from DS, where each party first broadcasts a message output by **Gen**, and then runs **Sample** on input all the parties' messages.

3.1 Security

In this section we formalise the definition of distributed samplers with relation to different security flavours, namely, semi-malicious and active. We always assume that we deal with a static adversary who can corrupt up to $n - 1$ out of the n parties. We recall that a protocol has semi-malicious security if it remains secure even if the corrupt parties behave semi-honestly, but the adversary can select their random tapes.

Definition 3.2 (Distributed Sampler with Semi-Malicious Security). *A distributed sampler $(\text{Gen}, \text{Sample})$ has semi-malicious security if there exists a PPT simulator Sim such that, for every set of corrupt parties $C \subsetneq [n]$ and corresponding randomness $(\rho_i)_{i \in C}$, the following two distributions are computationally indistinguishable:*

$$\left\{ \begin{array}{l} (U_i)_{i \in [n]} \\ (\rho_i)_{i \in C}, R \end{array} \left| \begin{array}{ll} \rho_i \stackrel{\$}{\leftarrow} \{0, 1\}^{L(\lambda)} & \forall i \in H \\ U_i \leftarrow \text{Gen}(\mathbb{1}^\lambda, i; \rho_i) & \forall i \in [n] \\ R \leftarrow \text{Sample}(U_1, U_2, \dots, U_n) & \end{array} \right. \right\} \quad \text{and}$$

$$\left\{ \begin{array}{l} (U_i)_{i \in [n]} \\ (\rho_i)_{i \in C}, R \end{array} \left| \begin{array}{l} R \stackrel{\$}{\leftarrow} \mathcal{D}(\mathbb{1}^\lambda) \\ (U_i)_{i \in H} \stackrel{\$}{\leftarrow} \text{Sim}(\mathbb{1}^\lambda, C, R, (\rho_i)_{i \in C}) \end{array} \right. \right\}$$

Observe that this definition implies that, even in the simulation, the relation

$$R = \text{Sample}(U_1, U_2, \dots, U_n)$$

holds with overwhelming probability. In other words, security requires that $(\text{Gen}, \text{Sample})$ securely implements the functionality that samples from \mathcal{D} and outputs the result to all of the parties.

Observe that the previous definition can be adapted to passive security by simply sampling the randomness of the corrupted parties inside the game in the real world and generating it using the simulator in the ideal world.

We now define actively secure distributed samplers. Here, to handle the challenges introduced by a rushing adversary, we model security by defining an ideal functionality in the universal composability (UC) framework [Can01], and require that the protocol Π_{DS} securely implements this functionality.

Definition 3.3 (Distributed Sampler with Active Security). *Let $\text{DS} = (\text{Gen}, \text{Sample})$ be a distributed sampler for the distribution \mathcal{D} . We say that DS has active security if the one-round protocol Π_{DS} securely implements the functionality $\mathcal{F}_{\mathcal{D}}$ (see Fig. 5) against a static and active adversary corrupting up to $n - 1$ parties.*

Remark 3.4 (Distributed Samplers with a CRS or Random Oracle). Our constructions with active security rely on a setup assumption in the form of a common reference string (CRS) and random oracle. For a CRS, we assume the algorithms $\text{Gen}, \text{Sample}$ are implicitly given the CRS as input, which is modelled

$\mathcal{F}_{\mathcal{D}}$

Initialisation. On input `Init` from every honest party and the adversary, the functionality activates and sets $Q := \emptyset$. (Q will be used to keep track of queries.) If all the parties are honest, the functionality outputs $R \stackrel{\$}{\leftarrow} \mathcal{D}(\mathbb{1}^\lambda)$ to every honest party and sends R to the adversary, then it halts.

Query. On input `Query` from the adversary, the functionality samples $R \stackrel{\$}{\leftarrow} \mathcal{D}(\mathbb{1}^\lambda)$ and creates a fresh label `id`. It sends (id, R) to the adversary and adds the pair to Q .

Output. On input $(\text{Output}, \widehat{\text{id}})$ from the adversary, the functionality retrieves the only pair $(\text{id}, R) \in Q$ with $\text{id} = \widehat{\text{id}}$. If such pair does not exist, the functionality does nothing. Otherwise, it outputs R to every honest party and terminates.

Abort. On input `Abort` from the adversary, the functionality outputs \perp to every honest party and terminates.

Fig. 5. Distributed Sampler Functionality

as being sampled by an ideal setup functionality. As usual, the random oracle is modelled as an external oracle that may be queried by any algorithm or party, and programmed by the simulator in the security proof.

Observe that this definition allows the adversary to request several samples R from the functionality, and then select the one it likes the most. Our definition must allow this in order to deal with a rushing adversary who might wait for the messages $(U_i)_{i \in H}$ of all the honest parties and then locally re-generate the corrupt parties' messages $(U_i)_{i \in C}$, obtaining a wide range of possible outputs. Finally, it can broadcast the corrupt parties' messages that lead to the output it likes the most. This makes distributed samplers with active security rather useless when the distribution \mathcal{D} has low entropy, i.e. when there exists a polynomial-size set S such that $\mathcal{D}(\mathbb{1}^\lambda) \in S$ with overwhelming probability. Indeed, in such cases, the adversary is able to select its favourite element in the image of \mathcal{D} .

On the usefulness of distributed samplers with a CRS. Our distributed samplers with active security require a CRS for NIZK proofs. Since one of the main goals of the construction is avoid trusted setup in multiparty protocols, assuming the existence of a CRS, which itself is some form of setup, may seem wrong.

We highlight, however, that some types of CRS are much easier to generate than others. A CRS that depends on values which must remain secret (e.g. an RSA modulus with unknown factorization, or an obfuscated program which contains a secret key) is difficult to generate. However, assuming the security of trapdoor permutations [FLS90], bilinear maps [GOS06], learning with errors [PS19] or indistinguishability obfuscation [BP15], we can construct NIZK proofs where the CRS is just a random string of bits, i.e. a URS. In the random oracle model, such a CRS can even be generated without any interaction. So, the CRS required by our constructions is the simplest, weakest kind of CRS setup.

4 A Construction with Semi-Malicious Security

We now present the main construction of this paper: a distributed sampler with semi-malicious security based on polynomially secure MHE with private evaluation and indistinguishability obfuscation. In Section 5, we explain how to upgrade this construction to achieve active security.

The basic idea. Our goal is to generate a random sample R from the distribution \mathcal{D} . The natural way to do it is to produce a random bit string s and feed it into \mathcal{D} . We want to perform the operation in an encrypted way as we need to preserve the privacy of s . A DS implements the functionality that provides samples from the underlying distribution, but not the randomness used to obtain them, so no information about s can be leaked.

We guarantee that any adversary corrupting up to $n - 1$ parties is not able to influence the choice of s by XORing n bit strings of the same length, the i -th one of which is independently sampled by the i -th party P_i . Observe that we are dealing with a semi-malicious adversary, so we do not need to worry about corrupted parties adaptively choosing their shares after seeing those of the honest players.

Preserving the privacy of the random string. To preserve the privacy of s , we rely on an MHE scheme with private evaluation $\text{pMHE} = (\text{Enc}, \text{PrivEval}, \text{FinDec})$. Each party P_i encrypts s_i , publishing the corresponding ciphertext c_i and keeping the private key sk_i secret. As long as the honest players do not reveal their partial decryption keys, the privacy of the random string s is preserved. Using the homomorphic properties of the MHE scheme, the parties are also able to obtain partial plaintexts of R without any interactions. However, we run into an issue: in order to finalise the decryption, the construction would require an additional round of communication where the partial plaintexts are broadcast.

Reverting to a one-round construction. We need to find a way to perform the final decryption without additional interaction, while at the same time preserving the privacy of the random string s . That means revealing a very limited amount of information about the private keys $\text{sk}_1, \text{sk}_2, \dots, \text{sk}_n$, so that it is only possible to retrieve R , revealing nothing more.

Inspired by [HIJ⁺17], we build such a precise tool by relying on indistinguishability obfuscation: in the only round of interaction, each party P_i additionally publishes an obfuscated *evaluation program* EvProg_i containing the private key sk_i . When given the ciphertexts of the other parties, EvProg_i evaluates the circuit producing the final result R and outputs the partial decryption with relation to sk_i . Using the evaluation programs, the players are thus able to retrieve R by feeding the partial plaintexts into pMHE.FinDec .

Dealing with the leakage about the secret keys. At first glance, the solution outlined in the previous paragraph seems to be secure. However, there are some sneaky issues we need to deal with.

In this warm-up construction, we aim to protect the privacy of the random string s by means of the reusable semi-malicious security of the MHE scheme with private evaluation. To rely on this assumption, no information on the secret keys must be leaked. However, this is not the case here, as the private keys are part of the evaluation programs.

In the security proof, we are therefore forced to proceed in two steps: first, we must remove the secret keys from the programs using obfuscation, and then we can apply reusable semi-malicious security. The first task is actually trickier than it may seem. iO states we cannot distinguish between the obfuscation of two equivalent programs. Finding a program with the same input-output behaviour as EvProg_i without it containing any information about sk_i is actually impossible, as any output of the program depends on the private key. We cannot even hard-code the partial decryptions under sk_i for all possible inputs into the obfuscated program as that would require storing an exponential amount of information, blowing up the size of EvProg_i .

In [HIJ⁺17], while constructing an NI-MPC protocol based on multi-key FHE and iO , the authors deal with an analogous issue by progressively changing the behaviour of the program input by input, first hard-coding the output corresponding to a specific input and then using the simulatability of partial decryptions to remove any dependency on the multi-key FHE secret key. Unfortunately, in our context, this approach raises additional problems. First of all, in contrast with some multi-key FHE definitions, MHE does not support simulatability of partial decryptions. Additionally, since the procedure of [HIJ⁺17] is applied input by input, the security proof would require exponentially many hybrids. In that case, security can be argued only if transitions between subsequent hybrids cause a subexponentially small increase in the adversary’s advantage. In other words, we would need to rely on subexponentially secure primitives even if future research shows that iO does not. Finally, we would still allow the adversary to compute several outputs without changing the random strings $(s_h)_{h \in H}$ selected by the honest parties. Each of the obtained values leaks some additional information about the final output of the distributed sampler. In [HIJ⁺17], this fact did not constitute an issue as this type of leakage is intrinsically connected to the notion of NI-MPC.

Bounding the leakage: key generation programs. To avoid the problems described above, we introduce the idea of *key generation programs*. Each party P_i publishes an obfuscated program KGProg_i which encrypts a freshly chosen string s_i , keeping the corresponding partial decryption key secret.

The randomness used by KGProg_i is produced via a puncturable PRF F taking as a nonce the key generation programs of the other parties. In this way, any slight change in the programs of the other parties leads to a completely unrelated string s_i , ciphertext c_i and key sk_i . It is therefore possible to protect the privacy of s_i using a polynomial number of hybrids, as we need only worry about a single combination of inputs. Specifically, we can remove any information about sk_i from EvProg_i and hard-code the partial plaintext d_i corresponding to $(c_j)_{j \in [n]}$. At that point, we can rely on the reusable semi-malicious security of the

MHE scheme with private evaluation, removing any information about s_i from c_i and d_i and programming the final output to be a random sample R from \mathcal{D} .

The introduction of the key generation programs requires minimal modifications to the evaluation programs. In order to retrieve the MHE private key, EvProg_i needs to know the same PRF key K_i used by KGProg_i . Moreover, it now takes as input the key generation programs of the other parties, from which it will derive the MHE ciphertexts needed for the computation of R . Observe that EvProg_i will also contain KGProg_i , which will be fed into the other key generation programs in a nested execution of obfuscated circuits.

Compressing the inputs. The only problem with the construction above, is that we now have a circularity issue: we cannot actually feed one key generation program as input to another key generation program, since the programs are of the same size. This holds even if we relied on obfuscation for Turing machines, since to prove security, we would need to puncture the PRF keys in the nonces, i.e. the key generation programs of the other parties. The point at which the i -th key is punctured, which is at least as big as the program itself, must be hard-coded into KGProg_i , which is clearly too small.

Instead of feeding entire key generation programs into KGProg_i , we can input their hash, which is much smaller. This of course means that there now exist different combinations of key generation programs leading to the same MHE ciphertext-key pair (c_i, sk_i) , and the adversary could try to extract information about sk_i by looking for collisions. The security of the hash function should, however, prevent this attack. The only issue is that iO does not really get along with this kind of argument based on collision-resistant hashing. We instead rely on the more iO -friendly notion of a *somewhere statistically binding* hash function $\text{SSB} = (\text{Gen}, \text{Hash})$ [HW15].

Final construction. We now present the formal description of our semi-maliciously secure DS. The algorithms Gen and Sample , as well as the unobfuscated key generation program \mathcal{P}_{KG} and evaluation program $\mathcal{P}_{\text{Eval}}$, can be found in Fig. 6. In the description, we assume that the puncturable PRF F outputs pseudorandom strings (r_1, r_2, r_3) where each of r_1, r_2 and r_3 is as long as the randomness needed by \mathcal{D} , pMHE.Enc , and HE.PrivEval respectively. Moreover, we denote by B the maximum number of blocks in the messages fed into SSB.Hash .

Theorem 4.1. *If $\text{SSB} = (\text{Gen}, \text{Hash})$ is a somewhere statistically binding hash function, $\text{pMHE} = (\text{Enc}, \text{PrivEval}, \text{FinDec})$ is a MHE scheme with private evaluation, iO is an indistinguishability obfuscator and (F, Punct) is a puncturable PRF, the construction in Fig. 6 is an n -party distributed sampler with semi-malicious security for the distribution \mathcal{D} .*

Proof. We prove the security of the construction in Fig. 6 in a sequence of hybrids. In the initial hybrid (hybrid 0), we start with the real game; in the last hybrid (hybrid 6), we produce a simulated sampler share on behalf of one of the honest parties, which leads the parties to output an R sampled at random from

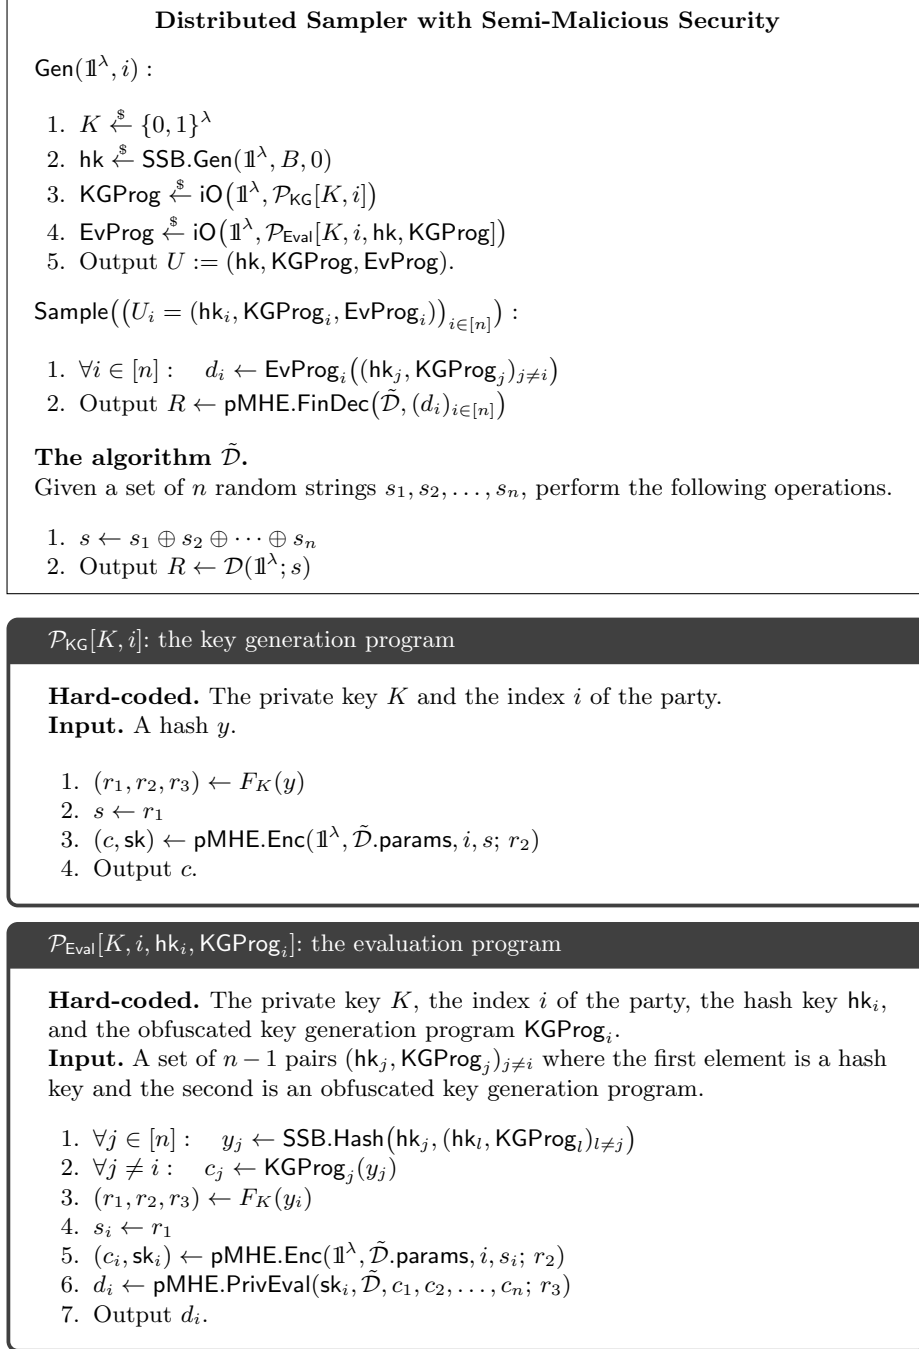


Fig. 6. A Distributed Sampler with Semi-Malicious Security

\mathcal{D} . We choose one honest party h , and throughout hybrids 1 – 6, we modify only how the sampler share U_h is produced. The rest of the honest parties continue to produce their sampler shares as per the Gen protocol in Fig. 6.

Because the simulator has access to all parties' random tapes, it of course knows all their secrets. In the following, we refer to $U_i = (\text{hk}_i, \text{KGProg}_i, \text{EvProg}_i)$ as the sampler share produced by party i for $i \neq h$. We also refer to the secret keys K_i contained in those programs (also known to the simulator). For whatever values $U_h = (\text{hk}_h, \text{KGProg}_h, \text{EvProg}_h)$ the simulator produces on behalf of party h in a given hybrid, we let \hat{s}_i denote the share of the randomness generated by KGProg_i (on the appropriate nonce y), and $(\hat{c}_i, \hat{\text{sk}}_i)$ denote the encryption of that randomness and the corresponding partial decryption key produced by KGProg_i . Finally, we let \hat{r}_2^i denote the random string input in pMHE.Enc by KGProg_i for the generation of \hat{c}_i .

Hybrid 0: This is the initial hybrid, where the simulator, on behalf of every honest party i , generates a sampler share U_i as per the Gen algorithm in Fig. 6.

Hybrid 1: In this hybrid, the simulator, on behalf of honest party h , punctures the key K_h at the relevant point (the hash of $(\text{hk}_j, \text{KGProg}_j)_{j \neq h}$ produced by the other parties j), but programs the appropriate output at that point into the programs. By the correctness of F puncturing, the input-output behaviour of both programs is the same as it was in the previous hybrid. Therefore, by the security of iO , this hybrid is indistinguishable from the previous one.

More specifically, let $(\text{hk}_j, \text{KGProg}_j)_{j \neq h}$ be the hash key and obfuscated key generation program of every other party. (The simulator knows these — even for corrupt parties — since she gets to see the randomness tape of corrupt parties in the definition of semi-malicious security.) The simulator does the following during Gen on behalf of party h , where the text in red indicates what changed since the previous hybrid:

1. $K_h \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\text{hk}_h \xleftarrow{\$} \text{SSB.Gen}(\mathbb{1}^\lambda, B, 0)$
3. $\hat{y}_h \leftarrow \text{SSB.Hash}(\text{hk}_h, (\text{hk}_j, \text{KGProg}_j)_{j \neq h})$
4. $\hat{K}_h \leftarrow \text{Punct}(K_h, \hat{y}_h)$
5. $(\hat{s}_h, \hat{r}_2, \hat{r}_3) \leftarrow F_{K_h}(\hat{y}_h)$
6. $(\hat{c}_h, \hat{\text{sk}}_h) \leftarrow \text{pMHE.Enc}(\mathbb{1}^\lambda, \tilde{\mathcal{D}}.\text{params}, h, \hat{s}_h; \hat{r}_2)$
7. $\text{KGProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{KG}}^1[\hat{K}_h, h, \hat{y}_h, \hat{c}_h])$ (see Fig. 7)
8. $\text{EvProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{Eval}}^1[\hat{K}_h, h, \text{hk}_h, \text{KGProg}_h, \hat{y}_h, \hat{c}_h, \hat{\text{sk}}_h, \hat{r}_3])$ (see Fig. 8)
9. Output $U_h := (\text{hk}_h, \text{KGProg}_h, \text{EvProg}_h)$.

Next, for every l from 0 to the length of the input to SSB.Hash , we proceed first to Hybrid 2.1.1, then to Hybrid 2.1.2.

Hybrid 2.1.1: In this hybrid, the simulator, on behalf of honest party h , makes the hash key hk_h statistically binding at index l (whereas before it was statistically binding at index $l - 1$). This hybrid is indistinguishable from the previous one by the index hiding property of SSB .

Hybrid 2.1.2: In this hybrid, the simulator, on behalf of honest party h , changes the evaluation program EvProg_h to only use the hardcoded key and

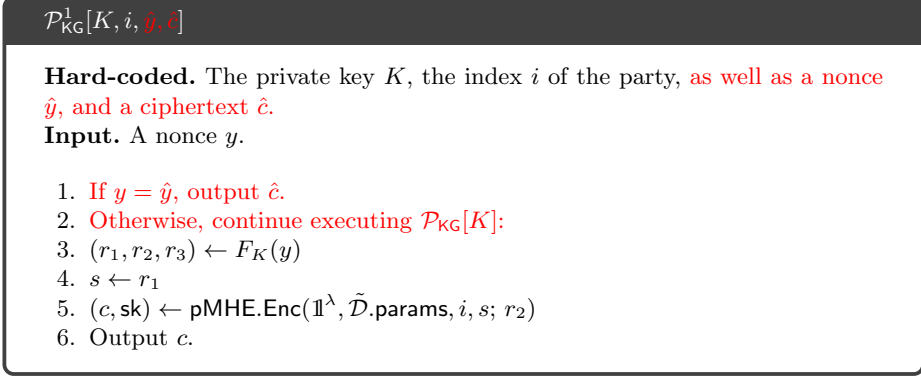


Fig. 7. The Key Generation Program

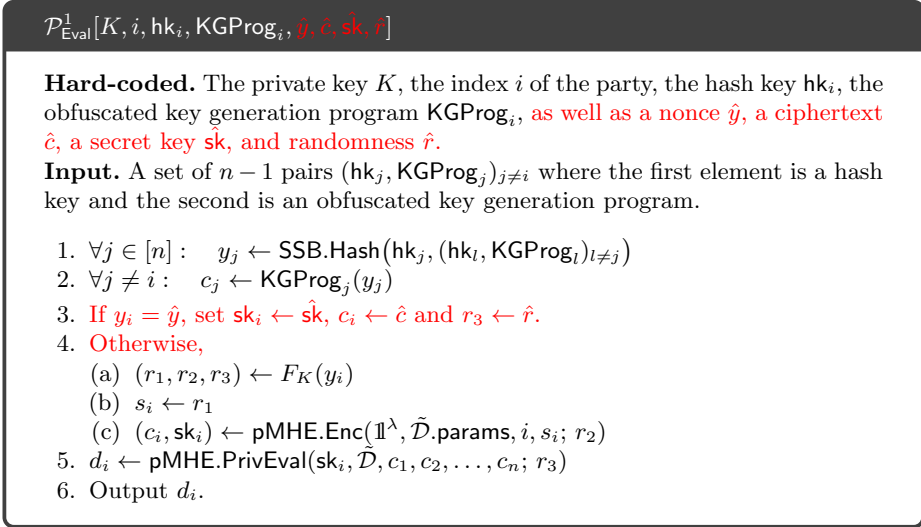


Fig. 8. The Evaluation Program

ciphertext if the first l blocks of the input coincide with a hardcoded reference input. (The simulated party h now obfuscates $\mathcal{P}_{\text{Eval}}^2$ (Fig. 9) instead of $\mathcal{P}_{\text{Eval}}^1$ (Fig. 8).) By the fact that SSB.Hash is statistically binding at l , the input-output behaviour of both programs is the same as it was in the previous hybrid. Therefore, by the security of iO , this hybrid is indistinguishable from the previous one.

More specifically, the simulator does the following during Gen on behalf of party h , where the text in red indicates what changed since the previous hybrid:

1. $K_h \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\text{hk}_h \xleftarrow{\$} \text{SSB.Gen}(\mathbb{1}^\lambda, B, l)$
3. $\hat{y}_h \leftarrow \text{SSB.Hash}(\text{hk}_h, (\text{hk}_j, \text{KGProg}_j)_{j \neq h})$
4. $\hat{K}_h \leftarrow \text{Punct}(K_h, \hat{y}_h)$
5. $(\hat{s}_h, \hat{r}_2, \hat{r}_3) \leftarrow F_{K_h}(\hat{y}_h)$
6. $(\hat{c}_h, \hat{\text{sk}}_h) \leftarrow \text{pMHE.Enc}(\mathbb{1}^\lambda, \tilde{\mathcal{D}}.\text{params}, h, \hat{s}_h; \hat{r}_2)$
7. $\text{KGProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{KG}}^1[\hat{K}_h, h, \hat{y}_h, \hat{c}_h])$ (see Fig. 7)
8. $\hat{w} \leftarrow (\text{hk}_j, \text{KGProg}_j)_{j \neq h}$
9. $\text{EvProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{Eval}}^2[\hat{K}_h, h, \text{hk}_h, \text{KGProg}_h, \hat{y}_h, \hat{c}_h, \hat{\text{sk}}_h, \hat{r}_3, \hat{w}])$ (see Fig. 9)
10. Output $U_h := (\text{hk}_h, \text{KGProg}_h, \text{EvProg}_h)$.

$\mathcal{P}_{\text{Eval}}^2[K, i, \text{hk}_i, \text{KGProg}_i, \hat{y}, \hat{\text{pk}}, \hat{c}, \hat{\text{sk}}, \hat{r}, \hat{w}]$

Hard-coded. The private key K , the index i of the party, the hash key hk_i , the obfuscated key generation program KGProg_i , as well as a nonce \hat{y} , a ciphertext \hat{c} , a secret key $\hat{\text{sk}}$, randomness \hat{r} , and a hardcoded input \hat{w} .

Input. A set of $n - 1$ pairs $(\text{hk}_j, \text{KGProg}_j)_{j \neq i}$ where the first element is a hash key and the second is an obfuscated key generation program.

1. $\forall j \in [n] : y_j \leftarrow \text{SSB.Hash}(\text{hk}_j, (\text{hk}_l, \text{KGProg}_l)_{l \neq j})$
2. $\forall j \neq i : c_j \leftarrow \text{KGProg}_j(y_j)$
3. If $y_i = \hat{y}$ and the first l blocks of \hat{w} and $(\text{hk}_j, \text{KGProg}_j)_{j \neq h}$ coincide, set $\text{sk}_i \leftarrow \hat{\text{sk}}, c_i \leftarrow \hat{c}$ and $r_3 \leftarrow \hat{r}$.
4. Otherwise,
 - (a) $(r_1, r_2, r_3) \leftarrow F_K(y_i)$
 - (b) $s_i \leftarrow r_1$
 - (c) $(c_i, \text{sk}_i) \leftarrow \text{pMHE.Enc}(\mathbb{1}^\lambda, \tilde{\mathcal{D}}.\text{params}, i, s_i; r_2)$
5. $d_i \leftarrow \text{pMHE.PrivEval}(\text{sk}_i, \tilde{\mathcal{D}}, c_1, c_2, \dots, c_n; r_3)$
6. Output d_i .

Fig. 9. The Evaluation Program

Hybrid 3: At this point, EvProg_h only uses the hardcoded key and ciphertext if the *entire* input matches the hardcoded reference input. Since that is the case,

EvProg_h will only ever use the hardcoded decryption key on one ciphertext; so, we can remove the hardcoded decryption key entirely, and instead hardcode a partial decryption value. In this hybrid, the simulator, on behalf of party h , replaces $\mathcal{P}_{\text{Eval}}^2$ (Fig. 9) with $\mathcal{P}_{\text{Eval}}^3$ (Fig. 10) which does exactly this. This hybrid is indistinguishable from the previous one by the security of iO .

More specifically, the simulator does the following during **Gen** on behalf of party h , where the text in red indicates what changed since the previous hybrid:

1. $K_h \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\text{hk}_h \xleftarrow{\$} \text{SSB.Gen}(\mathbb{1}^\lambda, B, B)$
3. $\hat{y}_h \leftarrow \text{SSB.Hash}(\text{hk}_h, (\text{hk}_j, \text{KGProg}_j)_{j \neq h})$
4. $\hat{K}_h \leftarrow \text{Punct}(K_h, \hat{y}_h)$
5. $(\hat{s}_h, \hat{r}_2, \hat{r}_3) \leftarrow F_{K_h}(\hat{y}_h)$
6. $(\hat{c}_h, \hat{\text{sk}}_h) \leftarrow \text{pMHE.Enc}(\mathbb{1}^\lambda, \tilde{\mathcal{D}}.\text{params}, h, \hat{s}_h; \hat{r}_2)$
7. $\text{KGProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{KG}}^1[\hat{K}_h, h, \hat{y}_h, \hat{c}_h])$ (see Fig. 7)
8. $\hat{d}_h \leftarrow \text{pMHE.PrivEval}(\hat{\text{sk}}_h, \tilde{\mathcal{D}}, \hat{c}_1, \hat{c}_2, \dots, \hat{c}_n; \hat{r}_3)$
9. $\hat{w} \leftarrow (\text{hk}_j, \text{KGProg}_j)_{j \neq h}$
10. $\text{EvProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{Eval}}^3[\hat{K}_h, h, \text{hk}_h, \text{KGProg}_h, \hat{w}, \hat{d}_h])$ (see Fig. 10)
11. Output $U_h := (\text{hk}_h, \text{KGProg}_h, \text{EvProg}_h)$.

$\mathcal{P}_{\text{Eval}}^3[K, i, \text{hk}_i, \text{KGProg}_i, \hat{w}, \hat{d}]$

Hard-coded. The private key K , the index i of the party, the hash key hk_i , the obfuscated key generation program KGProg_i , as well as a hardcoded input \hat{w} and a partial decryption \hat{d} .

Input. A set of $n - 1$ pairs $(\text{hk}_j, \text{KGProg}_j)_{j \neq i}$ where the first element is a hash key and the second is an obfuscated key generation program.

1. If $(\text{hk}_j, \text{KGProg}_j)_{j \neq i} = \hat{w}$, output \hat{d} .
2. Otherwise, $\forall j \in [n]: y_j \leftarrow \text{SSB.Hash}(\text{hk}_j, (\text{hk}_l, \text{KGProg}_l)_{l \neq j})$
3. $\forall j \neq i: c_j \leftarrow \text{KGProg}_j(y_j)$
4. $(r_1, r_2, r_3) \leftarrow F_K(y_i)$
5. $s_i \leftarrow r_1$
6. $(c_i, \text{sk}_i) \leftarrow \text{pMHE.Enc}(\mathbb{1}^\lambda, \tilde{\mathcal{D}}.\text{params}, i, s_i; r_2)$
7. $d_i \leftarrow \text{pMHE.PrivEval}(\text{sk}_i, \tilde{\mathcal{D}}, c_1, c_2, \dots, c_n; r_3)$
8. Output d_i .

Fig. 10. The Evaluation Program

Hybrid 4. In this hybrid, the simulator computes the final output R directly as $R \leftarrow \mathcal{D}(\mathbb{1}^\lambda; s)$ where $s = \hat{s}_1 \oplus \hat{s}_2 \oplus \dots \oplus \hat{s}_n$. (The simulator of course has all these values, as it has access to all parties' random tapes.) This hybrid is indistinguishable from the previous one by the correctness of obfuscation and MHE with private evaluation.

Hybrid 5. In this hybrid, the simulator, on behalf of party h , replaces the hardcoded output of F at \hat{y}_h with a truly random one. This hybrid is indistinguishable from the previous one by the security of F .

More specifically, the simulator does the following during **Gen** on behalf of party h , where the text in red indicates what changed since the previous hybrid:

1. $K_h \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\text{hk}_h \xleftarrow{\$} \text{SSB.Gen}(\mathbb{1}^\lambda, B, B)$
3. $\hat{y}_h \leftarrow \text{SSB.Hash}(\text{hk}_h, (\text{hk}_j, \text{KGProg}_j)_{j \neq h})$
4. $\hat{K}_h \leftarrow \text{Punct}(K_h, \hat{y}_h)$
5. **Sample $(\hat{s}_h, \hat{r}_2, \hat{r}_3)$ at random from the appropriate space**
6. $(\hat{c}_h, \hat{\text{sk}}_h) \leftarrow \text{pMHE.Enc}(\mathbb{1}^\lambda, \tilde{\mathcal{D}}.\text{params}, h, \hat{s}_h; \hat{r}_2)$
7. $\text{KGProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{KG}}^1[\hat{K}_h, h, \hat{y}_h, \hat{c}_h])$ (see Fig. 7)
8. $\hat{d}_h \leftarrow \text{pMHE.PrivEval}(\text{sk}_h, \tilde{\mathcal{D}}, \hat{c}_1, \hat{c}_2, \dots, \hat{c}_n; \hat{r}_3)$
9. $\hat{w} \leftarrow (\text{hk}_j, \text{KGProg}_j)_{j \neq h}$
10. $\text{EvProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{Eval}}^3[\hat{K}_h, h, \text{hk}_h, \text{KGProg}_h, \hat{w}, \hat{d}_h])$ (see Fig. 10)
11. Output $U_h := (\text{hk}_h, \text{KGProg}_h, \text{EvProg}_h)$.

Hybrid 6. In this hybrid, the simulator replaces the real ciphertext \hat{c}_h and partial decryption \hat{d}_h with simulated ones. The production of this simulated values does not require party h 's secret decryption key nor the plaintext \hat{s}_h ; it forces the final decryption to output the value R . Since the view of the adversary contains now no information about \hat{s}_h , the simulator can sample R at random from \mathcal{D} . This hybrid is indistinguishable from the previous one by the reusable semi-malicious security of the MHE scheme with private evaluation.

More specifically, the simulator does the following during **Gen** on behalf of party h , where the text in red indicates what changed since the previous hybrid:

1. $K_h \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\text{hk}_h \xleftarrow{\$} \text{SSB.Gen}(\mathbb{1}^\lambda, B, B)$
3. $\hat{y}_h \leftarrow \text{SSB.Hash}(\text{hk}_h, (\text{hk}_j, \text{KGProg}_j)_{j \neq h})$
4. $\hat{K}_h \leftarrow \text{Punct}(K_h, \hat{y}_h)$
5. **$(\tau, \hat{c}_h) \xleftarrow{\$} \text{pMHE.Sim}_1(\mathbb{1}^\lambda, \{h\}, \tilde{\mathcal{D}}.\text{params})$**
6. $\text{KGProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{KG}}^1[\hat{K}_h, h, \hat{y}_h, \hat{c}_h])$ (see Fig. 7)
7. **$R \xleftarrow{\$} \mathcal{D}(\mathbb{1}^\lambda)$**
8. **Retrieve the values $(\hat{s}_j)_{j \neq h}$ and the randomness $(\hat{r}_2^j)_{j \neq h}$ used by the other parties for the generation of $(\hat{c}_j)_{j \neq h}$**
9. **$(\tau, \hat{d}_h) \xleftarrow{\$} \text{pMHE.Sim}_2(\tau, \tilde{\mathcal{D}}, R, (\hat{s}_j, \hat{r}_2^j)_{j \neq h})$**
10. $\hat{w} \leftarrow (\text{hk}_j, \text{KGProg}_j)_{j \neq h}$
11. $\text{EvProg}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{Eval}}^3[\hat{K}_h, h, \text{hk}_h, \text{KGProg}_h, \hat{w}, \hat{d}_h])$ (see Fig. 10)
12. Output $U_h := (\text{hk}_h, \text{KGProg}_h, \text{EvProg}_h)$.

Since this hybrid gives us exactly the distribution we want our original Hybrid 0 to be indistinguishable from, this completes the proof. \square

Observe that a distributed sampler with semi-malicious security also has passive security.

5 Upgrading to Active Security

When moving from semi-malicious to active security, there are two main issues we need to tackle: corrupt parties publishing malformed shares of the sampler, and rushing adversaries. The former can be easily dealt with by adding NIZK proofs of well-formedness to the sampler shares (for this reason, our solution relies on a URS). Rushing adversaries are a more challenging problem, and to deal with this, we rely on a random oracle.

The problem of rushing. In the semi-maliciously secure construction described in Section 4, the randomness used to generate an honest party’s MHE ciphertexts and private keys is output by a PRF, which takes as input a nonce that depends on the key generation programs of all parties (including the corrupt ones). To prove security, we need to puncture the PRF key at that nonce, erasing any correlation between the MHE ciphertext and the PRF key. This can be done in the semi-malicious case, as the simulator knows the programs of the corrupted parties before it must produce those of the honest parties. In the actively secure case, we run into an issue. The adversary is able to adaptively choose the programs of the corrupted parties after seeing those of the other players, in what is called *rushing behaviour*. In the security proof, we would therefore need to puncture a PRF key without knowing the actual position where puncturing is needed.

Although the issue we described above is very specific, dealing with rushing behaviour is a general problem. In a secure distributed sampler, we can program the shares of the honest parties to output an ideal sample when used in conjunction with the shares of the corrupted players. Since the latter are unknown upon generation of the honest players’ shares, the immediate approach would be to program the outputs for every possible choice of the adversary. We run however into an incompressibility problem as we would need to store exponentially many ideal outputs in the polynomial-sized sampler shares.

5.1 Defeating Rushing

In this section, we present a compiler that allows us to deal with rushing behaviour without adding any additional rounds of interaction. This tool handles rushing behaviour not only for distributed samplers, but for a wide range of applications (including our public-key PCF in Section 6). Consider any single-round protocol with no private inputs, where `SendMsg` is the algorithm which party i runs to choose a message to send, and `Output` is an algorithm that determines each party’s output (from party i ’s state and all the messages sent). More concretely, we can describe any such one-round protocol using the following syntax:

`SendMsg`($\mathbb{1}^\lambda, i; r_i$) $\rightarrow \mathbf{g}_i$ generates party i ’s message \mathbf{g}_i , and
`Output`($i, r_i, (\mathbf{g}_j)_{j \in [n]}$) $\rightarrow \mathbf{res}_i$ produces party i ’s output \mathbf{res}_i .

(In the case of distributed samplers, `SendMsg` corresponds to `Gen`, and `Output` corresponds to `Sample`.)

We define modified algorithms (ARMsg, AROutput) such that the associated one-round protocol realizes an ideal functionality that first waits for the corrupted parties' randomness, and then generates the randomness and messages of the honest parties.

This functionality clearly denies the adversary the full power of rushing: the ability to choose corrupt parties' messages based on honest parties' messages. For this reason, we call it the *no-rush* functionality $\mathcal{F}_{\text{NoRush}}$. However, we do allow the adversary a weaker form of rushing behaviour: *selective sampling*. The functionality allows the adversary to re-submit corrupt parties' messages as many times as it wants, and gives the adversary the honest parties' messages in response (while hiding the honest parties' randomness). At the end, the adversary can select which execution she likes the most.

Definition 5.1 (Anti-Rusher). *Let (SendMsg, Output) be a one-round n -party protocol where SendMsg needs $L(\lambda)$ bits of randomness to generate a message. An anti-rusher for SendMsg is a one-round protocol (ARMsg, AROutput) implementing the functionality $\mathcal{F}_{\text{NoRush}}$ (see Fig. 11) for SendMsg against an active adversary.*

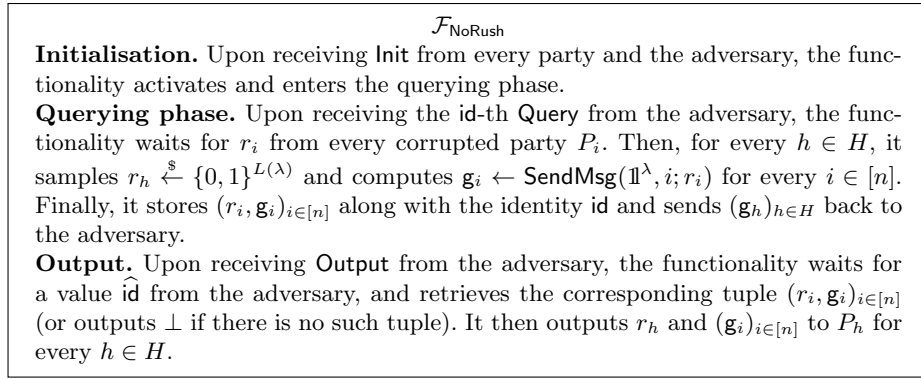


Fig. 11. The Anti-Rushing Functionality $\mathcal{F}_{\text{NoRush}}$

If (SendMsg, Output) = (Gen, Sample) is a distributed sampler with semi-malicious security, applying this transformation gives a distributed sampler with active security.

Intuition Behind our Anti-Rushing Compiler. We define (ARMsg, AROutput) as follows. When called by party i , ARMsg outputs an obfuscated program S_i ; this program takes as input a response of the random oracle, and uses it as a nonce for a PRF F_{K_i} . The program then feeds the resulting pseudorandom string r into SendMsg, and outputs whatever message SendMsg generates.

Our techniques are inspired by the *delayed backdoor programming* technique of Hofheinz *et al.* [HJK⁺16], used for adaptively secure universal samplers.

The trapdoor. In order to prove that our compiler realizes $\mathcal{F}_{\text{NoRush}}$ for `SendMsg`, a simulator must be able to force the compiled protocol to return given outputs of `SendMsg`, even *after* sending messages (outputs of `ARMsg`) on behalf of the honest parties.

Behind its usual innocent behaviour, the program S_i hides a trapdoor that allows it to secretly communicate with the random oracle. S_i owns a key k_i for a special authenticated encryption scheme based on puncturable PRFs. Every time it receives a random oracle response as input, S_i parses it as a ciphertext-nonce pair and tries to decrypt it. If decryption succeeds, S_i outputs the corresponding plaintext; otherwise, it resumes the usual innocent behaviour, and runs `SendMsg`. (The encryption scheme guarantees that the decryption of random strings fails with overwhelming probability; this trapdoor is never used accidentally, but it will play a crucial role in the proof.) Obfuscation conceals how the result has been computed as long as it is indistinguishable from a random `SendMsg` output.

The inputs fed into $(S_i)_{i \in [n]}$ are generated by querying the random oracle with the programs themselves and NIZKs proving their well-formedness. The random oracle response consists of a random nonce v and additional n blocks $(u_i)_{i \in [n]}$, the i -th one of which is addressed to S_i . The input to S_i will be the pair (u_i, v) . When the oracle tries to secretly communicate a message to S_i , u_i will be a ciphertext, whereas v will be the corresponding nonce.

Given a random oracle query, using the simulation-extractability of the NIZKs, the simulator can retrieve the secrets (in particular, the PRF keys) of the corrupted parties. It can then use this information to learn the randomness used to generate the corrupted parties' messages (i.e. their outputs of `SendMsg`). The simulator then needs only to encrypt these messages received from $\mathcal{F}_{\text{NoRush}}$ using $(k_i)_{i \in H}$, and include these ciphertexts in the oracle response.

Formal Description of our Anti-Rushing Compiler. We now formalise the ideas we presented in the previous paragraphs. Our anti-rushing compiler is described in Fig. 13. The unobfuscated program \mathcal{P}_{AR} is available in Fig. 12. We assume that its obfuscation needs $M(\lambda)$ bits of randomness. Observe that \mathcal{P}_{AR} is based on two puncturable PRFs F and F' , the first one of which is used to generate the randomness fed into `SendMsg`.

The second puncturable PRF is part of the authenticated encryption scheme used in the trapdoor. We assume that its outputs are naturally split into $2m$ λ -bit blocks, where $m(\lambda)$ is the size of an output of `SendMsg` (after padding). To encrypt a plaintext $(x^1, \dots, x^m) \in \{0, 1\}^m$ using the key k and nonce $v \in \{0, 1\}^\lambda$, we first expand v using F'_k . The ciphertext consists of m λ -bit blocks, the j -th one of which coincides with the $(2j + x^j)$ -th block output by F' . Decryption is done by reversing these operations. For this reason, we assume that the values $(u_i)_{i \in [n]}$ in the oracle responses are naturally split into m λ -bit chunks. Observe that if the j -th block of the ciphertext is different from both the $2j$ -th and the $(2j + 1)$ -th block output by the PRF, decryption fails.

Finally, let $\text{NIZK} = (\text{Gen}, \text{Prove}, \text{Verify}, \text{Sim}_1, \text{Sim}_2, \text{Extract})$ be a simulation-extractable NIZK for the relation \mathcal{R} describing the well-formedness of the obfuscated programs $(S_i)_{i \in [n]}$. Formally, a statement consists of the pair (S_i, i) , whereas the corresponding witness is the triple containing the PRF keys k_i and K_i hard-coded in S_i and the randomness used for the obfuscation of the latter.

$\mathcal{P}_{\text{AR}}[\text{SendMsg}, k, K, i]$

Hard-coded. The algorithm `SendMsg`, PRF keys k and K and the index i of the party.

Input. Oracle responses $(u, v) \in \{0, 1\}^{\lambda \cdot m(\lambda)} \times \{0, 1\}^\lambda$.

1. $(y_1^0, y_1^1, y_2^0, y_2^1, \dots, y_m^0, y_m^1) \leftarrow F'_k(v)$
2. For every $j \in [m]$ set

$$x^j \leftarrow \begin{cases} 0 & \text{if } y_j^0 = u^j, \\ 1 & \text{if } y_j^1 = u^j, \\ \perp & \text{otherwise.} \end{cases}$$
3. If $x^j \neq \perp$ for every $j \in [m]$, output (x^1, x^2, \dots, x^m) .
4. Set $r \leftarrow F_K(u, v)$.
5. Output $\mathbf{g}_i \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, i; r)$.

Fig. 12. The Anti-Rushing Program

Theorem 5.2. *If $(\text{SendMsg}, \text{Output})$ is a one-round n -party protocol, $\text{NIZK} = (\text{Gen}, \text{Prove}, \text{Verify}, \text{Sim}_1, \text{Sim}_2, \text{Extract})$ is a simulation-extractable NIZK with URS for the relation \mathcal{R} , iO is an indistinguishability obfuscator and (F, Punct) and (F', Punct') are two puncturable PRFs satisfying the properties described above, the protocol $\Pi_{\text{NoRush}} = (\text{ARMsg}, \text{AROutput})$ described in Fig. 13 realizes $\mathcal{F}_{\text{NoRush}}$ for `SendMsg` in the random oracle model with a URS.*

We prove Theorem 5.2 in Appendix A.

Theorem 5.3. *Suppose that $\text{DS} = (\text{Gen}, \text{Sample})$ is a semi-maliciously secure distributed sampler for the distribution \mathcal{D} . Assume that there exists an anti-rusher for DS.Gen . Then, there exists an actively secure distributed sampler for \mathcal{D} .*

On the novelty of this compiler. Observe that the idea of a compiler converting passive protocols into actively secure ones is not new. The most famous example is GMW [GMW87], which achieves this by adding ZK proofs proving the well-formedness of all the messages in the protocol. The novelty of our construction consists of doing this without increasing the number of rounds. GMW deals with rushing by requiring all the parties to commit to their randomness at the beginning of the protocol and then prove that all the messages in the interaction are

Anti-Rushing Compiler Π_{NoRush}

URS. The protocol needs a URS $\text{urs} \xleftarrow{\$} \text{NIZK.Gen}(\mathbb{1}^\lambda)$ for the NIZK proofs.

$\text{ARMsg}(\mathbb{1}^\lambda, i, \text{urs})$:

1. $k_i \xleftarrow{\$} \{0, 1\}^\lambda$
2. $K_i \xleftarrow{\$} \{0, 1\}^\lambda$
3. $w_i \xleftarrow{\$} \{0, 1\}^{M(\lambda)}$
4. $S_i \leftarrow \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}[\text{SendMsg}, k_i, K_i, i]; w_i)$ (see Fig. 12)
5. $\pi_i \xleftarrow{\$} \text{Prove}(\mathbb{1}^\lambda, \text{urs}, (S_i, i), (k_i, K_i, w_i))$
6. Output $\text{armsg}_i := (S_i, \pi_i)$.

$\text{AROutput}(i, (\text{armsg}_j = (S_j, \pi_j))_{j \in [n]}, \text{urs})$:

1. If there exists $j \in [n]$ such that $\text{Verify}(\text{urs}, \pi_j, (S_j, j)) = 0$, output \perp .
2. Query $(S_j, \pi_j)_{j \in [n]}$ to the random oracle \mathcal{H} to get $(v, (u_j)_{j \in [n]})$.
3. $\forall j \in [n]$: $\mathbf{g}_j \leftarrow S_j(u_j, v)$.
4. Output $(\mathbf{g}_j)_{j \in [n]}$ and $F_{K_i}(u_i, v)$.

Fig. 13. Anti-Rushing Compiler

consistent with the initial commitments. A passively secure one-round protocol would therefore be compiled, in the best case, into a 2-round one.

Although the techniques were inspired by [HJK⁺16], this work employs the ideas in a new context, generalising them to multiple players and applying them in multiparty protocols. Observe indeed that [HJK⁺16] devised the techniques to construct adaptively secure universal samplers. To some extent, we still use them to prevent the adversary from making adaptive choices.

6 Public-Key PCFs for Reverse-Samplable Correlations

We now consider the concept of a *distributed correlation sampler*, where the distribution \mathcal{D} produces *private, correlated* outputs R_1, R_2, \dots, R_n , where R_i is given only to the i -th party. This can also model the case where the distribution \mathcal{D} has only one output $R = R_1 = \dots = R_n$, which must be accessible only to the parties that took part in the computation (but not to outsiders; unlike with a distributed sampler).

PCGs and PCFs. The concept of distributed correlation samplers has been previously studied in the form of pseudorandom correlation generators (PCGs) [BCG18, BCG⁺19a, BCG⁺19b, BCG⁺20b] and pseudorandom correlation functions (PCFs) [BCG⁺20a, OSY21]. These are tailored to distributions with n outputs, each one addressed to a different player. Specifically, they consist of two algorithms (Gen, Eval): Gen is used to generate n short correlated seeds or

keys, one for each party. `Eval` is then used to locally expand the keys and non-interactively produce a large amount of correlated randomness, analogously to the non-correlated setting of a PRG (for PCG) or PRF (for PCF).

Both PCGs and PCFs implicitly rely on a trusted dealer for the generation and distribution of the output of `Gen`, which in practice can be realized using a secure multiparty protocol. The communication overhead of this computation should be small, compared with the amount of correlated randomness obtained from `Eval`.

If we consider a one-round protocol to distribute the output of `Gen`, the message of the i -th party and the corresponding randomness r_i act now as a kind of public/private key pair (r_i is necessary to retrieve the i -th output.) Such a primitive is called a *public-key PCF* [OSY21]. Orlandi *et al.* [OSY21] built public-key PCFs for the random OT and vector-OLE correlations based on Paillier encryption with a common reference string (a trusted RSA modulus). In this section, we will build public-key PCFs for general correlations, while avoiding trusted setups.

6.1 Correlation Functions and their Properties

Instead of considering single-output distributions \mathcal{D} , we now consider n -output correlations \mathcal{C} . We also allow different samples from \mathcal{C} to themselves be correlated by some secret parameters, which allows handling correlations such as vector-OLE and authenticated multiplication triples (where each sample depends on some fixed MAC keys). This is modelled by allowing each party i to input a *master secret* mk_i into \mathcal{C} . These additional inputs are independently sampled by each party using an algorithm `Secret`.

Some example correlations. Previous works have focussed on a simple class of *additive correlations*, where the outputs R_1, \dots, R_n form an additive secret sharing of values sampled from a distribution. This captures, for instance, oblivious transfer, (vector) oblivious linear evaluation and (authenticated) multiplication triples, which are all useful correlations for secure computation tasks. Vector OLE and authenticated triples are also examples requiring a master secret, which is used to fix a secret scalar or secret MAC keys used to produce samples. Assuming LWE, we can construct public-key PCFs for any additive correlation [BCG⁺20a], using homomorphic secret-sharing based on multi-key FHE [DHRW16]. However, we do not know how to build PCFs for broader classes of correlations, except for in the two-party setting and relying on subexponentially secure iO [DHRW16].

As motivation, consider the following important types of non-additive correlations:

- *Pseudorandom secret sharing.* This can be seen as a correlation that samples sharings of uniformly random values under some linear secret sharing scheme. Even for simple t -out-of- n threshold schemes such as Shamir, the best previous construction requires $\binom{n}{t}$ complexity [CDI05].

- *Garbled circuits.* In the two-party setting, one can consider a natural garbled circuit correlation, which for some circuit C , gives a garbling of C to one party, and all pairs of input wire labels to the other party. Having such a correlation allows preprocessing for secure 2-PC, where in the online phase, the parties just use oblivious transfer to transmit the appropriate input wire labels.³ Similarly, this can be extended to the multi-party setting, by for instance, giving n parties the garbled circuit together with a secret-sharing of the input wire labels.

For garbled circuits, it may also be useful to consider a variant that uses a master secret, if e.g. we want each garbled circuit to be sampled with a fixed offset used in the free-XOR technique [KS08].

Reverse-Samplable Correlations. The natural way to define a public-key PCF would be a one-round protocol implementing the functionality that samples from the correlation function \mathcal{C} and distributes the outputs. However, Boyle *et al.* [BCG⁺19b] prove that for PCGs, any construction satisfying this definition in the plain model would require that the messages be at least as long as the randomness generated, which negates one of the main advantages of using a PCF. Following the approach of Boyle *et al.*, in this section we adopt a weaker definition. We require that no adversary can distinguish the real samples of the honest parties from simulated ones which are *reverse sampled* based on the outputs of the corrupted players. This choice restricts the set of correlation functions to those whose outputs are efficiently reverse-samplable⁴. We formalise this property below.

Definition 6.1 (Reverse Samplable Correlation Function with Master Secrets). *An n -party correlation function with master secrets is a pair of PPT algorithms $(\text{Secret}, \mathcal{C})$ with the following syntax:*

- Secret takes as input the security parameter $\mathbb{1}^\lambda$ and the index of a party $i \in [n]$. It outputs the i -th party's master correlation secret mk_i .
- \mathcal{C} takes as input the security parameter $\mathbb{1}^\lambda$ and the master secrets $\text{mk}_1, \dots, \text{mk}_n$. It outputs n correlated values R_1, R_2, \dots, R_n , one for each party.

We say that $(\text{Secret}, \mathcal{C})$ is reverse samplable if there exists a PPT algorithm RSample such that, for every set of corrupted parties $C \subsetneq [n]$ and master secrets $(\text{mk}_i)_{i \in [n]}$ and $(\text{mk}'_h)_{h \in H}$ in the image of Secret , no PPT adversary is able to

³ Note that formally, in the presence of malicious adversaries, preprocessing garbled circuits in this way requires the garbling scheme to be adaptively secure [BHR12].

⁴ In the examples above, reverse-samplability is possible for pseudorandom secret-sharing, but not for garbled circuits, since we should not be able to find valid input wire labels when given only a garbled circuit.

distinguish between $\mathcal{C}(\mathbb{1}^\lambda, \text{mk}_1, \text{mk}_2, \dots, \text{mk}_n)$ and

$$\left\{ (R_1, R_2, \dots, R_n) \left| \begin{array}{l} \forall i \in C : \text{mk}'_i \leftarrow \text{mk}_i \\ (R'_1, R'_2, \dots, R'_n) \stackrel{\$}{\leftarrow} \mathcal{C}(\mathbb{1}^\lambda, \text{mk}'_1, \text{mk}'_2, \dots, \text{mk}'_n) \\ \forall i \in C : R_i \leftarrow R'_i \\ (R_h)_{h \in H} \stackrel{\$}{\leftarrow} \text{RSample}(\mathbb{1}^\lambda, C, (R_i)_{i \in C}, (\text{mk}_i)_{i \in C}, (\text{mk}_h)_{h \in H}) \end{array} \right. \right\}$$

Notice that indistinguishability cannot rely on the secrecy of the master secrets $(\text{mk}_i)_{i \in [n]}$ and $(\text{mk}'_h)_{h \in H}$, since the adversary could know their values. Furthermore, RSample does not take as input the same master secrets that were used for the generation of the outputs of the corrupted parties. The fact that indistinguishability holds in spite of this implies that the elements $(R_i)_{i \in C}$ leak no information about the master secrets of the honest players.

6.2 Defining Public Key PCFs

We now formalise the definition of public key PCF as it was sketched at the beginning of the section. We start by specifying the syntax, we will then focus our attention on security, in particular against semi-malicious and active adversaries.

Definition 6.2 (Public-Key PCF with Master Secrets). *A public-key PCF for the n -party correlation function with master secrets $(\text{Secret}, \mathcal{C})$ is a pair of PPT algorithms $(\text{Gen}, \text{Eval})$ with the following syntax:*

- **Gen** takes as input the security parameter $\mathbb{1}^\lambda$ and the index of a party $i \in [n]$, and outputs the PCF key pair $(\text{sk}_i, \text{pk}_i)$ of the i -th party. **Gen** needs $L(\lambda)$ bits of randomness.
- **Eval** takes as input an index $i \in [n]$, n PCF public keys, the i -th PCF private key sk_i and a nonce $x \in \{0, 1\}^{L(\lambda)}$. It outputs a value R_i corresponding to the i -th output of \mathcal{C} .

Every public-key PCF $(\text{Gen}, \text{Eval})$ for \mathcal{C} induces a one-round protocol $\Pi_{\mathcal{C}}$. This is the natural construction in which every party broadcasts pk_i output by **Gen**, and then runs **Eval** on all the parties' messages, its own private key and various nonces.

Definition 6.3 (Semi-Maliciously Secure Public-Key PCF for Reverse Samplable Correlation). *Let $(\text{Secret}, \mathcal{C})$ be an n -party, reverse samplable correlation function with master secrets. A public-key PCF $(\text{Gen}, \text{Eval})$ for $(\text{Secret}, \mathcal{C})$ is semi-maliciously secure if the following properties are satisfied.*

- **Correctness.** *No PPT adversary can win the game $\mathcal{G}_{\text{PCF-Corr}}(\lambda)$ (see Fig. 14) with noticeable advantage.*
- **Security.** *There exists a PPT extractor **Extract** such that for every set of corrupted parties $C \subsetneq [n]$ and corresponding randomness $(\rho_i)_{i \in C}$, no PPT adversary can win the game $\mathcal{G}_{\text{PCF-Sec}}^{C, (\rho_i)_{i \in C}}(\lambda)$ (see Fig. 15) with noticeable advantage.*

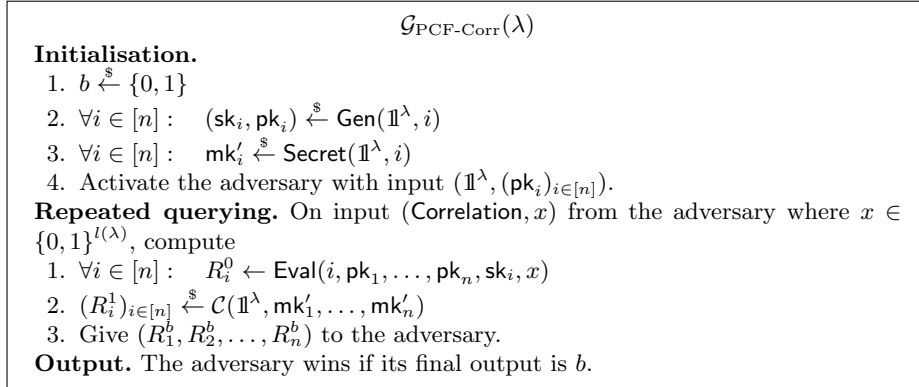


Fig. 14. Correctness Game for the Public-Key PCF

Correctness requires that the samples output by the PCF are indistinguishable from those produced by \mathcal{C} even if the adversary receives all the public keys. Security instead states that a semi-malicious adversary learns no information about the samples and the master secrets of the honest players except what can be deduced from the outputs of the corrupted parties themselves.

Like for distributed samplers, the above definition can be adapted to passive security by modifying the security game. Specifically, it would be sufficient to sample the randomness of the corrupted parties inside the game, perhaps relying on a simulator when $b = 1$.

In our definition, nonces are adaptively chosen by the adversary; however, in a *weak* PCF [BCG⁺20a], the nonces are sampled randomly or selected by the adversary ahead of time. We can define a weak public-key PCF similarly, and use the same techniques as Boyle *et al.* [BCG⁺20a] to convert a weak public-key PCF into a public-key PCF by means of a random oracle.

Active security. We define actively secure public-key PCFs using an ideal functionality, similarly to how we defined actively secure distributed samplers.

Definition 6.4 (Actively Secure Public-Key PCF for Reverse Samplable Correlation). *Let $(\text{Secret}, \mathcal{C})$ be an n -party reverse samplable correlation function with master secrets. A public-key PCF $(\text{Gen}, \text{Eval})$ for $(\text{Secret}, \mathcal{C})$ is actively secure if the corresponding one-round protocol $\Pi_{\mathcal{C}}$ implements the functionality $\mathcal{F}_{\mathcal{C}}^{\text{RSample}}$ (see Fig. 16) against a static and active adversary corrupting up to $n - 1$ parties.*

Any protocol that implements $\mathcal{F}_{\mathcal{C}}^{\text{RSample}}$ will require either a CRS or a random oracle; this is inherent for meaningful correlation functions, since the simulator needs to retrieve the values $(R_i)_{i \in C}$ in order to forward them to $\mathcal{F}_{\mathcal{C}}^{\text{RSample}}$. Therefore, some kind of trapdoor is needed.

$\mathcal{G}_{\text{PCF-Sec}}^{C, (\rho_i)_{i \in C}}(\lambda)$
<p>Initialisation.</p> <ol style="list-style-type: none"> 1. $b \stackrel{\\$}{\leftarrow} \{0, 1\}$ 2. $\forall h \in H : \rho_h \stackrel{\\$}{\leftarrow} \{0, 1\}^{L(\lambda)}$ 3. $\forall i \in [n] : (\text{sk}_i, \text{pk}_i) \leftarrow \text{Gen}(\mathbb{1}^\lambda, i; \rho_i)$ 4. $(\text{mk}_i)_{i \in C} \leftarrow \text{Extract}(C, \rho_1, \rho_2, \dots, \rho_n)$. 5. $\forall h \in H : \text{mk}'_h \stackrel{\\$}{\leftarrow} \text{Secret}(\mathbb{1}^\lambda, h)$ 6. Activate the adversary with $\mathbb{1}^\lambda$ and provide it with $(\text{pk}_i)_{i \in [n]}$ and $(\rho_i)_{i \in C}$. <p>Repeated querying. On input $(\text{Correlation}, x)$ from the adversary where $x \in \{0, 1\}^{l(\lambda)}$, compute</p> <ol style="list-style-type: none"> 1. $\forall i \in [n] : R_i^0 \leftarrow \text{Eval}(i, \text{pk}_1, \dots, \text{pk}_n, \text{sk}_i, x)$ 2. $\forall i \in C : R_i^1 \leftarrow R_i^0$ 3. $(R_h^1)_{h \in H} \stackrel{\\$}{\leftarrow} \text{RSample}(\mathbb{1}^\lambda, C, (R_i^1)_{i \in C}, (\text{mk}_i)_{i \in C}, (\text{mk}'_h)_{h \in H})$ 4. Give $(R_1^b, R_2^b, \dots, R_n^b)$ to the adversary. <p>Output. The adversary wins if its final output is b.</p>

Fig. 15. Security Game for the Public-Key PCF

Notice also that the algorithm `RSample` takes as input the master secrets of the corrupted parties. We can therefore assume that whenever the values $(R_i)_{i \in C}$ chosen by the adversary are inconsistent with $(\text{mk}_i)_{i \in C}$ or with \mathcal{C} itself, the output of the reverse sampler is \perp . As a consequence, an actively secure public-key PCF must not allow the corrupted parties to select these irregular outputs; otherwise distinguishing between real world and ideal world would be trivial.

6.3 Public-Key PCF with Trusted Setup

We will build our semi-maliciously secure public-key PCF by first relying on a trusted setup and then removing it by means of a distributed sampler. A public-key PCF with trusted setup is defined by Def. 6.2 to include an algorithm `Setup` that takes as input the security parameter $\mathbb{1}^\lambda$ and outputs a CRS. The CRS is then provided as an additional input to the evaluation algorithm `Eval`, but not to the generation algorithm `Gen`. (If `Gen` required the CRS, then substituting `Setup` with a distributed sampler would give us a two-round protocol, not a one-round protocol.)

We say that a public-key PCF with trusted setup is semi-maliciously secure if it satisfies Def. 6.3, after minor tweaks to the games $\mathcal{G}_{\text{PCF-Corr}}(\lambda)$ and $\mathcal{G}_{\text{PCF-Sec}}^{C, (\rho_i)_{i \in C}}(\lambda)$ to account for the modified syntax. Notice that in the latter, the extractor needs to be provided with the CRS but not with the randomness used to produce it. If that was not the case, we would not be able to use a distributed sampler to remove the CRS.

Definition of public key PCF with trusted setup. We formalise the concept of public key PCF with trusted setup describing its syntax and security proper-

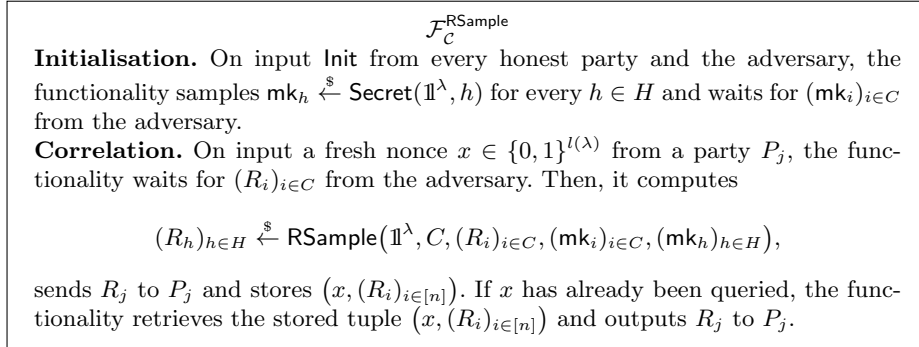


Fig. 16. The Actively Secure Public-Key PCF Functionality for Reverse Samplable Correlation

ties. The definitions closely resemble those of public key PCF (see Def. 6.2 and Def. 6.3). The only difference consists in the CRS generated by the setup algorithm, which is now provided as an additional input to Eval and Extract and whose value is always disclosed to the adversary in the security games.

Definition 6.5 (Public Key PCF with Trusted Setup). *A public key PCF with trusted setup for the n -party correlation function with master secret $(\text{Secret}, \mathcal{C})$ is a triple of PPT algorithms $(\text{Setup}, \text{Gen}, \text{Eval})$ with the following syntax:*

- **Setup** takes as input the security parameter $\mathbb{1}^\lambda$ and outputs a CRS S .
- **Gen** takes as input the security parameter $\mathbb{1}^\lambda$ and the index of a party $i \in [n]$, outputting the PCF key pair $(\text{sk}_i, \text{pk}_i)$ of the i -th party. The algorithm needs $L(\lambda)$ bits of randomness.
- **Eval** takes as input an index $i \in [n]$, a CRS S , n PCF public keys, one for each party, the PCF private key sk_i of the i -th party and a nonce $x \in \{0, 1\}^{l(\lambda)}$. The output is a value R_i corresponding to the i -th output of \mathcal{C} .

Definition 6.6 (Semi-Maliciously Secure Public Key PCF with Trusted Setup for Reverse Samplable Correlation). *Let $(\text{Secret}, \mathcal{C})$ be an n -party, reverse samplable correlation function with master secret. A public key PCF with trusted setup $(\text{Setup}, \text{Gen}, \text{Eval})$ for $(\text{Secret}, \mathcal{C})$ is semi-maliciously secure if the following properties are satisfied.*

- **Correctness.** No PPT adversary can win the game $\mathcal{G}_{\text{SetupCorr}}(\lambda)$ (see Fig. 17) with noticeable advantage.
- **Security.** There exists a PPT extractor Extract such that for every set of corrupted parties $C \subsetneq [n]$ and corresponding randomness $(\rho_i)_{i \in C}$, no PPT adversary can win the game $\mathcal{G}_{\text{SetupSec}}^{C, (\rho_i)_{i \in C}}(\lambda)$ (see Fig. 18) with noticeable advantage.

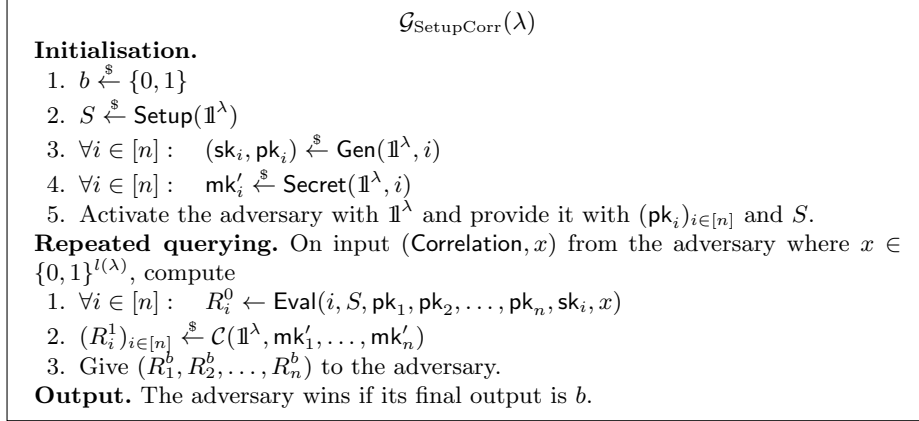


Fig. 17. Correctness Game for Public Key PCFs with Setup

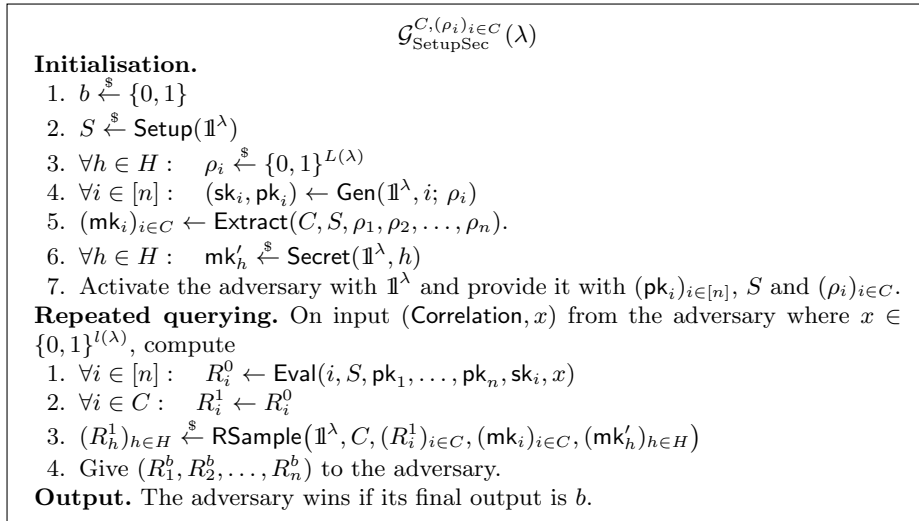


Fig. 18. Security Game for Public Key PCFs with Setup

Our public-key PCF with trusted setup. Our construction is based once again on iO . The key of every party i is a simple PKE pair $(\text{sk}_i, \text{pk}_i)$. The generation of the correlated samples and their distribution is handled by the CRS, which is an obfuscated program. Specifically, the latter takes as input the public keys of the parties and a nonce $x \in \{0, 1\}^{l(\lambda)}$. After generating the master secrets $\text{mk}_1, \text{mk}_2, \dots, \text{mk}_n$ using Secret and the correlated samples R_1, R_2, \dots, R_n using \mathcal{C} , the program protects their privacy by encrypting them under the provided public keys. Specifically, R_i and mk_i are encrypted using pk_i , making the i -th party the only one able to retrieve the underlying plaintext.

The randomness used for the generation of the samples, the master secrets and the encryption is produced by means of two puncturable PRF keys k and K , known to the CRS program. The CRS program is equipped with two keys: k and K . The first one is used to generate the master secrets; the input to the PRF is the sequence of all public keys $(\text{pk}_1, \text{pk}_2, \dots, \text{pk}_n)$. The master secrets remain the same if the nonce x varies. The second PRF key is used to generate the randomness fed into \mathcal{C} and the encryption algorithm; here, the PRF input consists of all the program inputs. As a result, any slight change in the inputs leads to completely unrelated ciphertexts and samples.

On the size of the nonce space. Unfortunately, in order to obtain semi-maliciously security, we need to assume that the nonce space is of polynomial size. In the security proof, we need to change the behaviour of the CRS program for all nonces. This is due to the fact that we cannot rely on the reverse samplability of the correlation function as long as the program contains information about the real samples of the honest players. If the number of nonces is exponential, our security proof would rely on a non-polynomial number of hybrids and therefore we would need to assume the existence of sub-exponentially secure primitives.

The formal description of our solution. Our public-key PCF with trusted setup for $(\text{Secret}, \mathcal{C})$ is described in Fig. 19 together with the program \mathcal{P}_{CG} used as a CRS.

Our solution relies on an IND-CPA PKE scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ and two puncturable PRFs F and F' . We assume that the output of the first one is naturally split into $n + 1$ blocks, the initial one as big as the randomness needed by \mathcal{C} , the remaining ones the same size as the random tape of PKE.Enc . We also assume that the output of F' is split into n blocks as big as the randomness used by Secret .

Theorem 6.7 (Public Key PCFs with Trusted Setup). *Let $(\text{Secret}, \mathcal{C})$ be an n -party, reverse samplable correlation function with master secrets. If $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ is an IND-CPA PKE scheme, iO is an indistinguishability obfuscator, (F, Punct) and (F', Punct') are puncturable PRFs with the properties described above and $l(\lambda)$ is $\text{polylog}(\lambda)$, the construction presented in Fig. 19 is a semi-maliciously secure public-key PCF with trusted setup for $(\text{Secret}, \mathcal{C})$.*

Furthermore, if PKE , iO , (F, Punct) and (F', Punct') are sub-exponentially secure, the public-key PCF with trusted setup is semi-maliciously secure even if $l(\lambda)$ is $\text{poly}(\lambda)$.

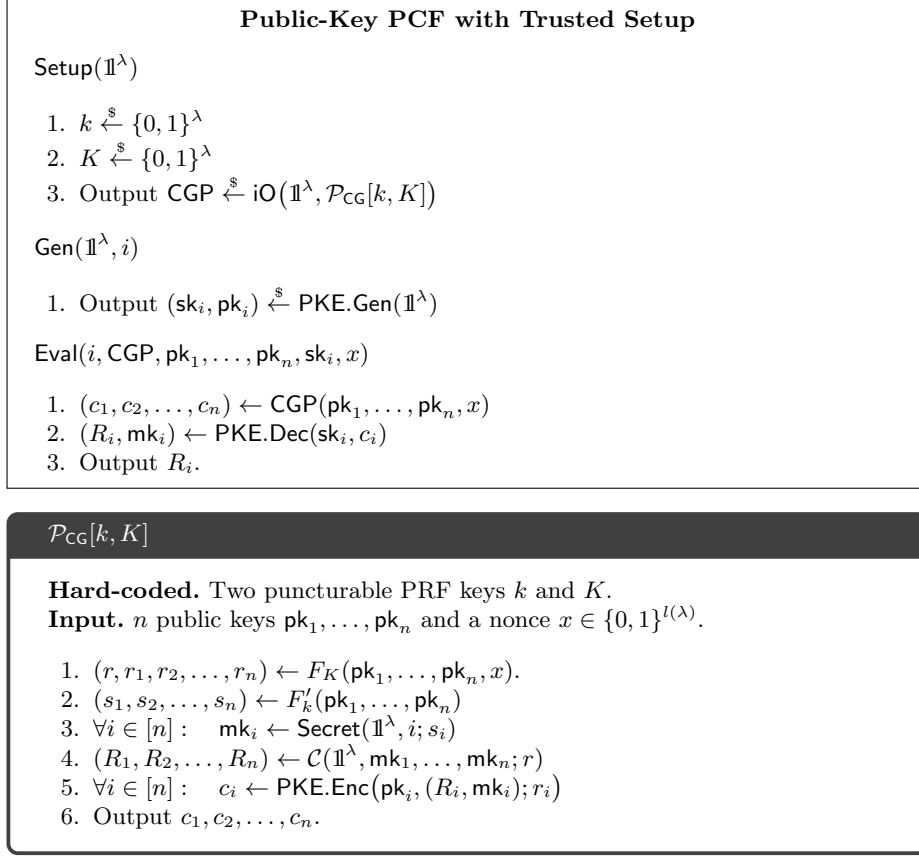


Fig. 19. A Public-Key PCF with Trusted Setup

In both cases, the size of the CRS and the PCF keys is $\text{poly}(l)$.

We prove Theorem 6.7 in Appendix B.

6.4 Our Public-Key PCFs

As mentioned in the previous section, once we obtain a semi-maliciously secure public-key PCF with trusted setup, we can easily remove the CRS using a distributed sampler. We therefore obtain a public-key PCF with security against semi-malicious adversaries. If the size of the CRS and the keys of the initial construction is logarithmic in the size of the nonce space, the key length after removing the setup is still polynomial in $l(\lambda)$.

Theorem 6.8 (Semi-Maliciously Secure Public Key PCFs). *Let $(\text{Secret}, \mathcal{C})$ be an n -party, reverse samplable correlation function with master secrets. Suppose that $\text{pkPCFS} = (\text{Setup}, \text{Gen}, \text{Eval})$ is a semi-maliciously secure public-key*

PCF with trusted setup for $(\text{Secret}, \mathcal{C})$. Moreover, assume that there exists a semi-maliciously secure n -party distributed sampler for pkPCFS.Setup . Then, public-key PCFs for $(\text{Secret}, \mathcal{C})$ with semi-malicious security exist.

We will not prove Theorem 6.8 formally. Security follows from the fact that distributed samplers implement the functionality that samples directly from the underlying distribution. From this point of view, it is fundamental that the randomness input into Setup is not given as input to the extractor of the public-key PCF pkPCFS .

Active security in the random oracle model. If we rely on a random oracle, it is easy to upgrade a semi-maliciously secure public-key PCF to active security. We can use an anti-rusher (see Section 5.1) to deal with rushing and malformed messages. If the key size of the semi-malicious construction is polynomial in $l(\lambda)$, after compiling with the anti-rusher, the key length is still $\text{poly}(l)$. The technique described above allows us to deduce the security of our solution from the semi-malicious security of the initial public-key PCF. The result is formalised by the following theorem. Again, we will not provide a formal proof.

Theorem 6.9 (Actively Secure Public Key PCFs in the Random Oracle Model). *Let $(\text{Secret}, \mathcal{C})$ be an n -party, reverse samplable correlation function with master secret. Assume that $\text{pkPCF} = (\text{Gen}, \text{Eval})$ is a semi-maliciously secure public-key PCFs for $(\text{Secret}, \mathcal{C})$ and suppose there exists an anti-rusher for the associated protocol. Then, actively secure public-key PCFs for $(\text{Secret}, \mathcal{C})$ exist.*

Active security from sub-exponentially secure primitives. So far, all our constructions rely on polynomially secure primitives. However, we often work in the random oracle model. We now show that it is possible to build actively secure public-key PCFs in the URS model assuming the existence of sub-exponentially secure primitives. Furthermore, these constructions come with no restrictions on the size of the nonce space.

Our solution is obtained by assembling a sub-exponentially and semimaliciously secure public-key PCF with trusted setup with a sub-exponentially and semi-maliciously secure distributed sampler. We add witness-extractable NIZKs proving the well-formedness of the messages. Like for our semi-malicious construction, if the size of the CRS and the keys of the public-key PCF with trusted setup is polynomial in the nonce length $l(\lambda)$, after composing with the DS, the key size remains $\text{poly}(l)$.

Theorem 6.10 (Actively Secure Public Key PCFs from Subexponentially Secure Primitives). *Let $(\text{Secret}, \mathcal{C})$ be an n -party, reverse samplable correlation function with master secret. Suppose that $\text{pkPCFS} = (\text{Setup}, \text{Gen}, \text{Eval})$ is a sub-exponentially and semi-maliciously secure public-key PCF with trusted setup for $(\text{Secret}, \mathcal{C})$. Assume that there exists a sub-exponentially and semi-maliciously secure n -party distributed sampler for pkPCFS.Setup . If there exist*

simulation-extractable NIZKs with URS proving the well-formedness of the sampler shares and the PCF public keys, there exists an actively secure public-key PCF for $(\text{Secret}, \mathcal{C})$ in the URS model.

We prove Theorem 6.10 in Appendix C.

7 Ideal Public Key PCFs and Distributed Universal Samplers

We now inspect the feasibility of ideal public key PCFs. The term is used to denote one-round constructions implementing the functionality that directly samples the outputs from adaptively chosen correlations and distributes them to the parties. In contrast with the other PCFs we described, ideal public key PCFs are not tailored to any specific correlation function; instead, the correlation function can be chosen on the fly. However, they can exist only in the random oracle model.

Defining ideal public key PCFs. We deal with generic correlations \mathcal{C} without master secrets. \mathcal{C} takes no inputs, and generates n correlated outputs. It does not need to satisfy any specific properties (in particular it is not required to be reverse samplable). Since the correlations supported by an ideal public key PCF are restricted to those whose description is polynomially bounded, we define the class of (n, ℓ, r, t) -correlations as the set of functions mapping r bits of randomness into n t -bit outputs and having an ℓ -bit description as a circuit.

The syntax of ideal public key PCFs is derived from that of their non-ideal counterparts (see Def. 6.2). The only difference is that the evaluation algorithm Eval takes as input the description of an (n, ℓ, r, t) -correlation instead of a nonce.

Definition 7.1 (Ideal Public-Key PCF). *Let $\ell(\lambda)$, $r(\lambda)$ and $t(\lambda)$ be polynomials. An ideal public-key PCF for (n, ℓ, r, t) -correlations is a pair of PPT algorithms $(\text{Gen}, \text{Eval})$ with the following syntax:*

- Gen takes as input the security parameter 1^λ and the index of a party $i \in [n]$, and outputs the PCF key pair $(\text{sk}_i, \text{pk}_i)$ of the i -th party.
- Eval takes as input an index $i \in [n]$, n PCF public keys, the i -th PCF private key sk_i and the description of an (n, ℓ, r, t) -correlation \mathcal{C} . It outputs a value R_i corresponding to the i -th output of \mathcal{C} .

Definition 7.2 (Ideal Public Key PCF with Active Security). *An ideal public key PCF $(\text{Gen}, \text{Eval})$ for (n, ℓ, r, t) -correlations satisfies active security if the corresponding one-round protocol $\Pi_{\mathcal{C}}$ implements the functionality $\mathcal{F}_{\mathcal{C}}^{\text{ideal}}$ (see Fig. 20) against a static and active adversary corrupting up to $n - 1$ parties.*

Like in the definition of the actively secure distributed sampler (see Def. 3.3), the adversary is allowed to request different samples stored under different labels. Afterwards, the adversary can specify a label of its choice, forcing the functionality to output the associated values to the honest players. We must allow

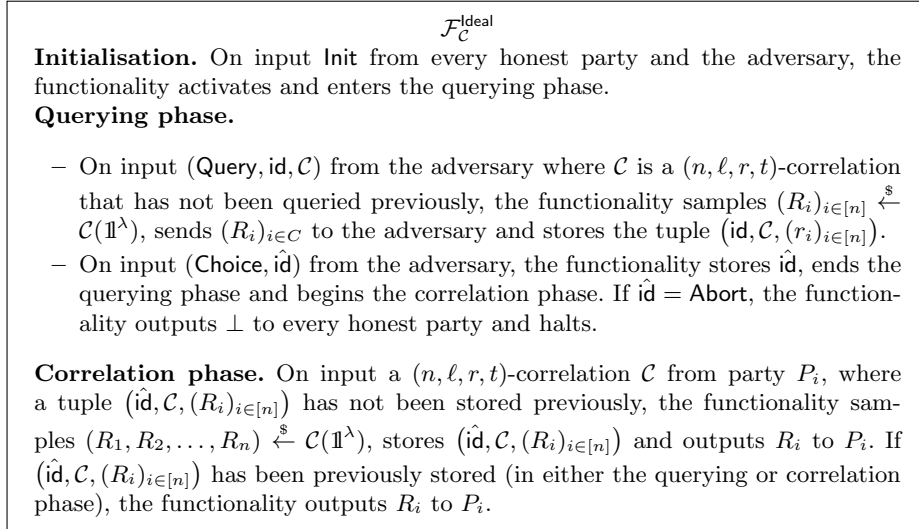


Fig. 20. The Functionality for Ideal Public Key PCFs with Active Security

this kind of influence in order to model rushing; an active adversary can always wait for the messages of the honest parties and adaptively choose the reply of the corrupted players. She is allowed to rerun the procedure as many time as it desire, repeatedly re-generating the messages of the corrupted parties and obtaining different collections of samples. The adversary can then use the messages that led to the most favourable results.

We will build our ideal public key PCFs upon a new primitive called a *distributed universal sampler*. We will present and analyse it in the following subsection.

7.1 Distributed Universal Samplers

A distributed universal sampler (DUS) generalises the concept of distributed sampler. Recall that a distributed sampler is tailored to output a single sample from some fixed distribution \mathcal{D} . In some applications, for instance when we need to sample from multiple distributions, chosen on-the-fly, this may be too restrictive.

What we want instead is analogous to a universal sampler (US) [HJK⁺16], where a trusted dealer first generates and publishes a sampler U . Later, the parties can use U to sample from arbitrary distributions, learning no additional information about the randomness used to generate the output. With a *distributed universal sampler*, we aim to remove the trusted dealer, sampling from generic distributions in a distributed way and with only one round of interaction.

Formally, (distributed) universal samplers do not support completely generic distributions, but are instead restricted to those whose descriptions are polynomially-

bounded. We therefore define the class of (ℓ, r, t) -distributions as the set of all distributions converting r bits of randomness into a t -bit output and having an ℓ -bit description as a circuit.

The syntax of a DUS is obtained by augmenting the `Sample` algorithm of a DS with an additional input, namely the description of the distribution from which to sample the output.

Definition 7.3 (Distributed Universal Sampler). *Let $\ell(\lambda)$, $r(\lambda)$ and $t(\lambda)$ be polynomials. An n -party distributed universal sampler for (ℓ, r, t) -distributions consists of a pair of PPT algorithms $(\text{Gen}, \text{Sample})$ with the following syntax.*

1. `Gen` is a probabilistic algorithm taking as input the security parameter $\mathbb{1}^\lambda$ and a party index $i \in [n]$ and outputting a sampler share U_i for party i . Suppose that the procedure needs $L(\lambda)$ bits of randomness.
2. `Sample` is a deterministic algorithm taking as input n shares of the sampler U_1, U_2, \dots, U_n and the description of an (ℓ, r, t) -distribution \mathcal{D} , outputting a sample R .

Similarly to a DS, any distributed universal sampler $\text{DUS} = (\text{Gen}, \text{Sample})$ naturally corresponds to a one-round protocol Π_{DUS} , where each party first broadcasts a message output by `Gen`, and then, for every required sample, runs `Sample` on input all the parties' messages and the desired distribution \mathcal{D} .

As in the setting of universal samplers [HJK⁺16], we can classify a DUS in two main ways: security for distributions chosen *selectively* by the adversary ahead of time, and security for adversaries who can *adaptively* choose distributions on-the-fly. We refer to the first class as *one-time security*, and the latter as *reusable security*.

One-Time Distributed Universal Samplers While reusable distributed universal samplers need a random oracle independently of the power of the adversary, it is possible to build one-time DUSs with semi-malicious security in the plain model. Indeed, we now consider a one-time, selective security definition, where the sampler may only be queried once, and on a distribution \mathcal{D} that is fixed ahead of time. We formalise the idea.

Definition 7.4 (Distributed Universal Sampler with One-Time Semi-Malicious Security). *A distributed universal sampler $(\text{Gen}, \text{Sample})$ satisfies one-time semi-malicious security if there exists a PPT simulator Sim such that, for every set of corrupted parties $C \subsetneq [n]$, corresponding randomness $(\rho_i)_{i \in C}$ and (ℓ, r, t) -distribution \mathcal{D} , the following two distributions are computationally indistinguishable.*

$$\left\{ \begin{array}{l} (U_i)_{i \in [n]}, (\rho_i)_{i \in C} \\ R \end{array} \middle| \begin{array}{l} \forall i \in H : \rho_i \stackrel{\$}{\leftarrow} \{0, 1\}^{L(\lambda)} \\ \forall i \in [n] : U_i \leftarrow \text{Gen}(\mathbb{1}^\lambda, i; \rho_i) \\ R \leftarrow \text{Sample}(U_1, U_2, \dots, U_n, \mathcal{D}) \end{array} \right\}$$

$$\left\{ \begin{array}{l} (U_i)_{i \in [n]}, (\rho_i)_{i \in C} \\ R \end{array} \middle| \begin{array}{l} R \stackrel{\$}{\leftarrow} \mathcal{D} \\ (U_i)_{i \in [n]} \stackrel{\$}{\leftarrow} \text{Sim}(\mathbb{1}^\lambda, C, \mathcal{D}, R, (\rho_i)_{i \in C}) \end{array} \right\}$$

Just as with a DS, we can adapt the above definition to passive security by sampling the randomness of the corrupted parties inside the game in the real world and by generating it using the simulator in the ideal world. In a definition for active security, we must also account for a rushing adversary, who may adaptively choose the sampler shares of the corrupted parties after seeing those of the honest parties. In other words, in the ideal world, the adversary would be allowed to select the final output R from a list of samples generated by the functionality. We do not consider this notion here, since our actively secure construction actually satisfies the stronger notion of a *reusable* DUS (see Section 7.1).

Universal samplers. We will build our one-time DUSs starting from their non-distributed version [HJK⁺16]. The corresponding definition is available in Section 2.6. A one-time universal sampler is a pair of PPT algorithms (US.Setup, US.Sample) with the same syntax as the reusable case. The main difference is that the random oracle is no longer needed. Security is defined by stating that no PPT adversary can distinguish the real samplers from fake ones specifically programmed to output an ideal sample R when used in conjunction with a distribution \mathcal{D} selected ahead of time. The property is required to hold for every (ℓ, r, t) -distribution \mathcal{D} . In other words, the main novelty is that we now program the sampler for one specific distribution selected ahead of time, whereas in the reusable case, we did that for multiple distributions adaptively chosen by the adversary.

Construction of a one-time DUS. Given distributed samplers and a one-time universal sampler, it is quite straightforward to build a semi-maliciously secure one-time DUS. We can simply substitute the trusted dealer generating the one-time sampler U with a semi-maliciously secure DS for US.Setup. Security follows from the programmability of one-time universal samplers and the fact that, by definition, a DS implements the functionality that directly samples from the underlying distribution. This idea is formalised in the following theorem. We will not, however, provide a formal proof.

Theorem 7.5 (One-Time Distributed Universal Samplers). *Suppose that (US.Setup, US.Sample) is a one-time universal sampler for (ℓ, r, t) -distributions. Assume that (DS.Gen, DS.Sample) is a semi-maliciously secure distributed sampler for US.Setup. Then, there exists a one-time distributed universal samplers for (ℓ, r, t) -distributions with semi-malicious security.*

Reusable Distributed Universal Samplers. In a reusable DUS, the shares output by Gen can be reused to obtain an arbitrary number of samples, from distributions that are chosen adaptively by the adversary. Note that, as is the case for (non-distributed) universal samplers [HJK⁺16], this notion is impossible to realize in the standard model, and our construction will use a random oracle.

We model security by requiring that the sampler can be used to obtain a one-round protocol that securely realizes an ideal sampling functionality, which can be queried adaptively.

Definition 7.6 (Reusable Distributed Universal Sampler). A distributed universal sampler $\text{DUS} = (\text{Gen}, \text{Sample})$ for (ℓ, r, t) -distributions satisfies reusable active security if the corresponding one-round protocol Π_{DUS} implements the functionality \mathcal{F}_{DUS} (see Fig. 21) against a static and active PPT adversary.

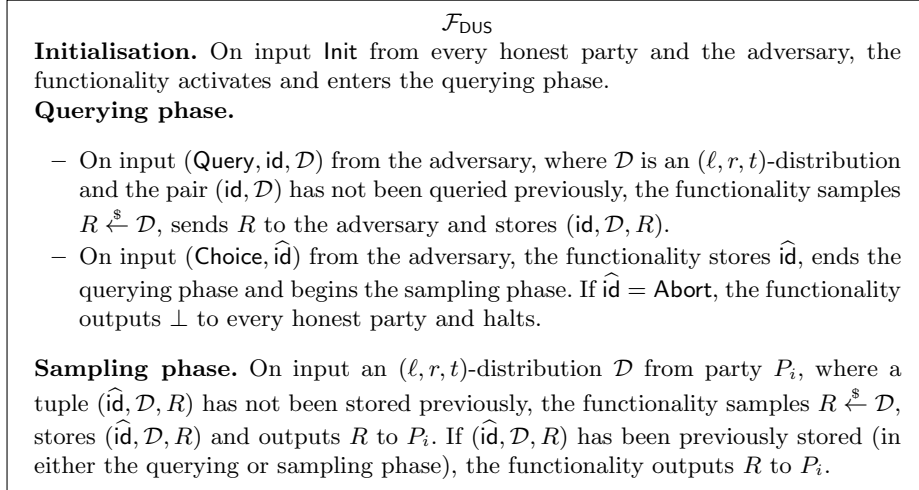


Fig. 21. The Reusable Distributed Universal Sampler Functionality for Active Security

Adaptively secure universal samplers. We construct a reusable DUS starting from an adaptively secure universal sampler. We briefly recall the corresponding definition [HJK⁺16], the formal version of which is available in Section 2.6. An adaptively secure universal sampler is a pair of PPT algorithms (US.Setup , US.Sample), the first one of which is used by a trusted dealer to generate a sampler U . By feeding U and (ℓ, r, t) -distributions \mathcal{D} to the second algorithm (and using the random oracle), anyone can obtain a sample R . Importantly, US.Sample is deterministic, so a set of parties can use U to generate public samples non-interactively. Security requires that no PPT adversary \mathcal{A} can distinguish the real sampler U and the original oracle responses from fake ones specifically programmed to output ideal samples from the (ℓ, r, t) -distributions chosen by \mathcal{A} on-the-fly. Hofheinz *et al.* [HJK⁺16] present an adaptively secure universal sampler for (ℓ, r, t) -distributions whose size is $\text{poly}(\ell, r, t)$.

Our reusable distributed universal sampler. It turns out that designing a reusable DUS is rather straightforward. It suffices to generate the adaptively secure sampler U using a DS for US.Setup . Observe that if we use the construction of [HJK⁺16], the size of the sampler shares is $\text{poly}(\ell, r, t)$. Security follows from the

adaptive programmability of US and the fact that the DS implements the functionality that directly samples from the associated distribution. We formalise this in the theorem below, again, given without proof.

Theorem 7.7 (Reusable Distributed Universal Samplers in the Random Oracle Model). *Suppose that $(\text{US.Setup}, \text{US.Sample})$ is an adaptively secure universal sampler for (ℓ, r, t) -distributions using a random oracle \mathcal{H} . Assume that there exists an actively secure distributed sampler for US.Setup . Then, there exists a reusable distributed universal sampler for (ℓ, r, t) -distributions with active security.*

7.2 Building Ideal Public Key PCFs upon Distributed Universal Samplers

If reusable distributed universal samplers exist, constructing ideal public key PCFs becomes easy. We can use the reusable DUS to produce correlated samples $(R_j)_{j \in [n]}$ and then protect their privacy using a PKE scheme. Specifically, each party P_i can generate a PKE pair $(\text{sk}_i, \text{pk}_i)$ and broadcast the public counterpart along with a reusable DUS share. In order to produce and deal correlated outputs $(R_j)_{j \in [n]}$ from a correlation function \mathcal{C} , it is sufficient to query the DUS with the distribution that samples from \mathcal{C} and outputs the encryption of R_j under pk_j for every $j \in [n]$. In this way, only the party j can retrieve the value of the j -th sample R_j . If we rely on the DUS built on top of the US of Hofheinz *et al.* [HJK⁺16], the key size of our public key PCF is $\text{poly}(\ell, r, n \cdot t)$.

Security. Proving the security of this construction is an easy task. We rely on UC composability, and substitute the DUS with the corresponding functionality \mathcal{F}_{DUS} (see Fig. 21). Then, by the IND-CPA security of the PKE scheme, we can substitute the ciphertexts addressed to the honest parties in the \mathcal{F}_{DUS} responses with the encryption of random values. We substitute the ciphertexts of the corrupt parties with encryptions of the elements provided by the functionality $\mathcal{F}_{\mathcal{C}}^{\text{ideal}}$ (see Fig. 20). Our result is formalised in the following theorem; we do not provide a formal proof.

Theorem 7.8 (Ideal Public Key PCFs in the Random Oracle Model). *If there exists an IND-CPA PKE scheme and an n -party reusable distributed universal sampler with active security, there exists an actively secure ideal public key PCFs for (n, ℓ, r, t) -correlations.*

References

- AJJM20. Prabhanjan Ananth, Abhishek Jain, Zhengzhong Jin, and Giulio Malavolta. Multi-key fully-homomorphic encryption in the plain model. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 28–57. Springer, Heidelberg, November 2020.

- BCCT12. Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In Shafi Goldwasser, editor, *ITCS 2012*, pages 326–349. ACM, January 2012.
- BCG⁺19a. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Peter Rindal, and Peter Scholl. Efficient two-round OT extension and silent non-interactive secure computation. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 291–308. ACM Press, November 2019.
- BCG⁺19b. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 489–518. Springer, Heidelberg, August 2019.
- BCG⁺20a. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Correlated pseudorandom functions from variable-density LPN. In *61st FOCS*, pages 1069–1080. IEEE Computer Society Press, November 2020.
- BCG⁺20b. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators from ring-LPN. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 387–416. Springer, Heidelberg, August 2020.
- BCGI18. Elette Boyle, Geoffroy Couteau, Niv Gilboa, and Yuval Ishai. Compressing vector OLE. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 896–912. ACM Press, October 2018.
- Bd94. Josh Cohen Benaloh and Michael de Mare. One-way accumulators: A decentralized alternative to digital signatures (extended abstract). In Tor Hellesest, editor, *EUROCRYPT'93*, volume 765 of *LNCS*, pages 274–285. Springer, Heidelberg, May 1994.
- BGG19. Sean Rowe, Ariel Gabizon, and Matthew D. Green. A multi-party protocol for constructing the public parameters of the pinocchio zk-SNARK. In Aviv Zohar, Ittay Eyal, Vanessa Teague, Jeremy Clark, Andrea Bracciali, Federico Pintore, and Massimiliano Sala, editors, *FC 2018 Workshops*, volume 10958 of *LNCS*, pages 64–77. Springer, Heidelberg, March 2019.
- BGI⁺01. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2001.
- BGI⁺14a. Amos Beimel, Ariel Gabizon, Yuval Ishai, Eyal Kushilevitz, Sigurd Meldgaard, and Anat Paskin-Cherniavsky. Non-interactive secure multiparty computation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 387–404. Springer, Heidelberg, August 2014.
- BGI14b. Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014.
- BGM17. Sean Rowe, Ariel Gabizon, and Ian Miers. Scalable multi-party computation for zk-SNARK parameters in the random beacon model. Cryptol-

- ogy ePrint Archive, Report 2017/1050, 2017. <https://eprint.iacr.org/2017/1050>.
- BHR12. Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 134–153. Springer, Heidelberg, December 2012.
- BP15. Nir Bitansky and Omer Paneth. ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 401–427. Springer, Heidelberg, March 2015.
- BW13. Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.
- Can01. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
- CDI05. Ronald Cramer, Ivan Damgård, and Yuval Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 342–362. Springer, Heidelberg, February 2005.
- CLTV15. Ran Canetti, Huijia Lin, Stefano Tessaro, and Vinod Vaikuntanathan. Obfuscation of probabilistic circuits and applications. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 468–497. Springer, Heidelberg, March 2015.
- DHRW16. Yevgeniy Dodis, Shai Halevi, Ron D. Rothblum, and Daniel Wichs. Spooky encryption and its applications. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 93–122. Springer, Heidelberg, August 2016.
- FLS90. Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st FOCS*, pages 308–317. IEEE Computer Society Press, October 1990.
- GGH⁺13. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
- GGM86. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1986.
- GMW87. Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.
- GO07. Jens Groth and Rafail Ostrovsky. Cryptography in the multi-string model. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 323–341. Springer, Heidelberg, August 2007.
- GOS06. Jens Groth, Rafail Ostrovsky, and Amit Sahai. Non-interactive zaps and new techniques for NIZK. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 97–111. Springer, Heidelberg, August 2006.
- GPSZ17. Sanjam Garg, Omkant Pandey, Akshayaram Srinivasan, and Mark Zhandry. Breaking the sub-exponential barrier in obfuscation. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017*,

- Part III*, volume 10212 of *LNCS*, pages 156–181. Springer, Heidelberg, April / May 2017.
- HJ⁺17. Shai Halevi, Yuval Ishai, Abhishek Jain, Ilan Komargodski, Amit Sahai, and Eylon Yogev. Non-interactive multiparty computation without correlated randomness. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part III*, volume 10626 of *LNCS*, pages 181–211. Springer, Heidelberg, December 2017.
- HJK⁺16. Dennis Hofheinz, Tibor Jager, Dakshita Khurana, Amit Sahai, Brent Waters, and Mark Zhandry. How to generate and use universal samplers. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 715–744. Springer, Heidelberg, December 2016.
- HW15. Pavel Hubacek and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In Tim Roughgarden, editor, *ITCS 2015*, pages 163–172. ACM, January 2015.
- JLS21. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, page 60–73, New York, NY, USA, 2021. Association for Computing Machinery.
- KPTZ13. Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 669–684. ACM Press, November 2013.
- KS08. Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 486–498. Springer, Heidelberg, July 2008.
- LZ17. Qipeng Liu and Mark Zhandry. Decomposable obfuscation: A framework for building applications of obfuscation from polynomial hardness. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 138–169. Springer, Heidelberg, November 2017.
- OSY21. Claudio Orlandi, Peter Scholl, and Sophia Yakubov. The rise of paillier: Homomorphic secret sharing and public-key silent OT. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 678–708. Springer, Heidelberg, October 2021.
- PS19. Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114. Springer, Heidelberg, August 2019.

A Proof of Theorem 5.2

Proof. The techniques used in this proof are inspired by [HJK⁺16]. We prove the security of the protocol Π_{NoRush} described in Fig. 13 by showing that it implements the functionality $\mathcal{F}_{\text{NoRush}}$ (see Fig. 11) in the UC-model [Can01]. We achieve this plan through a series of hybrids allowing us to transition from the real protocol Π_{NoRush} (Hybrid 0) to the composition of $\mathcal{F}_{\text{NoRush}}$ with a PPT

simulator (Hybrid 14). In all the stages, we assume that the keys k_h and K_h are uniformly sampled in $\{0, 1\}^\lambda$ for every $h \in H$. Moreover, we assume, without loss of generality, that we deal with adversaries that always query the elements $(S_i, \pi_i)_{i \in [n]}$ to the random oracle before broadcasting $(S_i, \pi_i)_{i \in C}$ to the honest parties.

Hybrid 0. This is the initial stage, corresponding to the real world. The simulator generates the URS, the programs of the honest parties $(S_h)_{h \in H}$, the corresponding NIZKs and the final outputs as per the protocol Π_{NoRush} . Moreover, it replies to the random oracle queries of the adversary sampling random strings. If any value is queried multiple times, the simulator takes care to answer always in the same way. Observe that with overwhelming probability, the final outputs are generated without using the trapdoor of the anti-rushing programs.

Hybrid 1. In this hybrid, we substitute the URS and the NIZKs proving the well-formedness of the programs of the honest parties with the outputs of the simulators NIZK.Sim_1 and NIZK.Sim_2 . In this way, we remove any information concerning the keys of the honest parties from $(\pi_h)_{h \in H}$. Hybrid 1 is indistinguishable from Hybrid 0 due to the multi-theorem zero-knowledge of NIZK.

Formally speaking, the simulator generates the URS and the anti-rushing messages $(S_h, \pi_h)_{h \in H}$ as follows (the red text indicates what changed since the last hybrid).

1. $\forall h \in H : S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}[\text{SendMsg}, k_h, K_h, h])$
2. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
3. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

Next, for every q from 1 to the number of random oracle queries issued by the adversary, we proceed from Hybrid $2.q$ to Hybrid $13.q$.

Hybrid $2.q$. In this hybrid, we schedule the q -th oracle response $(\hat{v}, (\hat{u}_i)_{i \in [n]})$ before generating the programs of the honest parties. Furthermore, for every honest party h , we puncture the key K_h in (\hat{u}_h, \hat{v}) and we store it in S_h . We also program the latter to output the appropriate result when (\hat{u}_h, \hat{v}) is input. Observe that with overwhelming probability, such input does not activate the trapdoor in S_h . By the correctness of puncturing, the input-output behaviour of the honest parties' programs is the same as in the previous hybrid. Hence, indistinguishability holds by the security of iO .

The formal steps performed by the simulator in order to generate $(S_h, \pi_h)_{h \in H}$ are described below.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in [n] : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H : \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H : \hat{r}_h \leftarrow F_{K_h}(\hat{u}_h, \hat{v})$
5. $\forall h \in H : \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
6. $\forall h \in H : S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^1[\text{SendMsg}, k_h, \hat{K}_h, h, \hat{u}_h, \hat{v}, \hat{g}_h])$ (see Fig. 22)
7. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
8. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

$\mathcal{P}_{\text{AR}}^1[\text{SendMsg}, k, K, i, \hat{u}, \hat{v}, \hat{g}]$

Hard-coded. The algorithm `SendMsg`, the PRF keys k and K and the index i of the party. Moreover, the scheduled oracle response (\hat{u}, \hat{v}) to the q -th query and the corresponding output \hat{g} .

Input. Oracle responses $(u, v) \in \{0, 1\}^{\lambda \cdot m(\lambda)} \times \{0, 1\}^\lambda$.

1. If $(u, v) = (\hat{u}, \hat{v})$, output \hat{g}
2. $(y_1^0, y_1^1, y_2^0, y_2^1, \dots, y_m^0, y_m^1) \leftarrow F'_k(v)$
3. For every $j \in [m]$ set

$$x^j \leftarrow \begin{cases} 0 & \text{if } y_j^0 = u^j, \\ 1 & \text{if } y_j^1 = u^j, \\ \perp & \text{otherwise.} \end{cases}$$

4. If $x^j \neq \perp$ for every $j \in [m]$, output (x^1, x^2, \dots, x^m) .
5. Set $r \leftarrow F_K(u, v)$.
6. Output $g \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, i; r)$.

Fig. 22. The Anti-Rushing Program

Hybrid 3. q . In this hybrid, on behalf of every honest party h , the simulator generates the element \hat{g}_h hard-coded into S_h using true randomness \hat{r}_h instead of the output of F . Moreover, if the anti-rushing messages $(\text{armsg}_i)_{i \in C}$ of the corrupted parties are valid and correspond to the q -th oracle query, for every $h \in H$, the simulator directly outputs \hat{r}_h to P_h , instead of using F_{K_h} . Observe that this hybrid is indistinguishable from the previous one due to the security of the puncturable PRF F .

The precise procedure used by the simulator to generate $(S_h, \pi_h)_{h \in H}$ is the following.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in [n]: \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H: \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H: \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
5. $\forall h \in H: \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
6. $\forall h \in H: S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^1[\text{SendMsg}, k_h, \hat{K}_h, h, \hat{u}_h, \hat{v}, \hat{g}_h])$ (see Fig. 22)
7. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
8. $\forall h \in H: \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

Hybrid 4. q . For every honest party h , we now puncture the PRF key k_h in \hat{v} . Furthermore, we hard-code into S_h the output of $F'_{k_h}(\hat{v})$ and we use it to compute the result when \hat{v} is input. Since the input-output behaviour of the program remains the same as in the previous hybrid, indistinguishability holds due to the security of iO .

The formal steps performed by the simulator for the generation of the the anti-rushing messages of the honest parties change as follows.

$\mathcal{P}_{\text{AR}}^2[\text{SendMsg}, k, K, i, \hat{u}, \hat{v}, \hat{g}, (\hat{y}_j^b)_{j,b}]$

Hard-coded. The algorithm `SendMsg`, the PRF keys k and K and the index i of the party. Moreover, the scheduled oracle response (\hat{u}, \hat{v}) to the q -th query and the corresponding output \hat{g} . **Finally, the PRF output $(\hat{y}_j^b)_{j,b}$.**

Input. Oracle responses $(u, v) \in \{0, 1\}^{\lambda \cdot m(\lambda)} \times \{0, 1\}^\lambda$.

1. If $(u, v) = (\hat{u}, \hat{v})$, output \hat{g}
2. If $v = \hat{v}$, for every $j \in [m]$, set

$$x^j \leftarrow \begin{cases} 0 & \text{if } \hat{y}_j^0 = u^j, \\ 1 & \text{if } \hat{y}_j^1 = u^j, \\ \perp & \text{otherwise.} \end{cases}$$

3. Otherwise, compute $(y_1^0, y_1^1, y_2^0, y_2^1, \dots, y_m^0, y_m^1) \leftarrow F'_k(v)$ and, for every $j \in [m]$, set

$$x^j \leftarrow \begin{cases} 0 & \text{if } y_j^0 = u^j, \\ 1 & \text{if } y_j^1 = u^j, \\ \perp & \text{otherwise.} \end{cases}$$

4. If $x^j \neq \perp$ for every $j \in [m]$, output (x^1, x^2, \dots, x^m) .
5. Set $r \leftarrow F_K(u, v)$.
6. Output $g \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, i; r)$.

Fig. 23. The Anti-Rushing Program

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in [n] : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H : \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H : \hat{k}_h \leftarrow \text{Punct}'(k_h, \hat{v})$
5. $\forall h \in H : (\hat{y}_{h,j}^b)_{j,b} \leftarrow F'_{k_h}(\hat{v})$
6. $\forall h \in H : \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
7. $\forall h \in H : \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
8. $\forall h \in H : S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^2[\text{SendMsg}, \hat{k}_h, \hat{K}_h, h, \hat{u}_h, \hat{v}, \hat{g}_h, (\hat{y}_{h,j}^b)_{j,b}])$ (see Fig. 23)
9. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
10. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

Hybrid 5.q. In this hybrid, on behalf of every honest party h , the simulator generates the values $(\hat{y}_{h,j}^b)_{j,b}$ sampling them uniformly in $\{0, 1\}^\lambda$ instead of using the PRF F' . This hybrid is indistinguishable from the previous one due to the security of the puncturable PRF F' .

The specific steps performed by the simulator for the generation of $(S_h, \pi_h)_{h \in H}$ are the following.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in [n] : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H : \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H : \hat{k}_h \leftarrow \text{Punct}'(k_h, \hat{v})$
5. $\forall h \in H, j \in [m] \text{ and } b \in \{0, 1\} : \hat{y}_{h,j}^b \xleftarrow{\$} \{0, 1\}^\lambda$
6. $\forall h \in H : \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
7. $\forall h \in H : \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
8. $\forall h \in H : \mathcal{S}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^2[\text{SendMsg}, \hat{k}_h, \hat{K}_h, h, \hat{u}_h, \hat{v}, \hat{g}_h, (\hat{y}_{h,j}^b)_{j,b}])$ (see Fig. 23)
9. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
10. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (\mathcal{S}_h, h))$

Hybrid 6.q. In this hybrid, we rely on an injective double-lengthening PRG $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$. Instead of hard-coding the values $(\hat{y}_{h,j}^b)_{j,b}$ in \mathcal{S}_h , the simulator now stores their images $(\hat{e}_{h,j}^b)_{j,b}$ under the PRG G , i.e. $\hat{e}_{h,j}^b = G(\hat{y}_{h,j}^b)$. Furthermore, when \hat{v} is input in \mathcal{S}_h , the program decodes now x^j by comparing $G(u^j)$ to $\hat{e}_{h,j}^0$ and $\hat{e}_{h,j}^1$. Observe that since G is injective, $u^j = \hat{y}_{h,j}^b$ if and only if $G(u^j) = \hat{e}_{h,j}^b$. In other words, the input-output behaviour of the programs of the honest parties did not change with respect to the previous hybrid. Therefore, indistinguishability holds by the security of iO .

The formal procedure performed by the simulator for the generation of $(\mathcal{S}_h, \pi_h)_{h \in H}$ is the following.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in [n] : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H : \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H : \hat{k}_h \leftarrow \text{Punct}'(k_h, \hat{v})$
5. $\forall h \in H, j \in [m] \text{ and } b \in \{0, 1\} : \hat{y}_{h,j}^b \xleftarrow{\$} \{0, 1\}^\lambda$
6. $\forall h \in H, j \in [m] \text{ and } b \in \{0, 1\} : \hat{e}_{h,j}^b \leftarrow G(\hat{y}_{h,j}^b)$
7. $\forall h \in H : \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
8. $\forall h \in H : \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
9. $\forall h \in H : \mathcal{S}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^3[\text{SendMsg}, \hat{k}_h, \hat{K}_h, h, \hat{u}_h, \hat{v}, \hat{g}_h, (\hat{e}_{h,j}^b)_{j,b}])$ (see Fig. 24)
10. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
11. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (\mathcal{S}_h, h))$

Hybrid 7.q. In this hybrid, instead of generating the values $(\hat{e}_{h,j}^b)_{j,b}$ using the PRG G , the simulator samples them uniformly in $\{0, 1\}^{2\lambda}$ for every $h \in H$. By the security of G , this hybrid is therefore indistinguishable from the previous one. Here, we are actually relying on the fact that, with overwhelming probability, the terms $(\hat{y}_{h,j}^b)_{j,b}$ are not used for the generation the first $q-1$ oracle responses. This is a consequence of the fact that, with overwhelming probability, \hat{v} is different from the nonces in the first $q-1$ oracle answers.

The procedure used by the simulator becomes now the following.

$\mathcal{P}_{\text{AR}}^3[\text{SendMsg}, k, K, i, \hat{u}, \hat{v}, \hat{g}, (\hat{e}_j^b)_{j,b}]$

Hard-coded. The algorithm `SendMsg`, PRF keys k and K and the index i of the party. Moreover, the scheduled oracle response (\hat{u}, \hat{v}) to the q -th query and the corresponding output \hat{g} . **Finally, the PRG outputs $(\hat{e}_j^b)_{j,b}$.**

Input. Oracle responses $(u, v) \in \{0, 1\}^{\lambda \cdot m(\lambda)} \times \{0, 1\}^\lambda$.

1. If $(u, v) = (\hat{u}, \hat{v})$, output \hat{g}
2. If $v = \hat{v}$, for every $j \in [m]$, set

$$x^j \leftarrow \begin{cases} 0 & \text{if } \hat{e}_j^0 = G(u^j), \\ 1 & \text{if } \hat{e}_j^1 = G(u^j), \\ \perp & \text{otherwise.} \end{cases}$$

3. Otherwise, compute $(y_1^0, y_1^1, y_2^0, y_2^1, \dots, y_m^0, y_m^1) \leftarrow F'_k(v)$ and, for every $j \in [m]$, set

$$x^j \leftarrow \begin{cases} 0 & \text{if } y_j^0 = u^j, \\ 1 & \text{if } y_j^1 = u^j, \\ \perp & \text{otherwise.} \end{cases}$$

4. If $x^j \neq \perp$ for every $j \in [m]$, output (x^1, x^2, \dots, x^m) .
5. Set $r \leftarrow F_K(u, v)$.
6. Output $g \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, i; r)$.

Fig. 24. The Anti-Rushing Program

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in [n]: \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H: \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H: \hat{k}_h \leftarrow \text{Punct}'(k_h, \hat{v})$
5. $\forall h \in H, j \in [m] \text{ and } b \in \{0, 1\}: \hat{e}_{h,j}^b \xleftarrow{\$} \{0, 1\}^{2\lambda}$
6. $\forall h \in H: \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
7. $\forall h \in H: \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
8. $\forall h \in H: S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^3[\text{SendMsg}, \hat{k}_h, \hat{K}_h, h, \hat{u}_h, \hat{v}, \hat{g}_h, (\hat{e}_{h,j}^b)_{j,b}])$ (see Fig. 24)
9. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
10. $\forall h \in H: \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

Hybrid 8.q. Starting from this hybrid, the program S_h generates the element \hat{g}_h using the trapdoor mechanism. Compared to the previous stage, the only thing that actually changes is how the values $(\hat{e}_{h,j}^b)_{j,b}$ are generated. Specifically, the simulator first encodes the element \hat{g}_h as a bit string \hat{x}_h . Then, for every $j \in [m]$, it sets $\hat{e}_{h,j}^{\hat{x}_h^j}$ to $G(\hat{u}_h^j)$. The remaining value $\hat{e}_{h,j}^{1-\hat{x}_h^j}$ is instead sampled uniformly in $\{0, 1\}^{2\lambda}$ as in the previous hybrid. Observe that in this way, even

if we remove the first line from $\mathcal{P}_{\text{AR}}^3$, S_h keeps outputting \hat{g}_h when (\hat{u}_h, \hat{v}) is provided as input. We call the program obtained in this way $\mathcal{P}_{\text{AR}}^4$.

It is possible to conclude that $\mathcal{P}_{\text{AR}}^3$ and $\mathcal{P}_{\text{AR}}^4$ have actually the same input-output behaviour. The claim is trivially verifiable when the input (u, v) coincides with the hard-coded pair (\hat{u}, \hat{v}) or $v \neq \hat{v}$. If instead $v = \hat{v}$ and $u \neq \hat{u}$, the matter is a little more complex.

Observe that the image of the PRG G has 2^λ elements and they are embedded into a space of cardinality $2^{2\lambda}$. Since the latter is much larger, with overwhelming probability, values uniformly sampled in $\{0, 1\}^{2\lambda}$ do not belong to the image of G . In particular, this holds for those elements among $(\hat{e}_{h,j}^b)_{j,b}$ that are sampled uniformly. For such values, there is no chance that $\hat{e}_{h,j}^b = G(\hat{u}_h^j)$. As a consequence, the output of $\mathcal{P}_{\text{AR}}^3$ is never generated using the trapdoor when $v = \hat{v}$. In the case of $\mathcal{P}_{\text{AR}}^4$ instead, the only input with $v = \hat{v}$ that activates the trapdoor is (\hat{u}, \hat{v}) due to the injectivity of G .

In this way, we have proven that Hybrid 8.q is indistinguishable from the previous one due to security of iO. The formal description of the steps performed by the simulator is available below.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in [n] : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H : \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H : \hat{k}_h \leftarrow \text{Punct}'(k_h, \hat{v})$
5. $\forall h \in H : \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
6. $\forall h \in H : \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
7. **For every $h \in H$, rewrite \hat{g}_h as an m -bit string \hat{x}_h .**
8. $\forall h \in H$ and $j \in [m] : \hat{e}_{h,j}^{\hat{x}_h^j} \leftarrow G(\hat{u}_h^j)$
9. $\forall h \in H$ and $j \in [m] : \hat{e}_{h,j}^{1-\hat{x}_h^j} \xleftarrow{\$} \{0, 1\}^{2\lambda}$
10. $\forall h \in H : S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^4[\text{SendMsg}, \hat{k}_h, \hat{K}_h, h, \hat{v}, (\hat{e}_{h,j}^b)_{j,b}])$ (see Fig. 25)
11. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
12. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

Hybrid 9.q. In this hybrid, the simulator changes the way it generates those values among $(\hat{e}_{h,j}^b)_{j,b}$ that were previously sampled uniformly in $\{0, 1\}^{2\lambda}$. For each of them, it indeed chooses a random λ -bit seed $\hat{y}_{h,j}^b$ and sets $\hat{e}_{h,j}^b \leftarrow G(\hat{y}_{h,j}^b)$. This hybrid is therefore indistinguishable from the previous one by the PRG security of G . The procedure describing the steps of the simulator is now the following.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in [n] : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H : \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H : \hat{k}_h \leftarrow \text{Punct}'(k_h, \hat{v})$
5. $\forall h \in H : \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$

$\mathcal{P}_{\text{AR}}^4[\text{SendMsg}, k, K, i, \hat{v}, (\hat{e}_j^b)_{j,b}]$

Hard-coded. The algorithm `SendMsg`, the PRF keys k and K and the index i of the party. **Moreover, the scheduled nonce \hat{v} for the q -th oracle query and the values $(\hat{e}_j^b)_{j,b}$.**

Input. Oracle responses $(u, v) \in \{0, 1\}^{\lambda \cdot m(\lambda)} \times \{0, 1\}^\lambda$.

1. If $v = \hat{v}$, for every $j \in [m]$, set

$$x^j \leftarrow \begin{cases} 0 & \text{if } \hat{e}_j^0 = G(u^j), \\ 1 & \text{if } \hat{e}_j^1 = G(u^j), \\ \perp & \text{otherwise.} \end{cases}$$

2. Otherwise, compute $(y_1^0, y_1^1, y_2^0, y_2^1, \dots, y_m^0, y_m^1) \leftarrow F'_k(v)$ and, for every $j \in [m]$, set

$$x^j \leftarrow \begin{cases} 0 & \text{if } y_j^0 = u^j, \\ 1 & \text{if } y_j^1 = u^j, \\ \perp & \text{otherwise.} \end{cases}$$

3. If $x^j \neq \perp$ for every $j \in [m]$, output (x^1, x^2, \dots, x^m) .
4. Set $r \leftarrow F_K(u, v)$.
5. Output $g \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, i; r)$.

Fig. 25. The Anti-Rushing Program

6. $\forall h \in H : \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
7. For every $h \in H$, rewrite \hat{g}_h as an m -bit string \hat{x}_h .
8. $\forall h \in H$ and $j \in [m] : \hat{y}_{h,j}^{\hat{x}_h} \leftarrow \hat{u}_h^j$
9. $\forall h \in H$ and $j \in [m] : \hat{y}_{h,j}^{1-\hat{x}_h} \xleftarrow{\$} \{0, 1\}^\lambda$
10. $\forall h \in H, j \in [m]$ and $b \in \{0, 1\} : \hat{e}_{h,j}^b \leftarrow G(\hat{y}_{h,j}^b)$
11. $\forall h \in H : S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^4[\text{SendMsg}, \hat{k}_h, \hat{K}_h, h, \hat{v}, (\hat{e}_{h,j}^b)_{j,b}])$ (see Fig. 25)
12. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
13. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

Hybrid 10. q . In this hybrid, instead of hard-coding the values $(\hat{e}_{h,j}^b)_{j,b}$ in S_h , the simulator stores their preimages $(\hat{y}_{h,j}^b)_{j,b}$ under the PRG G . Furthermore, when \hat{v} is input into S_h , the program decodes x^j by comparing u^j to $\hat{y}_{h,j}^0$ and $\hat{y}_{h,j}^1$. Observe that since G is injective, $u^j = \hat{y}_{h,j}^b$ if and only if $G(u^j) = \hat{e}_{h,j}^b$. In other words, the input-output behaviour of the programs of the honest parties did not change with respect to the previous hybrid. Therefore, indistinguishability holds by the security of `iO`.

The formal procedure performed by the simulator for the generation of $(S_h, \pi_h)_{h \in H}$ is the following.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$

$\mathcal{P}_{\text{AR}}^5[\text{SendMsg}, k, K, i, \hat{v}, (\hat{y}_j^b)_{j,b}]$

Hard-coded. The algorithm `SendMsg`, the PRF keys k and K and the index i of the party. Moreover, the scheduled nonce \hat{v} for the q -th oracle query and the values $(\hat{y}_j^b)_{j,b}$.

Input. Oracle responses $(u, v) \in \{0, 1\}^{\lambda \cdot m(\lambda)} \times \{0, 1\}^\lambda$.

1. If $v = \hat{v}$, for every $j \in [m]$, set

$$x^j \leftarrow \begin{cases} 0 & \text{if } \hat{y}_j^0 = u^j, \\ 1 & \text{if } \hat{y}_j^1 = u^j, \\ \perp & \text{otherwise.} \end{cases}$$

2. Otherwise, compute $(y_1^0, y_1^1, y_2^0, y_2^1, \dots, y_m^0, y_m^1) \leftarrow F'_k(v)$ and, for every $j \in [m]$, set

$$x^j \leftarrow \begin{cases} 0 & \text{if } y_j^0 = u^j, \\ 1 & \text{if } y_j^1 = u^j, \\ \perp & \text{otherwise.} \end{cases}$$

3. If $x^j \neq \perp$ for every $j \in [m]$, output (x^1, x^2, \dots, x^m) .
4. Set $r \leftarrow F_K(u, v)$.
5. Output $\mathbf{g} \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, i; r)$.

Fig. 26. The Anti-Rushing Program

2. $\forall i \in [n]: \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H: \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
4. $\forall h \in H: \hat{k}_h \leftarrow \text{Punct}'(k_h, \hat{v})$
5. $\forall h \in H: \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
6. $\forall h \in H: \hat{\mathbf{g}}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
7. For every $h \in H$, rewrite $\hat{\mathbf{g}}_h$ as an m -bit string \hat{x}_h .
8. $\forall h \in H$ and $j \in [m]: \hat{y}_{h,j}^{\hat{x}_h^j} \leftarrow \hat{u}_h^j$
9. $\forall h \in H$ and $j \in [m]: \hat{y}_{h,j}^{1-\hat{x}_h^j} \xleftarrow{\$} \{0, 1\}^\lambda$
10. $\forall h \in H: \mathbf{S}_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^5[\text{SendMsg}, \hat{k}_h, \hat{K}_h, h, \hat{v}, (\hat{y}_{h,j}^b)_{j,b}])$ (see Fig. 26)
11. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
12. $\forall h \in H: \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (\mathbf{S}_h, h))$

Hybrid 11. q . In this hybrid, we finally change the oracle response to the q -th query of the adversary, substituting the random values $(\hat{u}_h)_{h \in H}$ with the encryption of the elements $(\hat{\mathbf{g}}_h)_{h \in H}$.

Actually, the only thing that changes is the distribution of the terms $(\hat{y}_{h,j}^b)_{j,b}$. Indeed, they are not uniform in $\{0, 1\}^\lambda$ anymore, but they are generated by the simulator as $F'_{k_h}(\hat{v})$. Since $\hat{y}_{h,j}^{\hat{x}_h^j} = \hat{u}_h^j$ for every $j \in [m]$ and $h \in H$, we also modify the values $(\hat{u}_h)_{h \in H}$ accordingly.

Observe that this hybrid is indistinguishable from the previous one by the security of the puncturable PRF F' . The formal procedure used by the simulator for the generation of $(S_h, \pi_h)_{h \in H}$ and the q -th oracle response becomes the following.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in C : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H : \hat{k}_h \leftarrow \text{Punct}'(k_h, \hat{v})$
4. $\forall h \in H : \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
5. $\forall h \in H : \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
6. For every $h \in H$, rewrite \hat{g}_h as an m -bit string \hat{x}_h .
7. $\forall h \in H : (\hat{y}_{h,j}^b)_{j,b} \leftarrow F'_{k_h}(\hat{v})$
8. $\forall h \in H$ and $j \in [m] : \hat{u}_h^j \leftarrow \hat{y}_{h,j}^{\hat{x}_h^j}$
9. $\forall h \in H : \hat{K}_h \leftarrow \text{Punct}(K_h, (\hat{u}_h, \hat{v}))$
10. $\forall h \in H : S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}^5[\text{SendMsg}, \hat{k}_h, \hat{K}_h, h, \hat{v}, (\hat{y}_{h,j}^b)_{j,b}])$ (see Fig. 26)
11. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
12. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

Hybrid 12.q. In this hybrid, the simulator generates the programs of the honest parties by obfuscating \mathcal{P}_{AR} , as in the original protocol. We however keep replying to the q -th oracle query of the adversary as in Hybrid 11.q. Indistinguishability from the previous stage is guaranteed by the security of iO . As a matter of fact, the programs $\mathcal{P}_{\text{AR}}^5$ (as in Hybrid 11.q) and \mathcal{P}_{AR} have the same input-output behaviour.

Formally, the anti-rushing messages $(S_h, \pi_h)_{h \in H}$ are generated as follows.

1. $\forall h \in H : S_h \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{AR}}[\text{SendMsg}, k_h, K_h, h])$ (see Fig. 12)
2. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
3. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (S_h, h))$

The procedure for the generation of the q -th oracle response is instead the following.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in C : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall h \in H : \hat{r}_h \xleftarrow{\$} \{0, 1\}^{L(\lambda)}$
4. $\forall h \in H : \hat{g}_h \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, h; \hat{r}_h)$
5. For every $h \in H$, rewrite \hat{g}_h as an m -bit string \hat{x}_h .
6. $\forall h \in H : (\hat{y}_{h,j}^b)_{j,b} \leftarrow F'_{k_h}(\hat{v})$
7. $\forall h \in H$ and $j \in [m] : \hat{u}_h^j \leftarrow \hat{y}_{h,j}^{\hat{x}_h^j}$

Observe that the URS and the anti-rushing messages $(\text{armsg}_h)_{h \in H}$ of the honest parties can be now generated independently of the q -th oracle response. It is therefore possible to produce the q -th oracle answer on the fly, after receiving the corresponding query.

Hybrid 13.q. In this hybrid, the simulator generates the elements $(\hat{\mathbf{g}}_h)_{h \in H}$ hidden in the q -th oracle response using the functionality $\mathcal{F}_{\text{NoRush}}$. The operation is actually performed only if the q -th oracle query consists of n elements $(\mathbf{S}_i, \pi_i)_{i \in [n]}$ where the values $(\mathbf{S}_h, \pi_h)_{h \in H}$ coincide with the anti-rushing messages of the honest parties in the only round of interactions and, for every $i \in C$,

$$\text{NIZK.Verify}(\text{urs}, \pi_i, (\mathbf{S}_i, i)) = 1.$$

In such cases, the simulator extracts the witnesses from the NIZK proofs of the corrupted parties using NIZK.Extract , obtaining the corresponding PRF keys k_i and K_i . By the extractability of NIZK, the procedure is successful with overwhelming probability. If the q -th oracle query does not satisfy the properties described above, the response is generated as in the previous stage.

More precisely, the simulator now samples the nonce \hat{v} and the terms $(\hat{u}_i)_{i \in C}$ uniformly and retrieves the randomness \hat{r}_i used for the generation of $\hat{\mathbf{g}}_i = \mathbf{S}_i(\hat{u}_i, \hat{v})$ for every $i \in C$. Now, it is indeed possible to perform the operation as the PRF key K_i is no longer secret. Moreover, with overwhelming probability, the pair (\hat{u}_i, \hat{v}) does not activate the trapdoor in \mathbf{S}_i . The values $(\hat{\mathbf{g}}_h)_{h \in H}$ are finally generated by sending $(\hat{\mathbf{g}}_i, \hat{r}_i)_{i \in C}$ to $\mathcal{F}_{\text{NoRush}}$.

At the end, if the anti-rushing messages of the corrupted parties correspond to the q -th oracle query, the simulator retrieves the randomness of the honest parties $(\hat{r}_h)_{h \in H}$ by sending the label of the q -query to the functionality.

Observe that this hybrid is indistinguishable from the previous one by the witness extractability of NIZK. The formal steps used for the generation of the q -th oracle query become now the following.

1. $\hat{v} \xleftarrow{\$} \{0, 1\}^\lambda$
2. $\forall i \in C : \hat{u}_i \xleftarrow{\$} \{0, 1\}^{m \cdot \lambda}$
3. $\forall i \in C : (k_i, K_i, w_i) \leftarrow \text{NIZK.Extract}(\text{urs}, \tau, (\mathbf{S}_i, i), \pi_i)$
4. $\forall i \in C : \hat{r}_i \leftarrow F_{K_i}(\hat{u}_i, \hat{v})$
5. $\forall i \in C : \hat{\mathbf{g}}_i \leftarrow \text{SendMsg}(\mathbb{1}^\lambda, i; \hat{r}_i)$
6. **Send (Query, $(\hat{\mathbf{g}}_i, \hat{r}_i)_{i \in C}$ to $\mathcal{F}_{\text{NoRush}}$, obtaining $(\text{id}, (\hat{\mathbf{g}}_h)_{h \in H})$ as a reply.**
7. For every $h \in H$, rewrite $\hat{\mathbf{g}}_h$ as an m -bit string \hat{x}_h .
8. $\forall h \in H : (\hat{y}_{h,j}^b)_{j,b} \leftarrow F'_{k_h}(\hat{v})$
9. $\forall h \in H$ and $j \in [m] : \hat{u}_h^j \leftarrow \hat{y}_{h,j}^{\hat{x}_h^j}$

Hybrid 14. This stage corresponds to the ideal world and it just formalises what has been achieved through the series of hybrids we described above.

At this point, the simulator generates the anti-rushing messages of the honest parties and the URS as in Hybrid 1. Moreover, it replies to all the oracle queries as in Hybrid 13.q, receiving ideal $(\mathbf{g}_h)_{h \in H}$ from the functionality $\mathcal{F}_{\text{NoRush}}$. We recall that each of these samples is associated with a label id , so there is a one-to-one correspondence between labels and oracle queries. When the adversary selects the anti-rushing messages $(\mathbf{S}_i, \pi_i)_{i \in C}$ of the corrupted parties, the simulator sends the label of the corresponding oracle query to the functionality. The latter will take care of outputting the associated randomness to the honest parties. \square

B Proof of Theorem 6.7

Proof. We show that the public key PCF with trusted setup described in Fig. 19 is semi-maliciously secure. We prove both correctness and security in one go. As a matter of fact, correctness can be regarded as the special case of security in which $C = \emptyset$. Observe that in such case, we can always assume that `RSample` samples directly from \mathcal{C} using randomly chosen master secrets $(\mathbf{mk}'_i)_{i \in [n]}$.

We proceed by a sequence of 12 indistinguishable hybrids (some of them repeated for every possible nonce value) going from the real world (Hybrid 0) to the ideal world (Hybrid 12). The size of the nonce space will affect the proof only on the number of reductions needed. Specifically, the number of hybrids will be polynomial if and only if the cardinality of the nonce space is polynomial. In the other cases, we will need to assume the existence of sub-exponentially secure primitives.

We always assume that the PRF keys k and K hard-coded into the correlation generation program `CGP` are randomly sampled in $\{0, 1\}^\lambda$. Moreover, in every stage, we assume that the PKE pairs of the parties are all generated according to the protocol, using the randomness parametrising the game in the case of the corrupted players. We denote the i -th pair by $(\hat{\mathbf{sk}}_i, \hat{\mathbf{pk}}_i)$.

Hybrid 0. This is the initial hybrid and corresponds to the execution of the security game with $b = 0$. The challenger creates the correlation generation program by obfuscating \mathcal{P}_{CG} (see Fig. 19). Also the keys of the honest parties are generated following the protocol. Those of the corrupted players are instead derived using the randomness parametrising the game. Finally, the challenger replies to all the sampling queries using `Eval`.

Hybrid 1. In this hybrid, we puncture the PRF key k in the list of public keys of the players. We also store, in the correlation generation program `CGP`, the master secrets corresponding to the punctured position and we use them to compute the output when the public keys of the parties are fed into it. In this way, the input-output behaviour of `CGP` does not change with respect to the previous hybrid. Therefore, indistinguishability follows from the security of obfuscation.

The formal steps performed by the challenger for the generation of `CGP` are now the following (the red text highlights what changed since the last stage).

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\mathbf{pk}}_1, \hat{\mathbf{pk}}_2, \dots, \hat{\mathbf{pk}}_n))$
2. $(\hat{s}_1, \hat{s}_2, \dots, \hat{s}_n) \leftarrow F'_k(\hat{\mathbf{pk}}_1, \hat{\mathbf{pk}}_2, \dots, \hat{\mathbf{pk}}_n)$
3. $\forall i \in [n]: \hat{\mathbf{mk}}_i \leftarrow \text{Secret}(\mathbb{1}^\lambda, i; \hat{s}_i)$
4. `CGP` $\stackrel{s}{\leftarrow}$ `iO` $(\mathbb{1}^\lambda, \mathcal{P}_{\text{CG}}^1[\hat{k}, K, (\hat{\mathbf{pk}}_i, \hat{\mathbf{mk}}_i)_{i \in [n]})]$ (see Fig. 27)

Hybrid 2. In this hybrid, we change how we produce the master secrets $(\hat{\mathbf{mk}}_i)_{i \in [n]}$. Specifically, instead of generating the randomness of `Secret` by means of F'_k , we sample it uniformly. By the security of the puncturable PRF, this hybrid is indistinguishable from the previous one.

Formally, the challenger generates `CGP` as follows.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\mathbf{pk}}_1, \hat{\mathbf{pk}}_2, \dots, \hat{\mathbf{pk}}_n))$

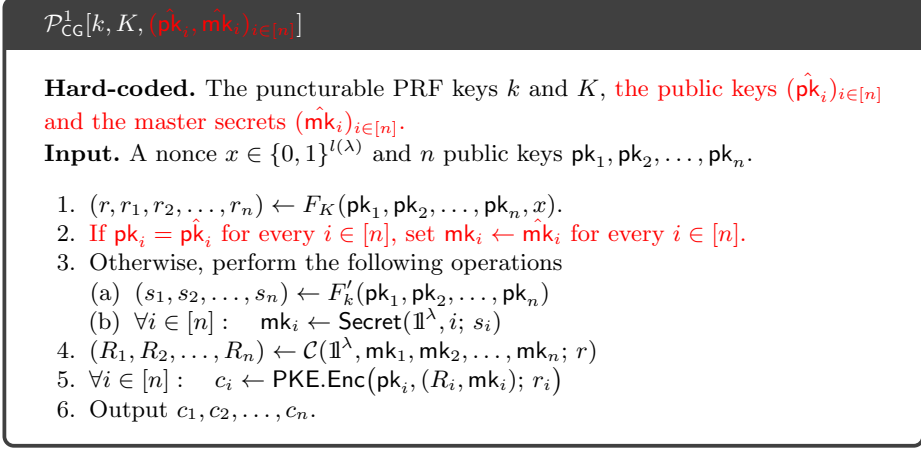


Fig. 27. The Correlation Generation Program

2. $\forall i \in [n]: \hat{\mathbf{mk}}_i \stackrel{\$}{\leftarrow} \text{Secret}(\mathbb{1}^\lambda, i)$
3. $\text{CGP} \stackrel{\$}{\leftarrow} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{CG}^1[\hat{k}, K, (\hat{\mathbf{pk}}_i, \hat{\mathbf{mk}}_i)_{i \in [n]}])$ (see Fig. 27)

We now consider the nonce space $\{0, 1\}^{l(\lambda)}$ and we order it using the lexicographical order $<_{\text{lex}}$. Let ϵ denote the minimum and Ω the maximum. We apply the series of hybrids from 3 to 10 for every nonce \hat{x} , starting from ϵ and following the lexicographical order.

Hybrid 3. \hat{x} . In this hybrid, the challenger samples additional master secrets $(\mathbf{mk}'_h)_{h \in H}$ for the honest parties and hard-codes them into CGP along with \hat{x} and H . When the nonce x input in CGP is strictly smaller than \hat{x} and the provided public keys coincide with the ones of the parties, the program generates the samples $(R_i)_{i \in [n]}$ substituting $\hat{\mathbf{mk}}_h$ with \mathbf{mk}'_h for every $h \in H$. Furthermore, when the nonce x is strictly smaller than \hat{x} , the challenger replies to the correlation queries using RSample , providing it with $(\hat{\mathbf{mk}}_i)_{i \in [n]}$.

Observe that for $\hat{x} = \epsilon$, the input-output behaviour of CGP has not changed with respect to Hybrid 2. Moreover, the challenger never uses RSample to reply to the correlation queries. If instead $\hat{x} \neq \epsilon$, we will see that the input-output behaviour of CGP has not changed with respect to the previous hybrid either and the challenger replies to the correlation queries as it did before. We conclude that indistinguishability holds in both cases due to the security of the obfuscator.

The formal steps used by the challenger for the generation of CGP are now the following.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\mathbf{pk}}_1, \hat{\mathbf{pk}}_2, \dots, \hat{\mathbf{pk}}_n))$
2. $\forall i \in [n]: \hat{\mathbf{mk}}_i \stackrel{\$}{\leftarrow} \text{Secret}(\mathbb{1}^\lambda, i)$
3. $\forall h \in H: \mathbf{mk}'_h \stackrel{\$}{\leftarrow} \text{Secret}(\mathbb{1}^\lambda, h)$
4. $\text{CGP} \stackrel{\$}{\leftarrow} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{CG}^{\hat{x}, 2}[\hat{k}, K, (\hat{\mathbf{pk}}_i, \hat{\mathbf{mk}}_i)_{i \in [n]}, (\mathbf{mk}'_h)_{h \in H}, H, \hat{x}])$ (see Fig. 28)

$$\mathcal{P}_{CG}^{\hat{x},2}[k, K, (\hat{pk}_i, \hat{mk}_i)_{i \in [n]}, (\hat{mk}'_h)_{h \in H}, H, \hat{x}]$$

Hard-coded. The puncturable PRF keys k and K , the public keys $(\hat{pk}_i)_{i \in [n]}$ and the master secrets $(\hat{mk}_i)_{i \in [n]}$ and $(\hat{mk}'_h)_{h \in H}$, the set of honest parties H and the nonce \hat{x} .

Input. A nonce $x \in \{0, 1\}^{l(\lambda)}$ and n public keys pk_1, pk_2, \dots, pk_n .

1. $(r, r_1, r_2, \dots, r_n) \leftarrow F_K(pk_1, pk_2, \dots, pk_n, x)$.
2. If $pk_i = \hat{pk}_i$ for every $i \in [n]$ and $x \geq_{\text{lex}} \hat{x}$, set $mk_i \leftarrow \hat{mk}_i$ for every $i \in [n]$.
3. If $pk_i = \hat{pk}_i$ for every $i \in [n]$ and $x <_{\text{lex}} \hat{x}$, set $mk_h \leftarrow \hat{mk}'_h$ for every $h \in H$ and $mk_i \leftarrow \hat{mk}_i$ for every $i \in C$.
4. Otherwise, perform the following operations
 - (a) $(s_1, s_2, \dots, s_n) \leftarrow F'_k(pk_1, pk_2, \dots, pk_n)$
 - (b) $\forall i \in [n]: mk_i \leftarrow \text{Secret}(\mathbb{1}^\lambda, i; s_i)$
5. $(R_1, R_2, \dots, R_n) \leftarrow \mathcal{C}(\mathbb{1}^\lambda, mk_1, mk_2, \dots, mk_n; r)$
6. $\forall i \in [n]: c_i \leftarrow \text{PKE.Enc}(pk_i, (R_i, mk_i); r_i)$
7. Output c_1, c_2, \dots, c_n .

Fig. 28. The Correlation Generation Program

The reply to $(\text{Correlation}, x)$ when $x <_{\text{lex}} \hat{x}$ is instead computed as follows.

1. $(c_i)_{i \in [n]} \leftarrow \text{CGP}(\hat{pk}_1, \hat{pk}_2, \dots, \hat{pk}_n, x)$
2. $\forall i \in C: (R_i, mk_i) \leftarrow \text{PKE.Dec}(\hat{sk}_i, c_i)$
3. $(R_h)_{h \in H} \stackrel{\$}{\leftarrow} \text{RSample}(\mathbb{1}^\lambda, C, (R_i)_{i \in C}, (\hat{mk}_i)_{i \in C}, (\hat{mk}'_h)_{h \in H})$

Hybrid 4. \hat{x} . In this hybrid, we puncture the PRF key K in the tuple consisting of the public keys of the parties and the nonce \hat{x} , i.e. $(\hat{pk}_1, \hat{pk}_2, \dots, \hat{pk}_n, \hat{x})$. Moreover, we program CGP to output the appropriate ciphertexts when the punctured position is given as input. Since the input-output behaviour of the program has not changed with respect to the previous hybrid, indistinguishability follows from the security of iO.

The formal procedure used by the challenger for the generation of CGP is now the following.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{pk}_1, \hat{pk}_2, \dots, \hat{pk}_n))$
2. $\forall i \in [n]: \hat{mk}_i \stackrel{\$}{\leftarrow} \text{Secret}(\mathbb{1}^\lambda, i)$
3. $\forall h \in H: \hat{mk}'_h \stackrel{\$}{\leftarrow} \text{Secret}(\mathbb{1}^\lambda, h)$
4. $\hat{K} \leftarrow \text{Punct}(K, (\hat{pk}_1, \hat{pk}_2, \dots, \hat{pk}_n, \hat{x}))$
5. $(\hat{r}, \hat{r}_1, \hat{r}_2, \dots, \hat{r}_n) \leftarrow F_K(\hat{pk}_1, \hat{pk}_2, \dots, \hat{pk}_n, \hat{x})$
6. $(\hat{R}_i)_{i \in [n]} \leftarrow \mathcal{C}(\mathbb{1}^\lambda, \hat{mk}_1, \hat{mk}_2, \dots, \hat{mk}_n; \hat{r})$
7. $\forall i \in [n]: \hat{c}_i \leftarrow \text{PKE.Enc}(\hat{pk}_i, (\hat{R}_i, \hat{mk}_i); \hat{r}_i)$
8. $\text{CGP} \stackrel{\$}{\leftarrow} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{CG}^{\hat{x},3}[\hat{k}, \hat{K}, (\hat{pk}_i, \hat{mk}_i)_{i \in [n]}, (\hat{mk}'_h)_{h \in H}, H, \hat{x}, (\hat{c}_i)_{i \in [n]}])$ (see Fig. 29)

Hybrid 5. \hat{x} . In this hybrid, we change how we generate the samples $(\hat{R}_i)_{i \in [n]}$ and encrypt them. Specifically, instead of producing the randomness using the

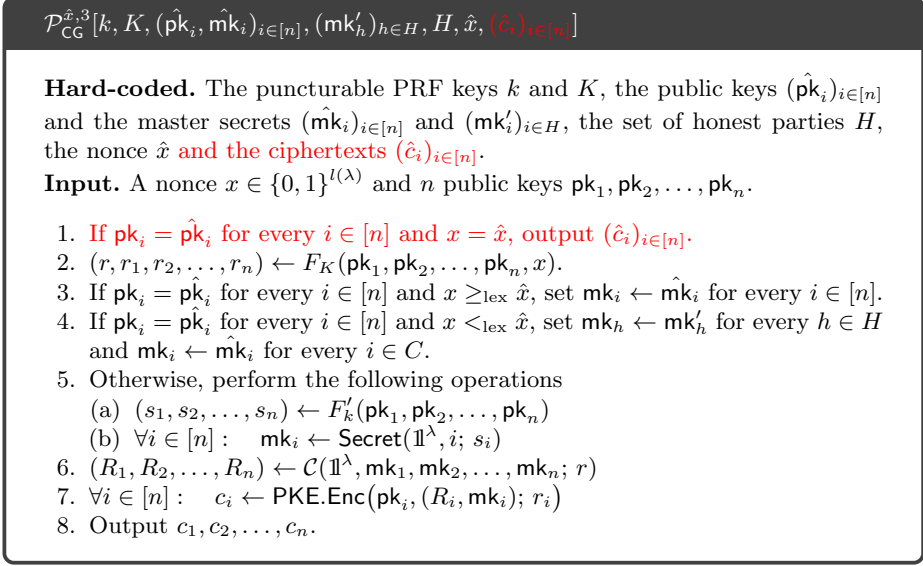


Fig. 29. The Correlation Generation Program

PRF F_K , we sample it uniformly. Observe that this hybrid and the previous one are indistinguishable by the security of the puncturable PRF F .

The formal procedure used by the challenger to generate CGP becomes the following.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\mathbf{pk}}_1, \hat{\mathbf{pk}}_2, \dots, \hat{\mathbf{pk}}_n))$
2. $\forall i \in [n]: \hat{\mathbf{mk}}_i \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, i)$
3. $\forall h \in H: \mathbf{mk}'_h \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, h)$
4. $\hat{K} \leftarrow \text{Punct}(K, (\hat{\mathbf{pk}}_1, \hat{\mathbf{pk}}_2, \dots, \hat{\mathbf{pk}}_n, \hat{x}))$
5. $(\hat{R}_i)_{i \in [n]} \xleftarrow{\$} \mathcal{C}(\mathbb{1}^\lambda, \hat{\mathbf{mk}}_1, \hat{\mathbf{mk}}_2, \dots, \hat{\mathbf{mk}}_n)$
6. $\forall i \in [n]: \hat{c}_i \xleftarrow{\$} \text{PKE.Enc}(\hat{\mathbf{pk}}_i, (\hat{R}_i, \hat{\mathbf{mk}}_i))$
7. $\text{CGP} \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{CG}^{\hat{x},3}[\hat{k}, \hat{K}, (\hat{\mathbf{pk}}_i, \hat{\mathbf{mk}}_i)_{i \in [n]}, (\mathbf{mk}'_h)_{h \in H}, H, \hat{x}, (\hat{c}_i)_{i \in [n]}])$ (see Fig. 29)

Hybrid 6. \hat{x} . In this hybrid, the challenger replies to the query $(\text{Correlation}, \hat{x})$ directly sending the values $(\hat{R}_i)_{i \in [n]}$ sampled by \mathcal{C} during the generation of PCG. This hybrid is indistinguishable from the previous one by the correctness of iO and PKE .

Hybrid 7. \hat{x} . In this hybrid, we change how we produce the samples $(\hat{R}_i)_{i \in [n]}$. Specifically, we first generate $(R'_i)_{i \in [n]}$ using the correlation function \mathcal{C} and substituting \mathbf{mk}'_h to $\hat{\mathbf{mk}}_h$ for every $h \in H$. Then, we obtain $(\hat{R}_h)_{h \in H}$ by feeding the original master secrets $(\hat{\mathbf{mk}}_i)_{i \in [n]}$ and $(R'_i)_{i \in C}$ into RSample . Finally, we set \hat{R}_i to R'_i for every $i \in C$. Observe that this hybrid is indistinguishable from the previous one by the reverse samplability of $(\text{Secret}, \mathcal{C})$.

The formal steps performed by the challenger for the generation of CGP become the following.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n))$
2. $\hat{K} \leftarrow \text{Punct}(K, (\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n, \hat{x}))$
3. $\forall i \in [n]: \hat{\text{mk}}_i \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, i)$
4. $\forall h \in H: \text{mk}'_h \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, h)$
5. $\forall i \in C: \text{mk}'_i \leftarrow \hat{\text{mk}}_i$
6. $(R'_i)_{i \in [n]} \xleftarrow{\$} \mathcal{C}(\mathbb{1}^\lambda, \text{mk}'_1, \text{mk}'_2, \dots, \text{mk}'_n)$
7. $(\hat{R}_h)_{h \in H} \xleftarrow{\$} \text{RSample}(\mathbb{1}^\lambda, C, (R'_i)_{i \in C}, (\hat{\text{mk}}_i)_{i \in C}, (\hat{\text{mk}}_h)_{h \in H})$
8. $\forall i \in C: \hat{R}_i \leftarrow R'_i$
9. $\forall i \in [n]: \hat{c}_i \xleftarrow{\$} \text{PKE.Enc}(\hat{\text{pk}}_i, (\hat{R}_i, \hat{\text{mk}}_i))$
10. $\text{CGP} \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{CG}}^{\hat{x};3}[\hat{k}, \hat{K}, (\hat{\text{pk}}_i, \hat{\text{mk}}_i)_{i \in [n]}, (\text{mk}'_h)_{h \in H}, H, \hat{x}, (\hat{c}_i)_{i \in [n]}])$ (see Fig. 29)

Hybrid 8. \hat{x} . In this stage, instead of hard-coding into CGP the encryption of $(\hat{R}_h, \hat{\text{mk}}_h)$, for every $h \in H$, we store the encryption of (R'_h, mk'_h) . The challenger, however, replies to the query $(\text{Correlation}, \hat{x})$ by feeding the actual samples of the corrupted parties and the original master secrets $(\hat{\text{mk}}_i)_{i \in [n]}$ into RSample . Observe that the challenger does not need to know the secret-keys of the honest parties to reply to the queries $(\text{Correlation}, x)$ with $x \neq \hat{x}$. The knowledge of the keys K and k permits indeed to recompute the samples $(R_i)_{i \in [n]}$. This fact allows us to reduce the indistinguishability between Hybrid 7 and 8 to the IND-CPA security of PKE.

The formal steps performed by the challenger for the generation of CGP become the following.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n))$
2. $\hat{K} \leftarrow \text{Punct}(K, (\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n, \hat{x}))$
3. $\forall i \in [n]: \hat{\text{mk}}_i \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, i)$
4. $\forall h \in H: \text{mk}'_h \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, h)$
5. $\forall i \in C: \text{mk}'_i \leftarrow \hat{\text{mk}}_i$
6. $(R'_i)_{i \in [n]} \xleftarrow{\$} \mathcal{C}(\mathbb{1}^\lambda, \text{mk}'_1, \text{mk}'_2, \dots, \text{mk}'_n)$
7. $\forall i \in [n]: \hat{c}_i \xleftarrow{\$} \text{PKE.Enc}(\hat{\text{pk}}_i, (R'_i, \text{mk}'_i))$
8. $\text{CGP} \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{CG}}^{\hat{x};3}[\hat{k}, \hat{K}, (\hat{\text{pk}}_i, \hat{\text{mk}}_i)_{i \in [n]}, (\text{mk}'_h)_{h \in H}, H, \hat{x}, (\hat{c}_i)_{i \in [n]}])$ (see Fig. 29)

Furthermore, the response to $(\text{Correlation}, \hat{x})$ is now computed as follows.

1. $(c_i)_{i \in [n]} \leftarrow \text{CGP}(\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n, \hat{x})$
2. $\forall i \in C: (R_i, \text{mk}_i) \leftarrow \text{PKE.Dec}(\hat{\text{sk}}_i, c_i)$
3. $(R_h)_{h \in H} \xleftarrow{\$} \text{RSample}(\mathbb{1}^\lambda, C, (R_i)_{i \in C}, (\hat{\text{mk}}_i)_{i \in C}, (\hat{\text{mk}}_h)_{h \in H})$

Hybrid 9. \hat{x} . In this hybrid, we change how we generate the samples $(R'_i)_{i \in [n]}$ and we encrypt them. Specifically, instead of sampling the randomness for \mathcal{C} and PKE.Enc uniformly, we rely on the output of the PRF F_K when its nonce is

the position where we had previously punctured. Indistinguishability from the previous stage holds by the security of the puncturable PRF F .

The formal steps performed by the challenger for the generation of CGP become the following.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n))$
2. $\hat{K} \leftarrow \text{Punct}(K, (\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n, \hat{x}))$
3. $(\hat{r}, \hat{r}_1, \hat{r}_2, \dots, \hat{r}_n) \leftarrow F_K(\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n, \hat{x})$
4. $\forall i \in [n]: \hat{\text{mk}}_i \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, i)$
5. $\forall h \in H: \text{mk}'_h \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, h)$
6. $\forall i \in C: \text{mk}'_i \leftarrow \hat{\text{mk}}_i$
7. $(R'_i)_{i \in [n]} \leftarrow \mathcal{C}(\mathbb{1}^\lambda, \text{mk}'_1, \text{mk}'_2, \dots, \text{mk}'_n; \hat{r})$
8. $\forall i \in [n]: \hat{c}_i \leftarrow \text{PKE.Enc}(\hat{\text{pk}}_i, (R'_i, \text{mk}'_i); \hat{r}_i)$
9. $\text{CGP} \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{CG}}^{\hat{x}, 3}[\hat{k}, \hat{K}, (\hat{\text{pk}}_i, \hat{\text{mk}}_i)_{i \in [n]}, (\text{mk}'_h)_{h \in H}, H, \hat{x}, (\hat{c}_i)_{i \in [n]}])$ (see Fig. 29)

Hybrid 10. \hat{x} . In this hybrid, we do not puncture K anymore and we remove $(\hat{c}_i)_{i \in [n]}$ from CGP. When the public keys of the parties and the nonce \hat{x} are input into the program, we compute the output running the same procedure as if the nonce was strictly smaller than \hat{x} . Observe that the input-output behaviour of the program is the same as in the previous hybrid, therefore indistinguishability follows from the security of iO .

The formal procedure adopted by the challenger for the generation of CGP becomes the following.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n))$
2. $\forall i \in [n]: \hat{\text{mk}}_i \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, i)$
3. $\forall h \in H: \text{mk}'_h \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, h)$
4. $\text{CGP} \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{\text{CG}}^{\hat{x}, 4}[\hat{k}, K, (\hat{\text{pk}}_i, \hat{\text{mk}}_i)_{i \in [n]}, (\text{mk}'_h)_{h \in H}, H, \hat{x}])$ (see Fig. 30)

The next step of the proof is to repeat Hybrid 3-10 for the following value of the nonce \hat{x} . When the procedure has been applied to all the elements of the nonce space, we move to Hybrid 11.

Hybrid 11. In this stage, we change how we generate the master secrets $(\text{mk}'_h)_{h \in H}$ and $(\hat{\text{mk}}_i)_{i \in C}$. Specifically, instead of sampling the randomness fed into Secret uniformly, we rely on the output of the PRF F'_k when its nonce is the position where we had previously punctured. Observe that this hybrid is indistinguishable from the previous one by the security of the puncturable PRF F' .

The formal procedure used by the challenger for the generation of CGP becomes the following.

1. $\hat{k} \leftarrow \text{Punct}'(k, (\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n))$
2. $(\hat{s}_1, \hat{s}_2, \dots, \hat{s}_n) \leftarrow F'_k(\hat{\text{pk}}_1, \hat{\text{pk}}_2, \dots, \hat{\text{pk}}_n)$
3. $\forall i \in [n]: \text{mk}'_i \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, i; \hat{s}_i)$
4. $\forall i \in C: \hat{\text{mk}}_i \leftarrow \text{mk}'_i$

$\mathcal{P}_{CG}^{\hat{x},4}[k, K, (\hat{pk}_i, \hat{mk}_i)_{i \in [n]}, (mk'_h)_{h \in H}, H, \hat{x}]$

Hard-coded. The puncturable PRF keys k and K , the public keys $(\hat{pk}_i)_{i \in [n]}$ and the master secrets $(\hat{mk}_i)_{i \in [n]}$ and $(mk'_h)_{h \in H}$, the set of honest parties H and the nonce \hat{x} .

Input. A nonce $x \in \{0, 1\}^{l(\lambda)}$ and n public keys pk_1, pk_2, \dots, pk_n .

1. $(r, r_1, r_2, \dots, r_n) \leftarrow F_K(pk_1, pk_2, \dots, pk_n, x)$.
2. If $pk_i = \hat{pk}_i$ for every $i \in [n]$ and $x >_{\text{lex}} \hat{x}$, set $mk_i \leftarrow \hat{mk}_i$ for every $i \in [n]$.
3. If $pk_i = \hat{pk}_i$ for every $i \in [n]$ and $x \leq_{\text{lex}} \hat{x}$, set $mk_h \leftarrow mk'_h$ for every $h \in H$ and $mk_i \leftarrow \hat{mk}_i$ for every $i \in C$.
4. Otherwise, perform the following operations
 - (a) $(s_1, s_2, \dots, s_n) \leftarrow F'_k(pk_1, pk_2, \dots, pk_n)$
 - (b) $\forall i \in [n]: mk_i \leftarrow \text{Secret}(\mathbb{1}^\lambda, i; s_i)$
5. $(R_1, R_2, \dots, R_n) \leftarrow \mathcal{C}(\mathbb{1}^\lambda, mk_1, mk_2, \dots, mk_n; r)$
6. $\forall i \in [n]: c_i \leftarrow \text{PKE.Enc}(pk_i, (R_i, mk_i); r_i)$
7. Output c_1, c_2, \dots, c_n .

Fig. 30. The Correlation Generation Program

5. $\forall h \in H: \hat{mk}_h \xleftarrow{\$} \text{Secret}(\mathbb{1}^\lambda, h)$
6. $\text{CGP} \xleftarrow{\$} \text{iO}(\mathbb{1}^\lambda, \mathcal{P}_{CG}^{\Omega,4}[\hat{k}, K, (\hat{pk}_i, \hat{mk}_i)_{i \in [n]}, (mk'_h)_{h \in H}, H, \Omega])$ (see Fig. 30)

Hybrid 12. This is the final stage and corresponds to the ideal world. In this hybrid, we do not puncture k anymore and we generate CGP by obfuscating the original program \mathcal{P}_{CG} (see Fig. 19). Observe that the input-output behaviour of CGP is the same as in Hybrid 11. Indeed, there is no $x >_{\text{lex}} \hat{x}$ because \hat{x} has reached the maximum. Indistinguishability holds by the security of iO .

Also notice that the challenger replies to every query $(\text{Correlation}, x)$ using the following procedure.

1. $(c_i)_{i \in [n]} \leftarrow \text{CGP}(\hat{pk}_1, \hat{pk}_2, \dots, \hat{pk}_n, x)$
2. $\forall i \in C: (R_i, mk_i) \leftarrow \text{PKE.Dec}(\hat{sk}_i, c_i)$
3. $(R_h)_{h \in H} \xleftarrow{\$} \text{RSample}(\mathbb{1}^\lambda, C, (R_i)_{i \in C}, (\hat{mk}_i)_{i \in C}, (\hat{mk}_h)_{h \in H})$

Furthermore, observe that the master secrets $(\hat{mk}_h)_{h \in H}$ are sampled at random using Secret and they are independent of the public keys and CGP. Finally, we formalise the operations of the extractor.

$\text{Extract}(C, \text{CGP}, \rho_1, \dots, \rho_n)$

1. $\forall i \in [n]: (\hat{sk}_i, \hat{pk}_i) \leftarrow \text{PKE.Gen}(\mathbb{1}^\lambda; \rho_i)$
2. $(c_i)_{i \in [n]} \leftarrow \text{CGP}(\hat{pk}_1, \hat{pk}_2, \dots, \hat{pk}_n, \epsilon)$
3. $\forall i \in C: (R_i, mk_i) \leftarrow \text{PKE.Dec}(\hat{sk}_i, c_i)$
4. Output $(mk_i)_{i \in C}$.

Observe that in every hybrid the size of CGP is polynomial in the length $l(\lambda)$ of the nonces. Furthermore, the size of the keys is always independent of the size of the nonce space. \square

C Proof of Theorem 6.10

Proof. Let $\text{pkPCFS} = (\text{Setup}, \text{Gen}, \text{Eval})$ be the sub-exponentially and semi-maliciously secure public key PCF with trusted setup for $(\text{Secret}, \mathcal{C})$. Assume that $\text{DS} = (\text{Gen}, \text{Sample})$ is a sub-exponentially and semi-maliciously secure distributed sampler for pkPCFS.Setup . Moreover, suppose that the algorithms DS.Gen and pkPCFS.Gen need $L(\lambda)$ and $L'(\lambda)$ bits of randomness for their execution respectively. We rely on a PRG G mapping a λ -bit seed into a pseudo-random string of $L(\lambda) + L'(\lambda)$ bits. We assume that the output of G is naturally split into two blocks of length $L(\lambda)$ and $L'(\lambda)$ bits respectively. Finally, let $\text{NIZK} = (\text{Gen}, \text{Prove}, \text{Verify})$ be a simulation-extractable NIZK proving the well-formedness of sampler shares and PCF public keys. Specifically, in the relation corresponding to NIZK , the statement consists of a tuple (U, pk, i) , whereas the witness is a pair (s, sk) such that

$$U = \text{DS.Gen}(\mathbb{1}^\lambda, i; r), \quad (\text{sk}, \text{pk}) = \text{pkPCFS.Gen}(\mathbb{1}^\lambda, i; r'), \quad (r, r') = G(s).$$

Our actively secure public key PCF $\Pi_{\text{exp-}\mathcal{C}}$ is described in Fig. 31. We now prove that no PPT adversary is able to distinguish between $\Pi_{\text{exp-}\mathcal{C}}$ and the composition of $\mathcal{F}_{\mathcal{C}}^{\text{RSample}}$ with a PPT simulator we are going to present. The proof proceeds by a series of 6 indistinguishable hybrids (some of them repeated multiple times) going from $\Pi_{\text{exp-}\mathcal{C}}$ (Hybrid 0) to the ideal world (Hybrid 6).

$\Pi_{\text{exp-}\mathcal{C}}$

URS. The protocol needs a URS $\text{urs} \stackrel{\$}{\leftarrow} \text{NIZK.Gen}(\mathbb{1}^\lambda)$ for the NIZK proofs.

Initialisation. Each party P_i performs the following steps.

1. $s_i \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$
2. $(r_i, r'_i) \leftarrow G(s_i)$
3. $U_i \leftarrow \text{DS.Gen}(\mathbb{1}^\lambda, i; r_i)$
4. $(\text{sk}_i, \text{pk}_i) \leftarrow \text{pkPCFS.Gen}(\mathbb{1}^\lambda, i; r'_i)$
5. $\pi_i \stackrel{\$}{\leftarrow} \text{NIZK.Prove}(\mathbb{1}^\lambda, \text{urs}, (U_i, \text{pk}_i, i), (s_i, \text{sk}_i))$
6. Broadcast $(U_i, \text{pk}_i, \pi_i)$ and wait for a similar message from every other party.
7. If there exists $j \in [n]$ such that $\text{NIZK.Verify}(\text{urs}, \pi_j, (U_j, \text{pk}_j, j)) = 0$ abort.
8. $S \leftarrow \text{DS.Sample}(U_1, U_2, \dots, U_n)$

Correlation. On input a nonce $x \in \{0, 1\}^{l(\lambda)}$, each party P_i outputs $R_i \leftarrow \text{pkPCFS.Eval}(i, S, (\text{pk}_j)_{j \in [n]}, \text{sk}_i, x)$.

Fig. 31. Actively Secure Public Key PCF based on Sub-Exponentially Secure Primitives

Hybrid 0. This hybrid coincides with the real world. The simulator runs the protocol $\Pi_{\text{exp-}\mathcal{C}}$ on behalf of the honest parties. Specifically, it starts its execution producing the URS for the NIZK proofs, it generates the sampler shares and the

keys of the honest parties, proves their well-formedness and sends everything except the private keys to the adversary. Moreover, the simulator replies to the correlation queries using pkPCFS.Eval as in $\Pi_{\text{exp-C}}$.

Hybrid 1. In this hybrid, we change how we generate the URS and the NIZK proofs of the honest parties. Specifically, we substitute them with the output of the simulators NIZK.Sim_1 and NIZK.Sim_2 . Indistinguishability between Hybrid 0 and 1 follows from the multi-theorem zero-knowledge of NIZK. Formally, the steps performed by the simulator for the generation of the messages of the honest parties are the following (the red text indicates what changed since the last hybrid).

1. $\forall h \in H : (r_h, r'_h) \leftarrow G(s_h)$
2. $\forall h \in H : U_h \xleftarrow{\$} \text{DS.Gen}(\mathbb{1}^\lambda, h)$
3. $\forall h \in H : (\text{sk}_h, \text{pk}_h) \leftarrow \text{pkPCFS.Gen}(\mathbb{1}^\lambda, h; r'_h)$
4. $(\text{urs}, \tau) \xleftarrow{\$} \text{NIZK.Sim}_1(\mathbb{1}^\lambda)$
5. $\forall h \in H : \pi_h \xleftarrow{\$} \text{NIZK.Sim}_2(\text{urs}, \tau, (U_h, \text{pk}_h, h))$

Hybrid 2. In this hybrid, we change how we generate the randomness $(r_h, r'_h)_{h \in H}$ of the honest parties. Specifically, instead of expanding the seed s_h , we sample r_h and r'_h uniformly in $\{0, 1\}^{L(\lambda)}$ and $\{0, 1\}^{L'(\lambda)}$ respectively. Observe that this hybrid is indistinguishable from the previous one by the PRG security of G .

Next, we repeat the hybrids from 3 to 5 for every possible choice of the seeds of the corrupted parties $\hat{\rho} := (\hat{s}_i)_{i \in C}$. We follow the lexicographical order starting from the minimum.

Hybrid 3. $\hat{\rho}$. In this hybrid, we change how we generate the sampler shares of the honest parties. Let (\hat{r}_i, \hat{r}'_i) be $G(\hat{s}_i)$ for every $i \in C$. We rely on DS.Sim , providing it with the randomness $(\hat{r}_i)_{i \in C}$ and an element \hat{S} produced by pkPCFS.Setup .

We also change the way we reply to the correlation queries. At the beginning of its execution, the simulator samples a random mk'_h for every $h \in H$. When it receives the messages of the corrupted parties, the simulator extracts the witnesses from the NIZK proofs, obtaining the seeds $\rho = (s_i)_{i \in C}$ and the corresponding secret keys. The operation can be performed due to simulation-extractability. Furthermore, the simulator derives the master secrets $(\text{mk}_i)_{i \in C}$ of the corrupt parties by computing $(r_i, r'_i) \leftarrow G(s_i)$ for every $i \in C$ and running pkPCFS.Extract on $(r'_i)_{i \in [n]}$ and S . The reply to the correlation queries is then computed as follows.

- If $\rho = \hat{\rho}$, the simulator substitutes the output of DS.Sample with \hat{S} in pkPCFS.Eval . It answers with the results.
- If $\rho >_{\text{lex}} \hat{\rho}$, the simulator replies as in the real protocol.
- If $\rho <_{\text{lex}} \hat{\rho}$, the simulator extracts the outputs of the corrupted parties by relying on their private keys, feeds the obtained values into RSample along with $(\text{mk}_i)_{i \in C}$ and $(\text{mk}'_h)_{h \in H}$ and, at the end, answers with the results.

Observe that this hybrid is indistinguishable from the previous one by the semi-malicious security of DS . Notice that when $\hat{\rho}$ is minimum, the simulator

never relies on RSample . The formal steps used by the simulator to generate $(U_h)_{h \in H}$ become the following.

1. $\forall i \in C : (\hat{r}_i, \hat{r}'_i) \leftarrow G(\hat{s}_i)$
2. $\hat{S} \stackrel{\$}{\leftarrow} \text{pkPCFS.Setup}(\mathbb{1}^\lambda)$
3. $(U_h)_{h \in H} \stackrel{\$}{\leftarrow} \text{DS.Sim}(\mathbb{1}^\lambda, C, \hat{S}, (\hat{r}_i)_{i \in C})$

The reply to $(\text{Correlation}, x)$ is instead computed as follows.

1. $S \leftarrow \text{DS.Sample}(U_1, U_2, \dots, U_n)$
2. $\forall i \in C : (s_i, \text{sk}_i) \leftarrow \text{NIZK.Extract}(\text{urs}, \tau, (U_i, \text{pk}_i, i), \pi_i)$
3. $\forall i \in C : (r_i, r'_i) \leftarrow G(s_i)$
4. $(\text{mk}_i)_{i \in C} \leftarrow \text{pkPCFS.Extract}(C, S, r'_1, r'_2, \dots, r'_n)$
5. $\rho \leftarrow (s_i)_{i \in C}$
6. If $\rho = \hat{\rho}$, compute $R_h \leftarrow \text{pkPCFS.Eval}(h, \hat{S}, (\text{pk}_i)_{i \in [n]}, \text{sk}_h, x) \forall h \in H$.
7. If $\rho >_{\text{lex}} \hat{\rho}$, compute $R_h \leftarrow \text{pkPCFS.Eval}(h, S, (\text{pk}_i)_{i \in [n]}, \text{sk}_h, x) \forall h \in H$.
8. If $\rho <_{\text{lex}} \hat{\rho}$, compute
 - (a) $\forall i \in C : R_i \leftarrow \text{pkPCFS.Eval}(i, S, (\text{pk}_j)_{j \in [n]}, \text{sk}_i, x)$
 - (b) $(R_h)_{h \in H} \stackrel{\$}{\leftarrow} \text{RSample}(\mathbb{1}^\lambda, C, (R_i)_{i \in C}, (\text{mk}_i)_{i \in C}, (\text{mk}'_h)_{h \in H})$

Hybrid 4. $\hat{\rho}$. In this hybrid, we change how we generate the outputs of the honest parties when the seeds of the corrupted players coincide with $\hat{\rho}$. Specifically, we now reverse sample them. We can indeed retrieve the samples addressed to the corrupted parties using their private keys. The latter can be extracted from the corresponding NIZK proofs along with the seeds of the corrupted players. Using pkPCFS.Extract , it is also possible to derive the master secrets $(\text{mk}_i)_{i \in C}$.

Observe that this hybrid is indistinguishable from the previous one by the semi-malicious security of pkPCFS . As a matter of fact, if the randomness ρ chosen by the adversary is different from $\hat{\rho}$, the two stages are perfectly identical. If instead $\rho = \hat{\rho}$, we can reduce distinguishability between Hybrid 2. $\hat{\rho}$ and 3. $\hat{\rho}$ to the security game $\mathcal{G}_{\text{SetupSec}}^{C, (\hat{r}'_i)_{i \in C}}(\lambda)$.

When $\rho = \hat{\rho}$, the simulator replies to $(\text{Correlation}, x)$ using the following procedure.

1. $\forall i \in C : (s_i, \text{sk}_i) \leftarrow \text{NIZK.Extract}(\text{urs}, \tau, (U_i, \text{pk}_i, i), \pi_i)$
2. $\forall i \in C : (r_i, r'_i) \leftarrow G(s_i)$
3. $(\text{mk}_i)_{i \in C} \leftarrow \text{pkPCFS.Extract}(C, \hat{S}, r'_1, r'_2, \dots, r'_n)$
4. $\forall i \in C : R_i \leftarrow \text{pkPCFS.Eval}(i, \hat{S}, (\text{pk}_j)_{j \in [n]}, \text{sk}_i, x)$
5. $(R_h)_{h \in H} \stackrel{\$}{\leftarrow} \text{RSample}(\mathbb{1}^\lambda, C, (R_i)_{i \in C}, (\text{mk}_i)_{i \in C}, (\text{mk}'_h)_{h \in H})$

Hybrid 5. $\hat{\rho}$. In this hybrid, we revert to the original procedure for the generation of the sampler shares of the honest parties. Specifically, we do not rely anymore on DS.Sim , but we use DS.Gen . Observe that this hybrid is indistinguishable from the previous one by the semi-malicious security of DS .

If $\rho = \hat{\rho}$, the procedure used by the simulator to reply to $(\text{Correlation}, x)$ becomes the following.

1. $S \leftarrow \text{DS.Sample}(U_1, U_2, \dots, U_n)$

2. $\forall i \in C : (s_i, \text{sk}_i) \leftarrow \text{NIZK.Extract}(\text{urs}, \tau, (U_i, \text{pk}_i, i), \pi_i)$
3. $\forall i \in C : (r_i, r'_i) \leftarrow G(s_i)$
4. $(\text{mk}_i)_{i \in C} \leftarrow \text{pkPCFS.Extract}(C, \mathbf{S}, r'_1, r'_2, \dots, r'_n)$
5. $\forall i \in C : R_i \leftarrow \text{pkPCFS.Eval}(i, \mathbf{S}, (\text{pk}_j)_{j \in [n]}, \text{sk}_i, x)$
6. $(R_h)_{h \in H} \stackrel{\mathcal{S}}{\leftarrow} \text{RSample}(\mathbb{1}^\lambda, C, (R_i)_{i \in C}, (\text{mk}_i)_{i \in C}, (\text{mk}'_h)_{h \in H})$

The next step is to repeat Hybrid 3-5 for the next value $\hat{\rho}$ of the seeds of the corrupted parties. If $\hat{\rho}$ has reached the maximum, we move to Hybrid 6.

Hybrid 6. This hybrid corresponds to the ideal world and summarises what we have achieved so far. In this final stage, the simulator selects the sampler shares and the keys of the honest parties as in the original protocol, using however true randomness instead of expanding PRG seeds. The URS and the NIZK proofs are generated by relying on the simulators NIZK.Sim_1 and NIZK.Sim_2 as in Hybrid 1.

When the simulator receives the messages of the corrupted parties from the adversary, it extracts their seeds and private keys from the zero-knowledge proofs using NIZK.Extract . At that point, it has all the necessary information to retrieve the master secrets $(\text{mk}_i)_{i \in C}$ of the corrupted players by means of pkPCFS.Extract . The values are sent to $\mathcal{F}_C^{\text{RSample}}$.

Upon receiving any query $(\text{Correlation}, x)$, the simulator is also able to compute the outputs of the corrupted parties as it knows their private keys. So it is just sufficient to forward the results to the functionality. The latter will take care of the generation and distribution of the samples of the honest players using RSample . The formal description of the simulator is available in Fig. 32. \square

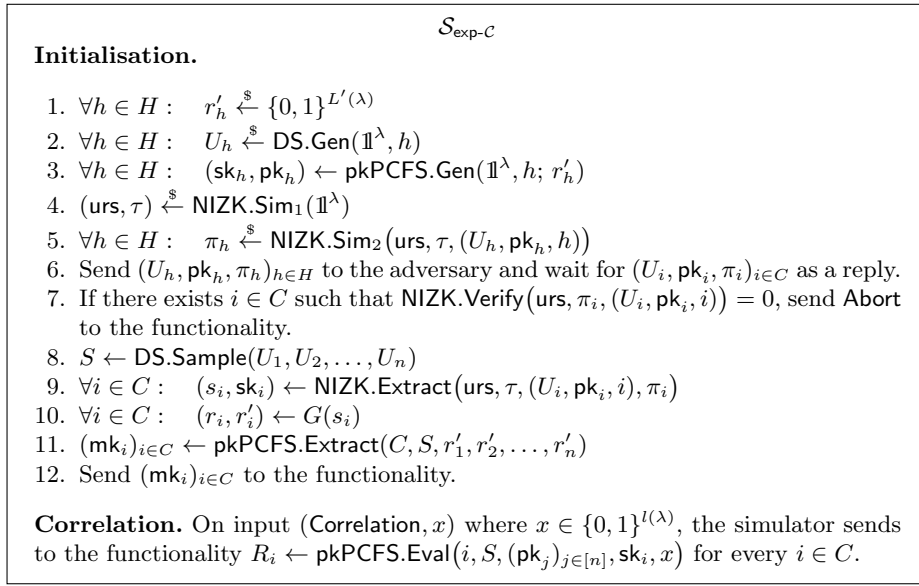


Fig. 32. Simulator for $\Pi_{\text{exp-C}}$