



HAL
open science

Breaking the Circuit Size Barrier for Secure Computation Under Quasi-Polynomial LPN

Geoffroy Couteau, Pierre Meyer

► **To cite this version:**

Geoffroy Couteau, Pierre Meyer. Breaking the Circuit Size Barrier for Secure Computation Under Quasi-Polynomial LPN. EUROCRYPT 2021 - Annual International Conference on the Theory and Applications of Cryptographic Techniques, Oct 2021, Zagreb, Croatia. pp.842-870, 10.1007/978-3-030-77886-6_29 . hal-03374180

HAL Id: hal-03374180

<https://hal.science/hal-03374180>

Submitted on 12 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Breaking the Circuit Size Barrier for Secure Computation under Quasi-Polynomial LPN ^{*}

Geoffroy Couteau¹ and Pierre Meyer²

¹ CNRS, IRIF, Université de Paris, France. couteau@irif.fr

² IRIF, Université de Paris, France and IDC Herzliya, Israel. pierre.meyer@ens-lyon.fr

Abstract. In this work we introduce a new (circuit-dependent) *homomorphic secret sharing* (HSS) scheme for any $\log / \log \log$ -local circuit, with communication proportional only to the width of the circuit and polynomial computation, which is secure assuming the super-polynomial hardness of *learning parity with noise* (LPN). At the heart of our new construction is a *pseudorandom correlation generator* (PCG) which allows two parties to locally stretch short seeds into pseudorandom instances of an arbitrary $\log / \log \log$ -local additive correlation.

Our main application, and the motivation behind this work, is a generic two-party secure computation protocol for every layered (boolean or arithmetic) circuit of size s with total communication $O(s / \log \log s)$ and polynomial computation, assuming the super-polynomial hardness of the standard learning parity with noise assumption (a circuit is layered if its nodes can be partitioned in layers, such that any wire connects adjacent layers). This expands the set of assumptions under which the ‘circuit-size barrier’ can be broken, for a large class of circuits. The strength of the underlying assumption is tied to the sublinearity factor: we achieve communication $O(s/k(s))$ under the $s^{2^{k(s)}}$ -hardness of LPN, for any $k(s) \leq (\log \log s)/4$.

Previously, the set of assumptions known to imply a PCG for correlations of degree $\omega(1)$ or generic secure computation protocols with sublinear communication was restricted to LWE, DDH, and a circularly secure variant of DCR.

Keywords: homomorphic secret sharing · multiparty computation · sublinear communication · learning parity with noise · pseudorandom correlation generators

1 Introduction

In this work, we present a novel (circuit dependent) *homomorphic secret sharing* (HSS) scheme for any $(\log / \log \log)$ -local circuit which is secure under the super-polynomial hardness of the learning parity with noise (LPN) assumption. The main application, and motivation for this work, is a new protocol for securely computing layered arithmetic and boolean circuits with communication sublinear in the circuit size, under the quasi-polynomial hardness of LPN.

Homomorphic Secret Sharing (HSS). An HSS is a compact secret sharing scheme equipped with homomorphism: the parties can locally convert compact shares of an input into (additive) shares of some function of it, without interaction. Compactness here means that the input shares should be much smaller than, and ideally independent of, the size of the evaluated circuit. More precisely, HSS for a circuit class allows the parties to homomorphically convert their shares for any circuit in the class. This powerful primitive has been instantiated for all circuits under LWE [BKS19], or for NC^1 under DDH [BGI16a] or DCR [FGJS17, OSY21, RS21], and for the class of constant degree polynomials from LPN [BCG⁺19b].

The circuit size barrier in secure computation. Secure computation allows mutually distrustful parties to securely compute a public function of their joint private inputs, concealing all information beyond the output. Since its introduction in the seminal works of Yao [Yao86], and Goldreich, Micali, and Wigderson [GMW87b, GMW87a], secure computation has received a constant attention.

^{*} An extended abstract of this work appeared in Advances in Cryptology – EUROCRYPT 2021, 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques. G. Couteau’s research is supported by ANR Project SCENE. This work was done while P. Meyer was a student of the École Normale Supérieure de Lyon, visiting the IRIF and the FACT Center. P. Meyer’s research is supported in part by ISF grant 1861/16, AFOSR Award FA9550-17-1-0069, and ERC Starting Grant 852952(HSS).

For a long time, however, all standard approaches to secure computation have been stuck at an intriguing *circuit-size barrier*, in that they require an amount of communication (at least) proportional to the size of the circuit being computed. In contrast, insecure computation only requires exchanging the inputs, which are usually considerably smaller than the entire circuit. Getting beyond this limitation has been a major challenge in secure computation. Early positive results required exponential computation [BFKR91, NN01], or were limited to very simple functions such as point functions [CGKS95, KO97, CG97] or constant-depth circuits [BI05].

The situation changed with the breakthrough result of Gentry [Gen09] on fully-homomorphic encryption (FHE), which led to optimal communication protocols in the computational setting [DFH12, AJL⁺12]. On the downside, the set of assumptions under which we know how to build FHE is very narrow; it is restricted to lattice-based assumptions such as LWE, and in particular does not include any of the traditional assumptions which were used in the 20th century. More recently, the elegant work of [BGI16a] showed for the first time that secure computation with sublinear communication could be based on assumptions not known to imply FHE, by building a two-party secure computation protocol under the DDH assumption, with communication $O(s/\log s)$ for *layered* circuits of size s .³ [FGJS17, OSY21, RS21] later followed this blueprint and switched out the DDH assumption for DCR assumption. It remains open whether secure computation with sublinear communication can be based on any other traditional and well-studied assumption, such as code-based assumptions.

1.1 Our Contribution

We show that circuit-dependent homomorphic secret sharing, *i.e.* HSS where the share generation requires knowing in advance the circuit to be evaluated homomorphically, for the class of log-local circuits exists, conditioned on (the quasi-polynomial hardness of) a well-studied 20th century assumption: the learning parity with noise (LPN) assumption [BFKL94]. Informally, the LPN assumption captures the hardness of solving an overdetermined system of linear equations over \mathbb{F}_2 , when a small subset of the equations is perturbed with a random noise. The LPN assumption has a long history in computational learning theory, where it emerged. Furthermore, our results only require a flavour of LPN where the adversary is given a very limited number of samples (typically, $O(n)$ equations in n indeterminates). In this regime, LPN is equivalent to the hardness of decoding random linear codes over \mathbb{F}_2 , which is the well-known *syndrome decoding* problem in the coding theory community, where it has been studied since the 60's [Pra62].

Details on the underlying assumption. In a bit more detail, given a security parameter λ , the (T, n, N, r) -LPN assumption with dimension $n = n(\lambda)$, number of samples $N = N(\lambda)$ and noise rate $r = r(\lambda)$ states that for every adversary Adv running in time at most $T = T(\lambda)$,

$$\Pr \left[A \stackrel{\$}{\leftarrow} \mathbb{F}_2^{N \times n}, e \stackrel{\$}{\leftarrow} \text{Ber}_r^N, s \stackrel{\$}{\leftarrow} \mathbb{F}_2^n : \text{Adv}(A, A \cdot s + e) = s \right] = \text{negl}(\lambda),$$

where Ber_r denotes the Bernoulli distribution which outputs 1 with probability r , and negl denote some negligible function. When T can be any polynomial (resp. any super-polynomial function), we say that we assume the polynomial (resp. quasi-polynomial, super-polynomial) hardness of LPN. For arithmetic circuits, we need to assume LPN over large fields, or equivalently syndrome decoding for random linear codes over large fields; this is also a well-founded and well-studied assumption, used in several previous works, e.g. [BCGI18, BCG⁺19b].

HSS for Any loglog-Depth Circuit. We introduce a new circuit-dependent HSS scheme for the class of any log log-depth circuits. We emphasise that unlike traditional forms of HSS, here the input-sharing phase depends on the homomorphic evaluation circuit: in other words it is an HSS scheme for any singleton class comprised of a single circuit of depth log log, not for the class of all such circuits simultaneously.

Main Theorem 1 (HSS for any loglog-Depth Circuit, Informal). *Let C be a size- s , n -input, m -output, $(\epsilon \cdot \log \log)$ -depth arithmetic circuit over \mathbb{F} (for some $\epsilon \leq 1/4$). If the \mathbb{F} -LPN assumption with super-polynomial dimension ℓ , $O(\ell)$ samples, and inverse super-polynomial rate holds then there exists a secure HSS scheme for the class $\{C\}$ with share size $n + O(m \cdot s \cdot \log s / c^{\log^{1-\epsilon} s - \log^{1-2\epsilon} s})$ (for some constant c) and computational complexity $O(m \cdot \text{poly}(s) \cdot (\log |\mathbb{F}|)^2)$.*

³ A depth- d circuit is layered if it can be divided into d layers such that any wire connects adjacent layers.

Restricting the circuit class to depth- k size- s circuits where $k(s) \leq \log \log s/4$ leads to quantitative improvements in the size of the shares, the computational complexity of expanding shares, and the strength of the LPN assumption.

Application to Sublinear Computation. Our HSS scheme has (non black-box) implications for sublinear computation. As in [BGI16a], our results holds for all layered (boolean or arithmetic) circuits, in the two-party setting.

Main Theorem 2 (Sublinear Computation of Layered Circuits, Informal). *For any layered arithmetic circuit C of polynomial size $s = s(\lambda)$ with n inputs and m outputs, for any function $k(s) \leq \log \log s - \log \log \log s + O(1)$, there exists a two party protocol for securely computing C in the honest-but-curious model, with total communication $2(n+m+s/k) \cdot \log |\mathbb{F}| + o(s/k)$ and computation bounded by $s^3 \cdot \text{polylog}(s) \cdot (\log |\mathbb{F}|)^2$ under a set of LPN assumptions, the exact nature of which depends on the sublinearity factor k .*

In particular, setting $k \leftarrow O(\log \log s)$ leads to a protocol with total communication $O(n+m+s/\log \log s)$, secure under the super-polynomial hardness of:

- \mathbb{F} -LPN with super-polynomial dimension ℓ , $O(\ell)$ samples, and inverse super-polynomial rate,
- \mathbb{F}_2 -LPN with super-polynomial dimension ℓ' , $O(\ell')$ samples, and inverse polynomial rate $1/s^{O(1)}$ (which is implied by the above if $\mathbb{F} = \mathbb{F}_2$).

Furthermore (but with a slightly different choice of parameters than the one described above), as k is reduced to an arbitrarily small $k = \omega(1)$, we need only assume the quasi-polynomial hardness of:

- \mathbb{F} -LPN with quasi-polynomial dimension ℓ , $O(\ell)$ samples, and inverse quasi-polynomial rate,
- \mathbb{F}_2 -LPN with quasi-polynomial dimension ℓ' , $O(\ell')$ samples, and inverse polynomial rate $1/s^{O(1)}$ (which is implied by the above if $\mathbb{F} = \mathbb{F}_2$).

and the computation is reduced to $O(s^{1+o(1)} \cdot (\log |\mathbb{F}|)^2)$.

Remark 1. While we require security against super-polynomial-time adversaries, this remains a relatively weak flavour of LPN where the dimension is very high, i.e. super-polynomial as well (and the adversary is allowed to run in time $O(\ell^2)$ where ℓ is the dimension), and the number of samples which the adversary gets is very limited, $O(\ell)$. On the other hand, we require a very small noise rate λ/N . For example, instantiating the above with $k = (\log \log s)/5$, we obtain a secure computation protocol with total communication $O(\ell + m + s/\log \log s)$ (sublinear in s) and polynomial computation, assuming that LPN is hard against adversaries running in super-polynomial time $\lambda^{O(\log \lambda)}$, with dimension $\ell = \lambda^{O(\log \lambda)}$, $N = 2\ell$ samples, and noise rate λ/N . More generally, for any super-constant function $\omega(1)$, there is a two-party protocol with communication $O(n+m+s/\log \omega(1))$ assuming the $\lambda^{\omega(1)}$ -hardness of LPN (i.e., the quasi-polynomial hardness of LPN).

We note that, in this regime of parameters, the best known attacks are the information set decoding attack [Pra62] and its variants (which only shave constant in the exponents, hence have the same asymptotic complexity), which require time $2^{O(\lambda)}$.⁴ Therefore, assuming hardness against $\lambda^{O(\log \lambda)}$ -time adversaries is a very plausible assumption.

Remark 2 (On the Generality of Layered Circuits). Our construction is restricted to the class of (boolean or arithmetic) layered circuits. This restriction stems from the blockwise structure of the construction, and was also present in the previous works of [BGI16a] and [Cou19]. As noted in [Cou19], layered circuits are a relatively large and general class of circuits, which furthermore capture many “real-world” circuits such as FFT-like circuits (used in signal processing, integer multiplication, or permutation networks [Wak68]), Symmetric crypto primitives (e.g. AES and algorithms that proceed in sequences of low-complexity rounds are naturally “layered by blocks”), or dynamic-programming algorithm (e.g. the Smith-Waterman distance, or the Levenshtein distance and its variants).

⁴ BKW and its variants [BKW00, Lyu05] do not improve over information set decoding attacks in this regime of parameters, due to the very low number of samples.

Generalisation to the malicious setting. Our result can directly be generalised to the malicious setting using a generic GMW-style compiler [GMW87a], which is communication preserving when instantiated with succinct zero-knowledge arguments [NN01]. Such arguments exist under collision-resistant hash functions; hence, Theorem 2 extends to the malicious setting as well, at the cost of further assuming collision-resistant hash functions (which is a mild assumption). We note that CRHFs have recently been built from (sub-exponentially strong) flavours of LPN [AHI⁺17, YZW⁺19, BLVW19].

1.2 Our Techniques

Our starting point is the construction of *pseudorandom generator* (PCG) from the work of [BCG⁺19b], under the LPN assumption. At a high level, a PCG allows to distributively generate long pseudorandom instances of a correlation. More precisely, a PCG for a correlation corr (seen as a distribution over pairs of elements) is a pair $(\text{Gen}, \text{Expand})$ where $\text{Gen}(1^\lambda)$ generates a pair of seeds (k_0, k_1) and $\text{Expand}(b, k_b)$ output a string R_b . A PCG must satisfy two properties: (correctness) (R_0, R_1) is indistinguishable from a random sample from corr , and (security) for $b \in \{0, 1\}$, the string R_b is indistinguishable, even given k_{1-b} , from a string R'_b sampled randomly conditioned on satisfying the correlation with R_{1-b} .

The technical contribution at the heart of this paper is to show that, under a certain LPN assumption, there exists a 2-party PCG for the following correlation, which we call *substrings tensor powers* (**stp**) correlation. It is (publicly) parametrised by

- a string length n ;
- subsets $S_1, \dots, S_{n_s} \in \binom{[n]}{\leq K}$ of at most $K = \log n / \log \log n$ many coordinates each;
- a tensor power parameter **tpp** (which can be super-constant, as high as K);

and generates additive shares of all the tensor powers of the prescribed substrings of a random string, *i.e.*

$$(\mathbf{r}, ((1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}})_{1 \leq i \leq n_s}), \text{ where } \mathbf{r} \in \mathbb{F}^n \text{ is (pseudo)random.}$$

In the above, $\mathbf{a}^{\otimes b}$ denotes a vector \mathbf{a} tensored with itself b times. In order to build shares of $(\mathbf{r}, \mathbf{r}^{\otimes 2})$ for some (pseudo)random $\mathbf{r} \in \mathbb{F}^n$ (the bilinear correlation), the PCG of [BCG⁺19b] uses a multi-point function secret sharing scheme (MPFSS) (defined in section 3.1) to give the parties small seeds which can be expanded locally to shares of $(\mathbf{e}, \mathbf{e}^{\otimes 2})$ for some random sparse vector $\mathbf{e} \in \mathbb{F}^n$. Thence, if H is some suitable public matrix the parties can get shares of $\mathbf{r} := H \cdot \mathbf{e}$, which is pseudorandom under LPN, and of $\mathbf{r}^{\otimes 2} = H^{\otimes 2} \cdot \mathbf{e}^{\otimes 2}$ by locally multiplying their shares of \mathbf{e} and $\mathbf{e}^{\otimes 2}$ by H and $H^{\otimes 2}$ respectively. The main issue in using this approach directly is that performing the expanding $\mathbf{r}^{\otimes \text{tpp}} = H[S_i]^{\otimes \text{tpp}} \cdot \mathbf{e}^{\otimes \text{tpp}}$ (where $H[S_i]$ —abusively—denotes the submatrix of H with only the rows indexed by elements of S_i) would require super-polynomial computation, as $H[S_i]$ has n columns.

The core idea of our work is to develop a very careful modified strategy. Instead of letting each \mathbf{r} be a (pseudo)random mask, we construct \mathbf{r} as a sum of $n \cdot \log n$ vectors \mathbf{r}_j , each associated with a public subset of at most K coordinates: these K coordinates are random, but all others are zero. The crucial property achieved by this construction is the following: with high probability, the sum of these sparse vectors will be pseudorandom, but every size- K substring of \mathbf{r} (and in particular S_1, \dots, S_{n_s}) will be expressible as a sum of ‘not too many’ of the \mathbf{r}_j . This allows the expanding to be done by raising to the tensor power **tpp** a matrix whose dimensions are both $K^{O(1)}$, and not n as before. Thus computation remains polynomial.

If we were to stop here, the size of the seeds would grow linearly with n_s , the number of subsets; this would violate the compactness requirement. Instead, we show that we can batch the subsets into n_s/β groups of at most β subsets each, for some parameter β to be refined, to reduce the share size and recover compactness, without harming computational efficiency. Indeed, so long as β is not too large, the substring of \mathbf{r} associated with the union of any β size- K subsets of coordinates will still be expressible as a sum of ‘not too many’ of the \mathbf{r}_j . Our computations reveal a sweet spot for the choice of β , for which the PCG seeds are compact and yet the complexity of expanding them remains polynomial.

1.3 Related Work

Pseudorandom correlation generators were first studied (under the name of cryptocapsules) in [BCG⁺17]. Constructions of PCGs for various correlations, under variants of the LPN assumptions, and applications of PCGs to low-communication secure computation, have been described in [BCG18, BCG⁺19b, BCG⁺19a, SGRR19, BCG⁺20b, BCG⁺20a].

Early works on sublinear-communication secure computation either incurred some exponential cost, or were restricted to very limited types of computations. The first protocols to break the circuit size barriers was shown in [BFKR91] (which gave a protocol with optimal communication, albeit with exponential computation and only for a number of parties linear in the input size). The work of [NN01] gave a sublinear protocol, but with exponential complexity. The work of [BI05] gives a low-communication protocol for constant-depth circuit, for a number of parties polylogarithmic in the circuit size, and the works of [CGKS95, KO97, CG97] gave sublinear protocols for the special case of point functions. The result of Gentry [Gen09] led to the first optimal communication protocols in the computational setting [DFH12, AJL⁺12] under LWE-style assumptions, for all circuits and without incurring any exponential cost. The work of [IKM⁺13] gave an optimal communication protocol in the correlated randomness model, albeit using an exponential amount of correlated randomness. More recently, [Cou19] constructed an unconditionally secure MPC protocol with sublinear communication for layered circuits, in the two-party setting, with a polynomial amount of correlated randomness. Finally, progress in breaking the circuit-size barrier for layered circuits in the computational setting is closely tied to the advances in HSS for super-constant depth circuits [BGI16a, FGJS17, OSY21, RS21].

2 Technical Overview

Notations. We say that a function $\text{negl}: \mathbb{N} \rightarrow \mathbb{R}^+$ is *negligible* if it vanishes faster than every inverse polynomial. For two families of distributions $X = \{X_\lambda\}$ and $Y = \{Y_\lambda\}$ indexed by a security parameter $\lambda \in \mathbb{N}$, we write $X \stackrel{c}{\approx} Y$ if X and Y are *computationally indistinguishable* (i.e. any family of circuits of size $\text{poly}(\lambda)$ has a negligible distinguishing advantage), $X \stackrel{s}{\approx} Y$ if they are *statistically indistinguishable* (i.e. the above holds for arbitrary, unbounded, distinguishers), and $X \equiv Y$ if the two families are identically distributed.

We usually denote matrices with capital letters (A, B, C) and vectors with bold lowercase (\mathbf{x}, \mathbf{y}). By default, vectors are assumed to be column vectors. If \mathbf{x} and \mathbf{y} are two (column) vectors, we use $\mathbf{x} \parallel \mathbf{y}$ to denote the (column) vector obtained by their concatenation. We write $\mathbf{x} \otimes \mathbf{y}$ to denote the tensor product between \mathbf{x} and \mathbf{y} , i.e., the vector of length $n_x n_y$ with coordinates $x_i y_j$ (where n_x is the length of \mathbf{x} and n_y is the length of \mathbf{y}). We write $\mathbf{x}^{\otimes 2}$ for $\mathbf{x} \otimes \mathbf{x}$, and more generally, $\mathbf{x}^{\otimes n}$ for the n -th tensor power of \mathbf{x} , $\mathbf{x} \otimes \mathbf{x} \otimes \dots \otimes \mathbf{x}$. Given a vector \mathbf{x} of length $|\mathbf{x}| = n$, the notation $\text{HW}(\mathbf{x})$ denotes the Hamming weight \mathbf{x} , i.e., the number of its nonzero entries. Let k be an integer. We let $\{0, 1\}^k$ denote the set of bitstrings of length k . For two strings (x, y) in $\{0, 1\}^k$, we denote by $x \oplus y$ their bitwise xor.

Circuits. An arithmetic circuit C with n inputs and m outputs over a field \mathbb{F} is a directed acyclic graph with two types of nodes: the *input nodes* are labelled according to variables $\{x_1, \dots, x_n\}$; the (*computation*) *gates* are labelled according to a base B of arithmetic functions. In this work, we will focus on arithmetic circuits with indegree two, over the standard basis $\{+, \times\}$. C contains m gates with no children, which are called *output gates*. If there is a path between two nodes (v, v') , we say that v is an *ancestor* of v' . In this work, we will consider a special type of arithmetic circuits, called *layered arithmetic circuits* (LBC). An LBC is an arithmetic circuit C whose nodes can be partitioned into $D = \text{depth}(C)$ layers (L_1, \dots, L_d) , such that any edge (u, v) of C satisfies $u \in L_i$ and $v \in L_{i+1}$ for some $i \leq d - 1$. Note that the width of a layered arithmetic circuit is also the maximal number of non-output gates contained in any single layer. Evaluating a circuit C on input $\mathbf{x} \in \mathbb{F}^n$ is done by assigning the coordinates of \mathbf{x} to the variables $\{x_1, \dots, x_n\}$, and then associating to each gate g of C (seen as an arithmetic function) the value obtained by evaluating g on the values associated to its parent nodes. The output of C on input \mathbf{x} , denoted $C(\mathbf{x})$, is the vector of values associated to the output gates.

2.1 PCG and HSS

Much like a PCG for the bilinear correlation yields an HSS for degree-two circuits [BCG⁺19b], given a PCG for the `stp` correlation with `tpp` = K , it is almost immediate to build an HSS scheme for any singleton class comprised of a log/loglog-local circuit C (which is the case in particular if its depth is at most $\log \log - \log \log \log$, since the gates have in-degree at most 2). Since the circuit to be homomorphically evaluated on the input shares is known, the `Share` procedure can depend on it (which is not usually the case for HSS). Let S_1, \dots, S_m be the subsets of inputs on which each output depends, and let K denote the locality of C ; we build a (circuit dependent) HSS scheme as follows:

- `HSS.Share`(\mathbf{x}): Generates compact PCG key (k_0, k_1) which expand to shares of $(\mathbf{r}, ((1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}})_{1 \leq i \leq m})$, set $\mathbf{x}' \leftarrow \mathbf{x} \oplus \mathbf{r}$, and give to each party P_σ a share $s_\sigma = (k_\sigma, \mathbf{x}')$.
- `HSS.Eval`(σ, s_σ): Expand s_σ and, for each $i = 1 \dots m$, extract a share of $(1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}}$. Use it to generate shares of the coefficients of the “degree- K polynomial” on $|S_i| \leq K$ variables P_i satisfying $P_i(X) = C(X - \mathbf{r}[S_i])$. Output the inner product of the vector of coefficient shares with the vector $(1_{\mathbb{F}} \parallel \mathbf{x}')^{\otimes K}$. (This linear product is a share of $P_i(\mathbf{x}')$.)

Correctness and security follow from inspection, along the same lines as [BCG⁺19b]. Usually, `HSS.Share` is given only a circuit class as auxiliary input, not a specific circuit, and the parties should be able to homomorphically evaluate any circuit in the class. In our case however the HSS is circuit-dependent, because the subsets S_1, \dots, S_m are intrinsically tied to the evaluated circuit. An alternative formulation is that our HSS scheme supports singleton circuit classes (or, more generally, local circuits with the same pattern of subsets).

2.2 Generating Correlated Randomness from a PCG

From now on, we set the number of parties to $N = 2$. The work of [BCG⁺19b, Section 6] provides a pseudorandom correlation generator under the LPN assumption, generates correlated (pseudo) random strings for the low-degree polynomial correlation, *i.e.* shares of $(\mathbf{r}, \mathbf{r}^{\otimes 2}, \dots, \mathbf{r}^{\otimes d})$ for some constant d , where \mathbf{r} is a (pseudo)random vector. With the construction from the previous paragraph, this yields an HSS for constant-depth circuits. Our goal is to design a PCG which would lead to an HSS for super-constant depth circuits. More specifically, and keeping our end application in mind, we would like for our PCG to have short enough seeds to lead to a *compact* HSS scheme (*i.e.*, shares of an input x should be at most $O(x)$). This is fundamental when using the scheme to generate correlated randomness in the protocol of [Cou19], which achieves sublinear communication in the correlated randomness model, and which is the starting point of our application to sublinear secure computation.

Our approach is therefore to directly plug in the construction of [BCG⁺19b] and see where it fails. Two issues emerge: the computation is super-polynomial, and the communication not sublinear. Below, we outline each of these issues, and explain how we overcome them.

First Issue: Too Many Polynomials. The first problem which appears when plugging the PCG of [BCG⁺19b] in the protocol of [Cou19] is that the latter requires distributing *many* shares of multivariate polynomials \hat{Q} – more precisely, s/k such polynomials (one for each coordinate of each first layer of a bloc). While the PCG of [BCG⁺19b] allows to compress pseudorandom pairs $(\mathbf{r}, Q(\mathbf{X} - \mathbf{r}))$ into short seeds, these seeds will still be of length at least $\omega(\log \lambda)$, where λ is the security parameter, for the PCG to have any hope of being secure. That means that even if we could manage to securely distribute all these seeds with optimal communication protocols, the overall communication would still be at the very least $\omega((s \log \lambda) / \log \log s)$, which cannot be sublinear since $\log \log s = o(\log \lambda)$ (as s is polynomial in λ).

We solve this first issue as follows: we fix a parameter β , and partition each \mathbf{y}_i into w/β subvectors, each containing β consecutive coordinates of \mathbf{y}_i . Then, the core observation is that a simple variant of the PCG of [BCG⁺19b] allows in fact to generate shares of $(\mathbf{r}, \mathbf{r}^{\otimes 2}, \dots, \mathbf{r}^{\otimes 2^k})$ for some pseudorandom r , where $\mathbf{r}^{\otimes j}$ denotes the tensor product of \mathbf{r} with itself j times (which we call from now on the *j -th tensor power of \mathbf{r}*): this correlation is enough to generate shares of all degree- 2^k polynomial in \mathbf{r} rather than a single one. We will build upon this observation to show how to generate a batch of β shares of multivariate polynomials from a single tensor-power correlation, thus reducing the number of PCG seeds required in the protocol by a factor of β , at the tolerable cost of slightly increasing the size of each seed.

Solution: Batching β Multivariate Polynomials. Consider the first length- β subvector of \mathbf{y}_{i+1} , which we denote \mathbf{v} . Observe that the entire subvector \mathbf{v} can depend on at most $\beta \cdot 2^k$ coordinates of \mathbf{y}_i , since each coordinate of \mathbf{v} depends on at most 2^k coordinates of \mathbf{y}_i . Therefore, we can now see the computation of \mathbf{v} from \mathbf{y}_i as evaluating β multivariate polynomials $(Q_1 \cdots, Q_\beta)$, where all multivariate polynomials take as input the same size- $(\beta 2^k)$ subset of coordinates of \mathbf{y}_i . To securely compute shares of \mathbf{v} from shares of \mathbf{y}_i , the parties can use the following type of correlated randomness: they will have shares of $(\mathbf{r}, \mathbf{r}^{\otimes 2}, \dots, \mathbf{r}^{\otimes 2^k})$, where \mathbf{r} is a random mask of length $\beta \cdot 2^k$. Consider the following polynomials:

$$(\hat{Q}_1(\mathbf{X}), \dots, \hat{Q}_\beta(\mathbf{X})) \stackrel{\text{def}}{=} (Q_1(\mathbf{X} - \mathbf{r}), \dots, Q_\beta(\mathbf{X} - \mathbf{r})).$$

Each coefficient of each \hat{Q} can be computed as a degree- 2^k multivariate polynomial in the coordinates of \mathbf{r} – or, equivalently, as a linear combination of the coordinates of $(\mathbf{r}, \mathbf{r}^{\otimes 2}, \dots, \mathbf{r}^{\otimes 2^k})$. Hence, given additive shares of $(\mathbf{r}, \mathbf{r}^{\otimes 2}, \dots, \mathbf{r}^{\otimes 2^k})$, the parties can locally compute additive shares of the coefficients of *all* the polynomials $(\hat{Q}_1, \dots, \hat{Q}_\beta)$. Using the PCG of [BCG⁺19b], the seeds for generating pseudorandom correlations of the form $(\mathbf{r}, \mathbf{r}^{\otimes 2}, \dots, \mathbf{r}^{\otimes 2^k})$ have length:

$$O\left(\lambda^{2^k} \cdot \log\left((\beta \cdot 2^k)^{2^k}\right)\right),$$

where λ is some security parameter related to the hardness of the underlying LPN assumption. Or more simply, using the fact the computational cost of generating the correlations contains the term $(\beta \cdot 2^k)^{2^k}$ which must remain polynomial in s . Therefore, the total number of bits which the parties have to distribute (for all $(d/k) \cdot (w/\beta) = s/(\beta k)$ such seeds) is $O((s/k) \cdot (\lambda^{2^k} \cdot \log s)/\beta)$.

Choosing the Parameter β . Suppose for simplicity that we already have at hand an MPC protocol allowing to securely distribute such seeds between the parties, with linear overhead over the total length of the seeds generated. This means that generating the full material will require a total communication of $c \cdot s \cdot \lambda^{2^k} \cdot \log s/(\beta k)$. By setting β to be larger than $c \cdot \lambda^{2^k} \cdot \log s$, the total communication will be upper bounded by $O(s/k) = O(s/\log \log s)$ when setting $k \leftarrow O(\log \log s)$, which is the highest our techniques will allow it to be pushed. The most important remaining question is whether we can execute this process in polynomial time given such a large β . Put more simply, the core issue is that the *computational complexity* of expanding short seeds to shares of $(\mathbf{r}, \mathbf{r}^{\otimes 2}, \dots, \mathbf{r}^{\otimes 2^k})$ with the PCG of [BCG⁺19b] contains a term of the form $(\beta \cdot 2^k)^{2^k}$. To make the computation polynomial, we must therefore ensure that β is at most $s^{O(2^{-k})}$, which is subpolynomial. Fortunately, this can be done by setting the security parameter λ of the underlying PCG to be $s^{O(2^{-2^k})}$. For instance, for any constant $\epsilon \in]0, 1[$, we can set $\lambda \leftarrow 2^{\log^\epsilon s}$, $k \leftarrow \log \log s/c_\epsilon$, and $\beta \leftarrow s^{O(2^{-k})}$ for some explicit constant $c_\epsilon > 2$, at the cost of now having to assume the *quasi-polynomial security* of the LPN assumption.

Second Issue: Too Much Communication. In the previous paragraphs, we focused on generating the appropriate correlated random coins using sublinear total communication. But doing so, we glossed over the fact that in the full protocol, the parties must *also* broadcast (shares of) values of the form $\mathbf{y} + \mathbf{r}$, where \mathbf{y} contains values of some layer, and \mathbf{r} is some mask. Recall that with the method which we just outlined, the parties must generate such a length- $(\beta 2^k)$ mask \mathbf{r} for the k -ancestors of each length- β subvector of each last layer of a block. Since there are d/k blocks, whose first layers contain w/β subvector each, and since each $\mathbf{y} + \mathbf{r}$ is of length $\beta \cdot 2^k$, this requires to communicate a total of $(d/k) \cdot (w/\beta) \cdot \beta 2^k = s \cdot 2^k/k$ values – and this cannot possibly be sublinear in s . In fact, this issue already appears in [Cou19], where it was solved as follows: rather than picking an independent mask for each vector of ancestors of a node on a layer (or, in our case, of a length- β block of nodes), pick a single \mathbf{r}_i to mask a full layer \mathbf{y}_i , and define the mask for the subset $S_{i,j}$ of ancestors of a target value $y_{i+1,j}$ to be $\mathbf{r}_i[S_{i,j}]$. This implies that the parties must now broadcast a single masked vector $\mathbf{y}_i + \mathbf{r}_i$ for each first layer of a block, reducing the overall communication back to $O(s/k)$. The correlated randomness which the parties must securely distribute now consists of tensor powers of many subsets of the coordinates of each mask.

Using the PCG of [BCG⁺19b] for ‘Subvectors Tensor Powers Correlations’. However, attempting to construct a PCG for generating this kind of correlated randomness from the PCG of [BCG⁺19b] blows up the computation to the point that it can no longer be polynomial. To explain this issue, we briefly recall the high level construction of the PCG of [BCG⁺19b]. To share a pseudorandom vector $(\mathbf{r}, \dots, \mathbf{r}^{\otimes 2^k})$ where \mathbf{r} is of length w , the PCG will first generate a *very sparse* vector \mathbf{r}' , with some number t of nonzero coordinates. Then, each $(\mathbf{r}')^{\otimes n}$ for some $n \leq 2^k$ is itself a t^n -sparse vector, of length w^n . Using multi-point function secret sharing (MPFSS, a primitive which was developed in a recent line of work [GI14, BGI15, BGI16b, BCGI18] and can be built from one way functions), one can compress shares of $(\mathbf{r}')^{\otimes n}$ to length- $t^n \cdot \log w$ seeds. Then, the final pseudorandom correlation is obtained by letting the parties locally compress \mathbf{r}' by multiplying it with a large public matrix H , giving a vector $\mathbf{r} = H \cdot \mathbf{r}'$. Similarly, $\mathbf{r}^{\otimes n}$ can be reconstructed by computing $H^{\otimes n} \cdot (\mathbf{r}')^{\otimes n} = (H \cdot \mathbf{r}')^{\otimes n} = \mathbf{r}^{\otimes n}$, using the multilinearity of tensor powers. The security relies on the fact that if H is a large compressing public random matrix, then its product with a random sparse noise vector \mathbf{r}' is indistinguishable from random, under the dual LPN assumption (which is equivalent to the standard LPN assumption). Concretely, one can think of \mathbf{r}' as being of length $2w$, and of H as being a matrix from $\mathbb{F}^{w \times 2w}$ which compresses \mathbf{r}' to a pseudorandom length- w vector.

Now, the issue with this construction is that even if we need only tensor powers of small subvectors (of length $\beta \cdot 2^k$ in our construction) of the vector \mathbf{r} , the computation for expanding the seed to these pseudorandom tensor powers will grow super-polynomially with the length of of *entire* vector w . Indeed, consider generating the 2^k -th tensor power of a subvector $\mathbf{r}[S]$ of \mathbf{r} , for some size- $\beta \cdot 2^k$ subset S of $[w]$. Then with the PCG of [BCG⁺19b], this requires computing $(H[S])^{\otimes 2^k} \cdot (\mathbf{r}'[S])^{\otimes 2^k}$, where the share of $(\mathbf{r}'[S])^{\otimes 2^k}$ are obtained from a short seed using MPFSS, and $H[S] \in \mathbb{F}^{|S| \times 2w}$ is the submatrix of H whose columns are indexed by S . The core issue becomes now visible: even though $H[S]$ has only $|S|$ rows, it still has $2w$ columns, and computing $H[S]^{\otimes 2^k}$ requires roughly $(|S| \cdot w)^{2^k}$ arithmetic operation. But since we want ultimately to have k be some increasing function of s , the above will contain a term of the form $w^{2^k} = w^{\omega(1)}$, where w (the circuit width) can be polynomial in the circuit size s , leading to an overall computational complexity of $s^{\omega(1)}$, which is super-polynomial.

Solution: Covering the Private Values with the Sum of Separable Masks. Our solution to circumvent the above problem is to generate \mathbf{r} as the sum of a certain number m of shorter masks $\mathbf{r}^1, \mathbf{r}^2, \dots$ which each only cover θ values (note that they may – and will – overlap). This way the 2^k -th tensor power of a subvector \mathbf{v} can be obtained from appropriate linear combinations of coordinates of the 2^k -th tensor power of the concatenation of *only* the \mathbf{r}^j which overlap with \mathbf{v} . The amount of computation grows super-polynomially in the length of this concatenated vector only (instead of w as before).

More formally, we have a list of w/β target subsets $S_1, \dots, S_{w/\beta}$ (each one corresponding to the $2^k\beta$ ancestors of a batch of β outputs) for which we want to compute the 2^k -th tensor power of $\mathbf{r}[S_i]$, for some random $\mathbf{r} \in \mathbb{F}^w$. We want to find M size- K sets $\alpha_1, \alpha_2, \dots, \alpha_M \in \binom{[w]}{K}$ such that each S_i intersects with a small number B of α_j s, while $\cup_{i=1}^M \alpha_i = [w]$. We associate each α_j with a vector $\mathbf{r}^j \in \mathbb{F}^K$: together they define a sparse subvector of \mathbb{F}^w . If we let \mathbf{r} be the sum of these sparse vectors, it is clear that for any $i \in [w/\beta]$, each element of $(1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes 2^k}$ can be obtained by a linear combination of the elements of the 2^k -th tensor power of the vector of size $(1 + BK)$ obtained by concatenating $(1_{\mathbb{F}})$ and the \mathbf{r}^j s such that $\alpha_j \cap S_i \neq \emptyset$. The amount of computation required is then of the order $(BK)^{2^k}$.

The problem of deterministically finding such subsets $\alpha_1, \dots, \alpha_M$ – which we call a *B-Good Cover* of $(S_i)_{i \in [w/\beta]}$ – turns out to be difficult in the general case. Fortunately, there is a straightforward probabilistic solution: choosing them independently and at random works with high probability. More specifically, taking $M \leftarrow O(w \cdot \ln w)$ i.i.d. uniformly random submasks covering $K \leftarrow \beta 2^k$ values each means that the $\beta 2^k$ ancestral inputs of any batch of β outputs will be covered by only a total of roughly $B = \log w$ submasks (the proof of this relies on standard concentration bounds). This effectively lifts the cost of the computation from being super-polynomial in w to being only super-polynomial in $\beta 2^k \log w$, which remains polynomial overall when setting β and k to be appropriately small.

2.3 Application to Sublinear Secure Computation

The work of [Cou19] gives a generic secure protocol with sublinear communication for layered circuits. It works in the *corruptible correlated randomness model*: before the protocol, a trusted dealer lets the

adversary choose the strings that the corrupted parties will get, samples the correlated random coins of the remaining parties afterwards, and distributes them to the parties. As shown in [BCG⁺19b], generating this corruptible randomness using a PCG leads to a secure protocol in the standard model. In a bit more detail, the parties use a generic secure protocol to generate the short seeds (k_0, k_1) then expand them locally; it might have a high overhead, but it will not be a bottleneck since the seeds are very small. We show that our new PCG can be used for just this purpose.

The general idea is to split a layered circuit of size s into carefully chosen blocks, each containing $O(\log \log s)$ consecutive layers. The precise block decomposition is detailed in [Cou19]. Using our PCG cast as an HSS scheme for $O(\log \log s)$ -depth circuits (with the duality described in section 2.1) allows the parties to evaluate the circuit in a block-by-clock fashion: for each block the parties start with additive shares of

- the inputs of the circuit;
- the values of the first layer of the block;

and, using HSS, compute additive shares of

- the outputs of the circuit which are in the block;
- the values of the last layer, which are also the values of the first layer of the next block.

Let us note that since the circuit and its blocks are publicly known to both parties, so the fact our HSS scheme is circuit-dependent is not an issue here. This block-by-block approach allows the parties to ‘skip’ a fraction $O(\log \log(s))$ of the gates when computing the circuit, by communicating at each block rather than at each gate. Unfortunately, combining all these blocks together involves pesky technicalities which prohibit a very modular approach and require us to consider the protocol in its entirety. Indeed, the inputs can appear arbitrarily many times—up to $O(s)$ even—across many blocks, so the randomness used to mask them has to be reused, and we cannot deal with each block using an independent instance of HSS. However, dealing with this problem does not require any additional insight, only more cumbersome notations.

In the above outline, we assumed that we had access to a sufficiently low-communication MPC protocol to distribute the generation of the seeds to our new PCG. To obtain our claimed result, it remains to show that this building block can be instantiated under the quasi-polynomial hardness of LPN. In fact, this MPC protocol needs not have linear communication in the seed size; it turns out that by tuning the parameters appropriately, any fixed polynomial in the seed size suffices to guarantee the existence of a “soft spot” for the parameters of our PCG such that we simultaneously get sublinear total communication $O(s/\log \log s)$ and polynomial computation. Distributing the generation procedure of our PCG essentially boils down to generating (many) seeds for a multi-point function secret sharing scheme, which itself boils down mainly to securely generating seeds for a standard length-doubling pseudorandom generator (PRG), and securely executing about $\log(\text{domsize})$ expansions of these short seeds, where domsize denotes the domain size of the MPFSS. Using a standard LPN-based PRG and GMW-style secure computation, instantiated with an LPN-based oblivious transfer protocol, suffices to securely generate the MPFSS seeds we need.

3 Preliminaries

3.1 Function Secret Sharing

Informally, an FSS scheme for a class of functions \mathcal{C} is a pair of algorithms $\text{FSS} = (\text{FSS.Gen}, \text{FSS.Eval})$ such that:

- FSS.Gen given a function $f \in \mathcal{C}$ outputs a pair of keys (K_0, K_1) ;
- FSS.Eval , given K_b and input x , outputs y_b such that y_0 and y_1 form additive shares of $f(x)$.

The security requirement is that each key K_b computationally hide f , except for revealing the input and output domains of f . Formally:

Definition 3 (Function Secret Sharing; adapted from [BGI16b]). *A 2-party function secret sharing (FSS) scheme for a class of functions $\mathcal{C} = \{f : I \rightarrow \mathbb{G}\}$ with input domain I and output domain an abelian group $(\mathbb{G}, +)$, is a pair of PPT algorithms $\text{FSS} = (\text{FSS.Gen}, \text{FSS.Eval})$ with the following syntax:*

- $\text{FSS.Gen}(1^\lambda, f)$, given security parameter λ and description of a function $f \in \mathcal{C}$, outputs a pair of keys (K_0, K_1) ;
- $\text{FSS.Eval}(b, K_b, x)$, given party index $b \in \{0, 1\}$, key K_b , and input $x \in I$, outputs a group element $y_b \in \mathbb{G}$.

Given an allowable leakage function $\text{Leak} : \{0, 1\}^* \rightarrow \{0, 1\}^*$, the scheme FSS should satisfy the following requirements:

- **Correctness:** For any $f : I \rightarrow \mathbb{G}$ in \mathcal{C} and $x \in I$, we have $\Pr[(K_0, K_1) \stackrel{\$}{\leftarrow} \text{FSS.Gen}(1^\lambda, f) : \sum_{b \in \{0, 1\}} \text{FSS.Eval}(b, K_b, x) = f(x)] = 1$.
- **Security:** For any $b \in \{0, 1\}$, there exists a PPT simulator Sim such that for any polynomial-size function sequence $f_\lambda \in \mathcal{C}$, the distributions $\{(K_0, K_1) \stackrel{\$}{\leftarrow} \text{FSS.Gen}(1^\lambda, f_\lambda) : K_b\}$ and $\{K_b \stackrel{\$}{\leftarrow} \text{Sim}(1^\lambda, \text{Leak}(f_\lambda))\}$ are computationally indistinguishable.

Our application of FSS requires applying the evaluation algorithm on *all inputs*. Following [BGI16b, BCGI18, BCG⁺19b, BCG⁺19a], given an FSS scheme $(\text{FSS.Gen}, \text{FSS.Eval})$, we denote by FSS.FullEval an algorithm which, on input a bit b , and an evaluation key K_b (which defines the input domain I), outputs a list of $|I|$ elements of \mathbb{G} corresponding to the evaluation of $\text{FSS.Eval}(b, K_b, \cdot)$ on every input $x \in I$ (in some predetermined order). Below, we recall some results from [BGI16b] on FSS schemes for useful classes of functions.

Distributed Point Functions A distributed point function (DPF) [GI14] is an FSS scheme for the class of point functions $f_{\alpha, \beta} : \{0, 1\}^\ell \rightarrow \mathbb{G}$ which satisfies $f_{\alpha, \beta}(\alpha) = \beta$, and $f_{\alpha, \beta}(x) = 0$ for any $x \neq \alpha$. A sequence of works [GI14, BGI15, BGI16b] has led to highly efficient constructions of DPF schemes from any pseudorandom generator (PRG).

Theorem 4 (PRG-based DPF [BGI16b]). *Given a PRG $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda+2}$, there exists a DPF for point functions $f_{\alpha, \beta} : \{0, 1\}^\ell \rightarrow \mathbb{G}$ with key size $\ell \cdot (\lambda + 2) + \lambda + \lceil \log_2 |\mathbb{G}| \rceil$ bits. For $m = \lceil \frac{\log |\mathbb{G}|}{\lambda+2} \rceil$, the key generation algorithm Gen invokes G at most $2(\ell + m)$ times, the evaluation algorithm Eval invokes G at most $\ell + m$ times, and the full evaluation algorithm FullEval invokes G at most $2^\ell(1 + m)$ times.*

FSS for Multi-Point Functions Similarly to [BCGI18, BCG⁺19b, BCG⁺19a], we use FSS for *multi-point functions*. A k -point function evaluates to 0 everywhere, except on k specified points. When specifying multi-point functions we often view the domain of the function as $[n]$ for $n = 2^\ell$ instead of $\{0, 1\}^\ell$.

Definition 5 (Multi-Point Function [BCGI18]). *An (n, t) -multi-point function over an abelian group $(\mathbb{G}, +)$ is a function $f_{S, \mathbf{y}} : [n] \rightarrow \mathbb{G}$, where $S = (s_1, \dots, s_t)$ is an ordered subset of $[n]$ of size t and $\mathbf{y} = (y_1, \dots, y_t) \in \mathbb{G}^t$, defined by $f_{S, \mathbf{y}}(s_i) = y_i$ for any $i \in [t]$, and $f_{S, \mathbf{y}}(x) = 0$ for any $x \in [n] \setminus S$.*

We assume that the description of S includes the input domain $[n]$ so that $f_{S, \mathbf{y}}$ is fully specified. A *Multi-Point Function Secret Sharing* (MPFSS) is an FSS scheme for the class of multi-point functions, where a point function $f_{S, \mathbf{y}}$ is represented in a natural way. We assume that an MPFSS scheme leaks not only the input and output domains but also the number of points t that the multi-point function specifies. An MPFSS can be easily obtained by adding t instances of a DPF.

3.2 Learning Parity with Noise

Our constructions rely on the Learning Parity with Noise assumption [BFKL93] (LPN) over a field \mathbb{F} (the most standard variant of LPN typically assumes $\mathbb{F} = \mathbb{F}_2$, but other fields can be considered). Unlike the LWE assumption, in LPN over \mathbb{F} the noise is assumed to have a small Hamming weight. Concretely, the noise is a random field element in a small fraction of the coordinates and 0 elsewhere. Given a field \mathbb{F} , $\text{Ber}_r(\mathbb{F})$ denote the distribution which outputs a uniformly random element of $\mathbb{F} \setminus \{0\}$ with probability r , and 0 with probability $1 - r$.

Definition 6 (LPN). For dimension $k = k(\lambda)$, number of samples (or block length) $q = q(\lambda)$, noise rate $r = r(\lambda)$, and field $\mathbb{F} = \mathbb{F}(\lambda)$, the \mathbb{F} -LPN(k, q, r) assumption states that

$$\{(A, \mathbf{b}) \mid A \stackrel{\$}{\leftarrow} \mathbb{F}^{q \times k}, \mathbf{e} \stackrel{\$}{\leftarrow} \text{Ber}_r(\mathbb{F})^q, \mathbf{s} \stackrel{\$}{\leftarrow} \mathbb{F}^k, \mathbf{b} \leftarrow A \cdot \mathbf{s} + \mathbf{e}\} \\ \stackrel{c}{\approx} \{(A, \mathbf{b}) \mid A \stackrel{\$}{\leftarrow} \mathbb{F}^{q \times k}, \mathbf{b} \stackrel{\$}{\leftarrow} \mathbb{F}^q\}$$

Here and in the following, all parameters are functions of the security parameter λ and computational indistinguishability is defined with respect to λ . Note that the search LPN problem, of finding the vector can be reduced to the decisional LPN assumption [BFKL93, AIK09]. In this paper, our protocols will mostly rely on a variant of LPN, called *exact LPN* (xLPN) [JKPT12]. In this variant, the noise vector \mathbf{e} is not sampled from $\text{Ber}_r(\mathbb{F})^q$, but it is sampled uniformly from the set $\text{HW}_{r,q}(\mathbb{F}^q)$ of length- q vectors over \mathbb{F} with *exactly* rq nonzero coordinates (in contrast, a sample from $\text{Ber}_r(\mathbb{F})^q$ has an *expected* number $r \cdot q$ of nonzero coordinates). While standard LPN is usually preferred since the Bernoulli distribution is convenient to analyze, xLPN is often preferred in concrete implementations, since it offers a potentially higher level of security for similar parameters (by avoiding weak instances with a low amount of noise). Furthermore, as outlined in [JKPT12], xLPN and LPN are equivalent: xLPN reduces to its search version using the sample-preserving reduction of [AIK07], and search-xLPN is easily seen to be polynomially equivalent to search-LPN.

Dual LPN. In our protocols, it will also prove convenient to work with the (equivalent) alternative *dual* formulation of LPN.

Definition 7 (Dual LPN). For dimension $k = k(\lambda)$, number of samples (or block length) $q = q(\lambda)$, noise rate $r = r(\lambda)$, and field $\mathbb{F} = \mathbb{F}(\lambda)$, the dual- \mathbb{F} -LPN(k, q, r) assumption states that

$$\{(H, \mathbf{b}) \mid H \stackrel{\$}{\leftarrow} \mathbb{F}^{q-k \times q}, \mathbf{e} \stackrel{\$}{\leftarrow} \text{Ber}_r(\mathbb{F})^q, \mathbf{b} \leftarrow H \cdot \mathbf{e}\} \\ \stackrel{c}{\approx} \{(H, \mathbf{b}) \mid H \stackrel{\$}{\leftarrow} \mathbb{F}^{q-k \times q}, \mathbf{b} \stackrel{\$}{\leftarrow} \mathbb{F}^q\}$$

Solving the dual LPN assumption is easily seen to be at least as hard as solving LPN: given a sample (A, \mathbf{b}) , define $H \in \mathbb{F}^{q-k \times q}$ to be the parity-check matrix of A (hence $H \cdot A = 0$), and feed $(H, H \cdot \mathbf{b})$ to the dual LPN solver. Note that the parity check matrix of a random matrix is distributed as a random matrix. Furthermore, when $\mathbf{b} = A \cdot \mathbf{s} + \mathbf{e}$, we have $H \cdot \mathbf{b} = H \cdot (A \cdot \mathbf{s} + \mathbf{e}) = H \cdot \mathbf{e}$. For discussions regarding existing attacks on LPN and their efficiency, we refer the reader to [BCGI18, BCG⁺19b].

3.3 Pseudorandom Correlation Generators

Pseudorandom correlation generators (PCG) have been introduced in [BCG⁺19b]. Informally, a pseudorandom correlation generator allows to generate pairs of short keys (or seeds) $(\mathbf{k}_0, \mathbf{k}_1)$ such that each key \mathbf{k}_σ can be expanded to a long string $R_\sigma = \text{Expand}(\sigma, \mathbf{k}_\sigma)$, with the following guarantees: given the key $\mathbf{k}_{1-\sigma}$, the string R_σ is indistinguishable from a random string sampled conditioned on satisfying the target correlation with the string $R_{1-\sigma} = \text{Expand}(1-\sigma, \mathbf{k}_{1-\sigma})$. We provide below the formal definition of pseudorandom correlation generators, after defining the notion of *reverse-sampleable correlation generator*.

Definition 8 (Correlation Generator). A PPT algorithm \mathcal{C} is called a correlation generator, if \mathcal{C} on input 1^λ outputs a pair of elements in $\{0, 1\}^n \times \{0, 1\}^n$ for $n \in \text{poly}(\lambda)$.

In order to define security, we require the notion of a reverse-sampleable correlation generator introduced in the following.

Definition 9 (Reverse-sampleable Correlation Generator). Let \mathcal{C} be a correlation generator. We say \mathcal{C} is reverse sampleable if there exists a PPT algorithm RSample such that for $\sigma \in \{0, 1\}$ the correlation obtained via:

$$\{(R'_0, R'_1) \mid (R_0, R_1) \stackrel{\$}{\leftarrow} \mathcal{C}(1^\lambda), R'_\sigma := R_\sigma, R'_{1-\sigma} \stackrel{\$}{\leftarrow} \text{RSample}(\sigma, R_\sigma)\}$$

is computationally indistinguishable from $\mathcal{C}(1^\lambda)$.

Definition 10 (Pseudorandom Correlation Generator (PCG)).

Let \mathcal{C} be a reverse-sampleable correlation generator. A pseudorandom correlation generator (PCG) for \mathcal{C} is a pair of algorithms (PCG.Gen, PCG.Expand) with the following syntax:

- PCG.Gen(1^λ) is a PPT algorithm that given a security parameter λ , outputs a pair of seeds (k_0, k_1) ;
- PCG.Expand(σ, k_σ) is a polynomial-time algorithm that given party index $\sigma \in \{0, 1\}$ and a seed k_σ , outputs a bit string $R_\sigma \in \{0, 1\}^n$.

The algorithms (PCG.Gen, PCG.Expand) should satisfy the following:

- **Correctness.** The correlation obtained via:

$$\{(R_0, R_1) \mid (k_0, k_1) \xleftarrow{\$} \text{PCG.Gen}(1^\lambda), R_\sigma \leftarrow \text{PCG.Expand}(\sigma, k_\sigma) \text{ for } \sigma \in \{0, 1\}\}$$

is computationally indistinguishable from $\mathcal{C}(1^\lambda)$.

- **Security.** For any $\sigma \in \{0, 1\}$, the following two distributions are computationally indistinguishable:

$$\begin{aligned} & \{(k_{1-\sigma}, R_\sigma) \mid (k_0, k_1) \xleftarrow{\$} \text{PCG.Gen}(1^\lambda), R_\sigma \leftarrow \text{PCG.Expand}(\sigma, k_\sigma)\} \text{ and} \\ & \{(k_{1-\sigma}, R_\sigma) \mid (k_0, k_1) \xleftarrow{\$} \text{PCG.Gen}(1^\lambda), R_{1-\sigma} \leftarrow \text{PCG.Expand}(\sigma, k_{1-\sigma}), \\ & \quad R_\sigma \xleftarrow{\$} \text{RSample}(\sigma, R_{1-\sigma})\} \end{aligned}$$

where RSample is the reverse sampling algorithm for correlation \mathcal{C} .

Note that the above definition is trivial to achieve in general: We can let PCG.Gen on input 1^λ return $(R_0, R_1) \leftarrow \mathcal{C}(1^\lambda)$, and simply define Expand to be the identity. Typically, we will be interested in non-trivial constructions of PCGs, in which the seed size is significantly shorter than the output size. A pseudorandom generator with image in $\{0, 1\}^n$ is a simple example for an expanding PCG for the equality correlation $\{(R, R) \mid R \in \{0, 1\}^n\}$.

4 Secure Computation from Super-Constant-Degree Low-Locality Polynomial Correlated Randomness

4.1 Block Decomposition of Layered Circuits

Given an arithmetic circuit C and an input vector \mathbf{x} , we call *value of the gate g on input \mathbf{x}* the value carried by the output wire of a given gate g of C during the evaluation of $C(\mathbf{x})$. The following decomposition of layered circuits is implicit in [Cou19]; for completeness, we give the proof here.

Lemma 11 (Block-Decomposition of Layered Circuits, Implicit in [Cou19]). *Let C be a layered arithmetic circuit over a field \mathbb{F} with n inputs and m outputs, of size s and depth $d = d(n)$. For any integer k , denoting $t = t(k) = \lceil d/k \rceil$, there exists $2t + 1$ integers $(s_0 = 0, s_1, \dots, s_{t-1}, s_t = 0)$, (m_0, \dots, m_{t-1}) , and functions (f_0, \dots, f_{t-1}) with $f_i : \mathbb{F}^n \times \mathbb{F}^{s_i} \rightarrow \mathbb{F}^{s_{i+1}} \times \mathbb{F}^{m_i}$, such that:*

- The algorithm A given below satisfies, for any input vector $\mathbf{x} \in \mathbb{F}^n$, $A(\mathbf{x}) = C(\mathbf{x})$ (that is, A computes C);

function $A(\mathbf{x})$

$\mathbf{x}_0 \leftarrow \mathbf{x}$

for $i = 0$ **to** $t - 1$ **do** $(\mathbf{x}_{i+1}, \mathbf{y}_i) \leftarrow f_i(\mathbf{x}_i)$

$\mathbf{y} \leftarrow \mathbf{y}_0 \parallel \dots \parallel \mathbf{y}_{t-1}$

return \mathbf{y}

- For any $i \in \llbracket 0, t - 1 \rrbracket$, $j \leq s_{i+1} + m_i$, the j -th output⁵ of $f_i : \mathbb{F}^n \times \mathbb{F}^{s_i} \mapsto \mathbb{F}^{s_{i+1}} \times \mathbb{F}^{m_i}$ can be computed by a multivariate polynomial $P_{i,j}$ over \mathbb{F}^{2^k} of degree $\deg P_{i,j} \leq 2^k$;
- $\sum_{i=0}^{t-1} s_i \leq s/k$ and $\sum_{i=0}^{t-1} m_i = m$.

The decomposition of a layered circuits into chunks computable by low-degree functions is illustrated on Figure 1.

⁵ i.e. the j^{th} coordinate of the image by f_i , seen as $f_i : \mathbb{F}^n \times \mathbb{F}^{s_i} \rightarrow \mathbb{F}^{s_{i+1} + m_i}$.




Fig. 1: Block Decomposition of a Circuit.

Proof. Let C be a layered boolean circuit with n inputs and m outputs, of size s and depth d , with layers (L_1, \dots, L_d) . For $i = 1$ to d , we let w_i denote the width of the layer L_i (that is, the number of computation gates it contains; note that $s = \sum_{i=1}^d w_i$). Fix an integer k and let $t = \lceil d/k \rceil$.

We start by considering t ‘chunks’ of layers, each containing k consecutive layers (the last may in fact contain fewer if $k \nmid d$). Observe there must exist a $j \in \{0, \dots, k-1\}$, such that the sum of the widths of the j^{th} layer of each chunk (with the convention that if the last chunk has fewer than j layer, its j^{th} one is empty; $w_i = 0$ if $i > d$) is at most s/k , *i.e.* $\sum_{i=1}^t w_{k \cdot (i-1) + j} \leq s/k$. Indeed otherwise $\forall j \in \{0, \dots, k\}$, $\sum_{i=0}^{t-1} w_{r_0+1+k \cdot (i-1)+j} > s/k$, so $s = \sum_{i=1}^d w_i \geq \sum_{i=r_0+1}^{d-r_1} w_i = \sum_{j=0}^{k-1} \sum_{i=1}^t w_{r_0+1+k \cdot (i-1)+j} > k \cdot s/k = s$, which is a contradiction. With j being fixed we now define $j_i \leftarrow k \cdot (i-1) + j$ for $i \in [t-1]$ and $j_t \leftarrow \min(d, k \cdot (t-1) + j)$.

Now, for each $0 \leq i \leq t$, we let B_i the block containing the consecutive layers $(L_{j_i}, \dots, L_{j_{i+1}})$. Note that the depth of each block is at most k . Let m_i denote the number of output nodes contained in B_i (note that $\sum_{i=0}^{t-1} m_i = m$). For each $1 \leq i \leq t$, set $s_i \leftarrow w_{j_i}$, and $s_0, s_d \leftarrow 0$. This decomposition into blocks is illustrated in fig. 1.

The intuition behind the decomposition of C is the following: each function f_i will take as input the n inputs \mathbf{x} to C , together with the s_i values of the gates in the first layer of B_i . It evaluates the layers of B_i , starting from the s_i values of the first layer (using the input \mathbf{x} when the layer contains an input node), and outputs the s_{i+1} values on the first layer of B_{i+1} , together with the m_i values of the output nodes contained in B_i . Given this decomposition, the correctness of algorithm A is guaranteed by definition: A simply corresponds to a ‘block-by-block’ evaluation of the circuit C , where each block evaluation outputs the current state \mathbf{x}_{i+1} (which must be given as input to the next block in addition to the input vector \mathbf{x}) together with the outputs of C contained in this block. Since each block has depth at most k , each output of f_i can be computed by multivariate polynomials with at most 2^k inputs, and of degree at most 2^k . \square

4.2 Securely Computing C in the Correlated Randomness Model

We represent in fig. 2 the ideal functionality for securely evaluating the layered arithmetic circuit C .

We represent on fig. 3 an ideal functionality for distributing (function-dependent) correlated randomness between the parties.

Theorem 12. *Let $k \leq \log \log s - \log \log \log s$. There exists a protocol Π_C which (perfectly) securely implements the N -party functionality \mathcal{F}_C in the $\mathcal{F}_{\text{corr}}$ -hybrid model, against a static, passive, non-aborting adversary corrupting at most $N-1$ out of N parties, with communication complexity upper bounded by $O(N \cdot (n + \frac{s}{k} + m) \cdot \log |\mathbb{F}|)$ and polynomial computation.*

The protocol follows closely the construction of [Cou19], with some tedious technical adaptations which are necessary to rely on the specific type of correlated randomness which we will manage to

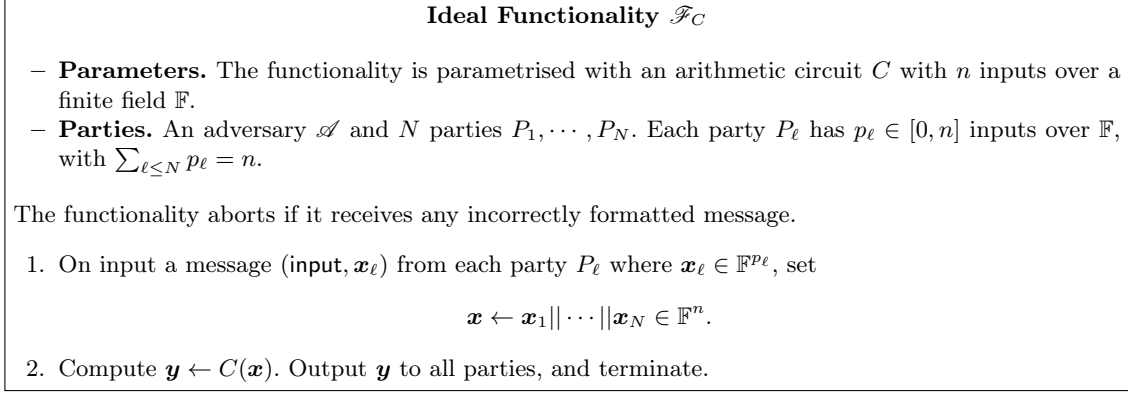


Fig. 2: Ideal functionality \mathcal{F}_C for securely evaluating an arithmetic circuit C among N parties.

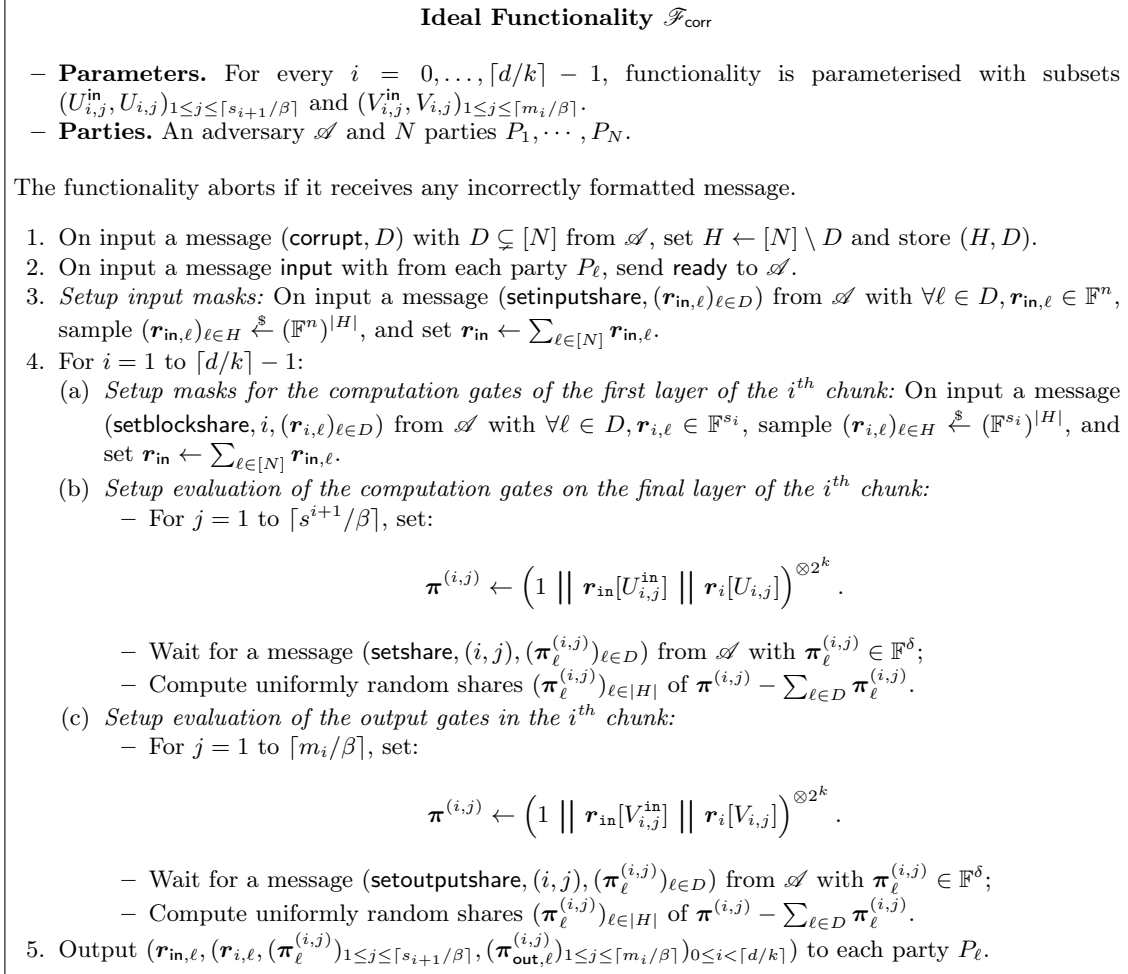


Fig. 3: Ideal corruptible functionality $\mathcal{F}_{\text{corr}}$ to deal out correlated randomness to the parties.

securely generate with low communication overhead. The rest of this section is dedicated to making Π_C explicit and to analysing its security.

In the sequel, we fix a layered arithmetic C with block decomposition:

$$(s_0, s_1, \dots, s_{t-1}), (m_0, \dots, m_{t-1}), \text{ and } (f_0, \dots, f_{t-1}).$$

We now proceed with the description of the protocol Π_C , which securely implements \mathcal{F}_C in the $\mathcal{F}_{\text{corr}}$ -hybrid model, with security against a static adversary passively corrupting at most $N-1$ parties. Fix parameters $\beta, k \in \mathbb{N}^*$. We let \mathbf{x}_ℓ denote the input vector of party P_ℓ over \mathbb{F} . We slightly abuse this notation and view each vector \mathbf{x}_ℓ as a length- n vector with zeroes at the positions where P_ℓ does not hold an input, so that the vectors \mathbf{x}_ℓ form additive shares of the input vector $\mathbf{x} = \sum_{\ell=1}^N \mathbf{x}_\ell$.

For each block B_i , we denote by \mathbf{u}^i the values on the nodes of the first layer of B_i , and by \mathbf{v}^i the values on all output nodes in B_i . Observe that by definition of f_i , we have $f_i(\mathbf{x}, \mathbf{u}^i) = (\mathbf{u}^{i+1}, \mathbf{v}^i)$. We further partition the outputs of f_i in subvectors $(\mathbf{u}_j^{i+1})_{1 \leq j \leq \lceil s_{i+1}/\beta \rceil}$ and $(\mathbf{v}_j^i)_{1 \leq j \leq \lceil m_i/\beta \rceil}$, such that each subvector has length at most β . Recall that each output of f_i depends on at most 2^k inputs; therefore, each subvector \mathbf{u}_j^{i+1} and \mathbf{v}_j^i depends on at most $\beta \cdot 2^k$ coordinates of $(\mathbf{x}, \mathbf{u}^i)$. For each subvector \mathbf{u}_j^{i+1} (resp. \mathbf{v}_j^i), we denote by $U_{i,j}^{\text{in}} \subset [n]$ (resp. $V_{i,j}^{\text{in}} \subset [n]$) the subset of coordinates of \mathbf{x} which influence \mathbf{u}_j^{i+1} (resp. \mathbf{v}_j^i), and by $U_{i,j} \subset [s_i]$ (resp. $V_{i,j} \subset [s_i]$) the subset of coordinates of \mathbf{u}^i which influence \mathbf{u}_j^{i+1} (resp. \mathbf{v}_j^i). Note that $|U_{i,j}^{\text{in}}| + |U_{i,j}| \leq \beta \cdot 2^k$ and $|V_{i,j}^{\text{in}}| + |V_{i,j}| \leq \beta \cdot 2^k$. This decomposition is illustrated in fig. 4.




Fig. 4: Decomposition of the i^{th} Block into low-degree Polynomials.

Initialisation.

- Each party P_ℓ sends input to $\mathcal{F}_{\text{corr}}$ and waits until it receives

$$(\mathbf{r}_{\text{in},\ell}, (\mathbf{r}_{i,\ell}, (\boldsymbol{\pi}_\ell^{(i,j)})_{1 \leq j \leq \lceil s_{i+1}/\beta \rceil}), (\boldsymbol{\pi}_{\text{out},\ell}^{(i,j)})_{1 \leq j \leq \lceil m_i/\beta \rceil})_{0 \leq i < t}.$$

- Each party P_ℓ broadcasts $\mathbf{z}_\ell^{\text{in}} \leftarrow \mathbf{x}_\ell + \mathbf{r}_{\text{in},\ell}$. All parties compute $\mathbf{z}^{\text{in}} \leftarrow \sum_\ell \mathbf{z}_\ell^{\text{in}} = \mathbf{x} + \mathbf{r}_{\text{in}}$.

i -th Block Evaluation. We now resume the description of the protocol, and assume that the protocol maintains the following invariant: at the beginning of the i -th block evaluation, each party P_ℓ holds an additive share \mathbf{u}_ℓ^i of the values \mathbf{u}^i on the first layer of the block. The block decomposition of f_i guarantees that each output of $f_i : \mathbb{F}^n \times \mathbb{F}^{s_i} \mapsto \mathbb{F}^{s_{i+1}} \times \mathbb{F}^{m_i}$ can be computed by a multivariate polynomial of degree at most 2^k . Let $\mathbf{Q}^{(i,j)}$ denote the vector of multivariate polynomials of degree at most 2^k such that $\mathbf{Q}^{(i,j)}(\mathbf{x}[U_{i,j}^{\text{in}}] \parallel \mathbf{u}^i[U_{i,j}]) = \mathbf{u}_j^{i+1} \in \mathbb{F}^{s_{i+1}}$. Similarly, let $\mathbf{Q}_{\text{out}}^{(i,j)}$ denote the vector of multivariate polynomials of degree at most 2^k such that $\mathbf{Q}_{\text{out}}^{(i,j)}(\mathbf{x}[V_{i,j}^{\text{in}}] \parallel \mathbf{u}^i[V_{i,j}]) = \mathbf{v}_j^i \in \mathbb{F}^{m_i}$.

- Each party P_ℓ broadcasts $\mathbf{z}_\ell^i \leftarrow \mathbf{u}_\ell^i + \mathbf{r}_{i,\ell}$. All parties compute $\mathbf{z}^i \leftarrow \sum_\ell \mathbf{z}_\ell^i = \mathbf{u}^i + \mathbf{r}_{i,\ell}$.
- Define the following vectors of multivariate polynomials:

$$\begin{aligned} \hat{\mathbf{Q}}^{(i,j)}(\mathbf{X}) &\leftarrow \mathbf{Q}^{(i,j)}(\mathbf{X} - (\mathbf{r}_{\text{in}}[U_{i,j}^{\text{in}}] \parallel \mathbf{r}_i[U_{i,j}])) \\ \hat{\mathbf{Q}}_{\text{out}}^{(i,j)}(\mathbf{X}) &\leftarrow \mathbf{Q}_{\text{out}}^{(i,j)}(\mathbf{X} - (\mathbf{r}_{\text{in}}[V_{i,j}^{\text{in}}] \parallel \mathbf{r}_i[V_{i,j}])). \end{aligned}$$

Observe that each coefficient of $\hat{\mathbf{Q}}^{(i,j)}(\mathbf{X})$ is itself a multivariate polynomial of degree at most 2^k in the coordinates of $(\mathbf{r}_{\text{in}}[U_{i,j}^{\text{in}}] \parallel \mathbf{r}_i[U_{i,j}])$. Therefore, all its coefficients can be computed as linear combinations of the coordinates of $\boldsymbol{\pi}^{(i,j)} = (1_{\mathbb{F}} \parallel \mathbf{r}_{\text{in}}[U_{i,j}^{\text{in}}] \parallel \mathbf{r}_i[U_{i,j}])^{\otimes 2^k}$. Similarly, all coefficients of $\hat{\mathbf{Q}}_{\text{out}}^{(i,j)}(\mathbf{X})$ can be computed as linear combinations of the coordinates of $\boldsymbol{\pi}_{\text{out}}^{(i,j)} = (1_{\mathbb{F}} \parallel \mathbf{r}_{\text{in}}[V_{i,j}^{\text{in}}] \parallel \mathbf{r}_i[V_{i,j}])^{\otimes 2^k}$.

- Each party P_ℓ computes shares $(\hat{\mathbf{Q}}_\ell^{(i,j)}(\mathbf{X}), \hat{\mathbf{Q}}_{\text{out},\ell}^{(i,j)}(\mathbf{X}))$ of $(\hat{\mathbf{Q}}^{(i,j)}(\mathbf{X}), \hat{\mathbf{Q}}_{\text{out}}^{(i,j)}(\mathbf{X}))$ from his shares $(\boldsymbol{\pi}_\ell^{(i,j)}, \boldsymbol{\pi}_{\text{out},\ell}^{(i,j)})$ of $(\boldsymbol{\pi}^{(i,j)}, \boldsymbol{\pi}_{\text{out}}^{(i,j)})$, and sets:

$$\begin{aligned} \mathbf{u}_\ell^{i+1} &\leftarrow \left(\hat{\mathbf{Q}}_\ell^{(i,j)}(\mathbf{z}^{\text{in}}[U_{i,j}^{\text{in}}], \mathbf{z}^i[U_{i,j}]) \right)_{1 \leq j \leq \lceil s_{i+1}/\beta \rceil} \\ \mathbf{v}_\ell^i &\leftarrow \left(\hat{\mathbf{Q}}_{\text{out},\ell}^{(i,j)}(\mathbf{z}^{\text{in}}[V_{i,j}^{\text{in}}], \mathbf{z}^i[V_{i,j}]) \right)_{1 \leq j \leq \lceil m_i/\beta \rceil}. \end{aligned}$$

Output Reconstruction. Each party P_ℓ sets and broadcasts \mathbf{v}_ℓ , where:

$$\mathbf{v}_\ell \leftarrow (\mathbf{v}_\ell^{(0,j)})_{j \leq q_0} \parallel (\mathbf{v}_\ell^{(1,j)})_{1 \leq j \leq \lceil m_1/\beta \rceil} \parallel \cdots \parallel (\mathbf{v}_\ell^{(\lceil d/k \rceil - 1, j)})_{1 \leq j \leq \lceil m_{\lceil d/k \rceil - 1}/\beta \rceil}.$$

All parties reconstruct and output $\mathbf{v} \leftarrow \sum_\ell \mathbf{v}_\ell$.

We now prove Theorem 12 by proving the above protocol satisfies the necessary requirements.

Proof. The efficiency of the protocol follows by inspection: the total length of all messages broadcast in Π_C is

$$N \cdot \left(n + \sum_{i=0}^{t-1} s_i + \sum_{i=0}^{t-1} m_i \right) \leq N \cdot \left(n + \frac{s}{k} + m \right).$$

We now analyse the security of Π_C . We describe in fig. 5 a simulator Sim which perfectly simulates an execution of Π_C .

Simulator Sim

Let $D \subsetneq [N]$ be the subset of statically corrupted parties, and $H = [N] \setminus D$ be the (nonempty) subset of honest parties.

- **Initialisation.**
 - Sim honestly simulates the functionality $\mathcal{F}_{\text{corr}}$ and stores the following values:
$$(\mathbf{r}_{\text{in},\ell}, (\mathbf{r}_{i,\ell}, (\boldsymbol{\pi}_\ell^{(i,j)})_{1 \leq j \leq \lceil s_{i+1}/\beta \rceil}), (\boldsymbol{\pi}_{\text{out},\ell}^{(i,j)})_{1 \leq j \leq \lceil m_i/\beta \rceil})_{0 \leq i < \lceil d/k \rceil}, \ell \in [N].$$
 - Sim broadcasts a random vector $\mathbf{z}_\ell^{\text{in}}$ on behalf of all honest parties. Let $\mathbf{z}^{\text{in}} \leftarrow \sum_{\ell \in H} \mathbf{z}_\ell^{\text{in}}$. For each $\ell \in D$, Sim extracts $\mathbf{x}_\ell \leftarrow \mathbf{z}_\ell^{\text{in}} - \mathbf{r}_{\text{in},\ell}$.
 - Sim sends $(\text{input}, \mathbf{x}_\ell)$ to \mathcal{F}_C on behalf of all corrupted parties, and receives an output \mathbf{v} .
- **i -th Block Evaluation.**

Sim broadcasts uniformly random $\mathbf{z}_\ell^i \stackrel{\$}{\leftarrow} \mathbb{F}^{s_i}$ on behalf of all honest parties, waits to receive $(\mathbf{z}_\ell^i)_{\ell \in D}$, and sets $\mathbf{z}^i \leftarrow \sum_{\ell \in [N]} \mathbf{z}_\ell^i$.
- **Output Phase.**

For each $\ell \in D$, Sim computes $\mathbf{v}_\ell^i \leftarrow \hat{\mathbf{Q}}_{\text{out},\ell}^{(i,j)}(\mathbf{z}^{\text{in}}[V_{i,j}^{\text{in}}], \mathbf{z}^i[V_{i,j}])$ and sets $\mathbf{v}_\ell \leftarrow \mathbf{v}_\ell^0 \parallel \cdots \parallel \mathbf{v}_\ell^{t-1}$. Eventually, Sim broadcasts uniformly random shares of $\mathbf{v} - \sum_{\ell \in D} \mathbf{v}_\ell$ on behalf of the honest parties.

Fig. 5: Simulator Sim for the sublinear protocol in the $\mathcal{F}_{\text{corr}}$ -hybrid model.

It remains to show that Sim perfectly simulates a run of the real protocol Π_C in the $\mathcal{F}_{\text{corr}}$ -hybrid model. But this follows almost immediately by inspection, as the view of all corrupted parties in Π_C contains exactly:

- the tuple $(\mathbf{r}_{\text{in},\ell}, (\mathbf{r}_{i,\ell}, (\boldsymbol{\pi}_\ell^{(i,j)})_{1 \leq j \leq \lceil s_{i+1}/\beta \rceil}), (\boldsymbol{\pi}_{\text{out},\ell}^{(i,j)})_{1 \leq j \leq \lceil m_i/\beta \rceil})_{0 \leq i < t}, \ell \in D$, which is chosen by the adversary;

- the vectors $(\mathbf{z}_\ell^{\text{in}} = \mathbf{x}_\ell + \mathbf{r}_{\text{in},\ell})_{\ell \in H}$, which are perfectly random since each $\mathbf{r}_{\text{in},\ell}$ is perfectly random, by definition of $\mathcal{F}_{\text{corr}}$;
- the vectors $(\mathbf{z}_\ell^i = \mathbf{u}_\ell^i + \mathbf{r}_\ell^i)_{\ell \in H}$, which are perfectly random since each \mathbf{r}_ℓ^i is perfectly random, by definition of $\mathcal{F}_{\text{corr}}$;
- the vectors $(\mathbf{v}_\ell)_{\ell \in H} = (\mathbf{v}_\ell^0 \parallel \cdots \parallel \mathbf{v}_\ell^{t-1})_{\ell \in H}$. By construction, and by definition of the $\hat{\mathbf{Q}}_{\mathbf{v},\ell}^{(i,j)}(\mathbf{X})$, the \mathbf{v}_ℓ satisfy:

$$\begin{aligned} \mathbf{v}_\ell^{(i,j)} &= \hat{\mathbf{Q}}_{\text{out},\ell}^{(i,j)}(\mathbf{w}[V_{i,j}^{\text{in}}], \mathbf{z}^i[V_{i,j}]) \\ &= \mathbf{Q}_{\text{out},\ell}^{(i,j)}((\mathbf{w}[V_{i,j}^{\text{in}}], \mathbf{z}^i[V_{i,j}]) - (\mathbf{r}[V_{i,j}^{\text{in}}], \mathbf{r}^i[V_{i,j}])) \\ &= \mathbf{Q}_{\text{out},\ell}^{(i,j)}(\mathbf{x}[V_{i,j}^{\text{in}}], \mathbf{u}^i[V_{i,j}]). \end{aligned}$$

Therefore, the \mathbf{v}_ℓ are uniformly random conditioned on:

$$\sum_{\ell \in H} \mathbf{v}_\ell^{(i,j)} = \mathbf{Q}_{\text{out}}^{(i,j)}(\mathbf{x}[V_{i,j}^{\text{in}}], \mathbf{u}^i[V_{i,j}]) - \sum_{\ell \in D} \mathbf{v}_\ell^{(i,j)}.$$

By definition of the $\mathbf{Q}_{\text{out}}^{(i,j)}(\mathbf{X})$, we have $\mathbf{Q}_{\text{out}}^{(i,j)}(\mathbf{x}[V_{i,j}^{\text{in}}], \mathbf{u}^i[V_{i,j}]) = \mathbf{v}^{(i,j)}$, where the $\mathbf{v}^{(i,j)}$ are such that $f_i(\mathbf{x}, \mathbf{u}^i) = (\mathbf{u}^{i+1}, (\mathbf{v}^{(i,j)})_{1 \leq j \leq \lceil m_i/\beta \rceil})$. Therefore, by definition of the f_i , the vectors \mathbf{v}_ℓ are uniformly random conditioned on:

$$\sum_{\ell \in H} \mathbf{v}_\ell = \mathbf{v} - \sum_{\ell \in D} \mathbf{v}_\ell, \text{ where } \mathbf{v} = C(\mathbf{x}).$$

This concludes the proof that Sim perfectly simulates Π_C . \square

5 Generating Correlated Randomness from LPN

In this section, we construct a protocol Π_{corr} , which implements the ideal functionality $\mathcal{F}_{\text{corr}}$ with small communication, under the quasi-polynomial LPN assumption. A very natural approach to realise a functionality that distributes correlated random coins using a small amount of communication is to rely on *pseudorandom correlation generators*, a primitive recently defined and constructed (for various types of correlations, and under a variety of assumptions) in [BCG⁺19b]. At a high level, [BCG⁺19b] suggests to distribute correlated randomness with the following approach:

- Use a generic secure computation protocol Π_{Gen} to distributively execute the PCG.Gen functionality of the pseudorandom correlation generator. Note that PCG.Gen outputs short seeds, much smaller than the correlated pseudo-random strings which can be stretched from these seeds. Therefore, Π_{Gen} can potentially have a relatively high communication overhead in its inputs and outputs, while maintaining the overall communication overhead of Π_{corr} small.
- Expand the distributively generated seeds locally using the Expand algorithm of the PCG. Each such string is guaranteed, by the security of the PCG, to be indistinguishable (from the viewpoint of the other parties) from a uniformly random string sampled conditioned on satisfying the target correlation with the expanded strings held by the other parties.

While this approach does not necessarily leads to a secure implementation of an ideal functionality generating correlated random coins, it was shown in [BCG⁺19b] (Theorem 19 in [BCG⁺19b]) that it provides a provably secure implementation for all *corruptible* ideal functionalities for distributing correlated random coins. Note that this property is satisfied by our functionality $\mathcal{F}_{\text{corr}}$. Our protocol Π_{corr} will follow this approach. We start by constructing a pseudorandom correlation generator for the type of correlated randomness produced by $\mathcal{F}_{\text{corr}}$, building upon an LPN-based construction of [BCG⁺19b].

5.1 Substrings Tensor Powers Correlations (stp)

We now describe our construction of a PCG for generating the type of correlated randomness produced by $\mathcal{F}_{\text{corr}}$. As all constructions of [BCG⁺19b], our construction will be restricted to the two-party setting; hence, we focus on $N = 2$ parties from now on. Abstracting out the unnecessary details, the

functionality $\mathcal{F}_{\text{corr}}$ does the following. It is parametrised with a vector length w , subsets $(S_i)_{1 \leq i \leq n_s} \in \binom{[w]}{\leq K}^{n_s}$, a tensor power parameter tpp , and generates shares of:

$$(\mathbf{r}, ((1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}})_{1 \leq i \leq n_s}), \text{ where } \mathbf{r} \in \mathbb{F}^w \text{ is random.}$$

We call \mathcal{C} the correlation generator associated with $\mathcal{F}_{\text{corr}}$, i.e. the PPT algorithm that, on input the security parameter in unary 1^λ , samples correlated random string as above (where the parameters (n_s, K, tpp) are functions of λ). It is straightforward to see that \mathcal{C} is a reverse-samplable correlation generator (see Definition 9), since it is an additive correlation: given any fixed share share_0 , a matching share can be reverse-sampled by sampling \mathbf{r} and setting $\text{share}_1 \leftarrow (\mathbf{r}, ((1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}})_{1 \leq i \leq n_s}) - \text{share}_0$. We call this type of correlated randomness a *subsets tensor powers* (stp). Below, we describe a pseudorandom correlation generator for such correlations.

5.2 Good Cover

Before we proceed with the description of a PCG to generate such correlations, we need to introduce a concept, that of a *good cover*. The notations in this subsection are completely self-contained, and may conflict with the parameters defined for the main protocol. In the course of our construction we will want to solve the following problem: given a vector \mathbf{v} of size n , a family $(S_i)_{i \in [t]} \in \mathcal{P}([n])^t$ of t (*short*) subsets of coordinates of \mathbf{v} , and a (*small*) bound $B > 0$, the problem is to find a family $(\mathbf{v}_j)_{j \in [M]}$ of some number m of size- K subvectors of \mathbf{v} such that:

1. The subvectors collectively cover \mathbf{v} ;
2. For each $i \in [t]$, there are at most B subvectors in $(\mathbf{v}_j)_{j \in [M]}$ whose coordinates intersect S_i .

We call such a family a *B-Good Cover* of $(\mathbf{v}, (S_i)_{i \in [t]})$. First of all we note that the values of the vectors and subvectors do not matter, so we will conflate them with sets and subsets (of coordinates) for simplicity, which leads to a more natural formulation.

Definition 13 (Good Cover – Set Formulation). *Let $n, B, K, t, q, M \in \mathbb{N}$ and $(S_i)_{i \in [t]} \in \binom{[n]}{\leq q}^t$ a family of t subsets of $[n]$ of size at most q each. A family $A = (\alpha^j)_{j \in [M]} \in \binom{[n]}{K}^M$ is a *B-Good Cover* of $(S_i)_{i \in [t]}$ if:*

1. A covers $[n]$: $\bigcup_{j=1}^M \alpha^j = [n]$
2. Each S_i intersects at most B elements of A : $\forall i \in [t], |\{j \in [M] : \alpha^j \cap S_i \neq \emptyset\}| \leq B$.

We abusively conflate the two views, where a good cover is just a family of subsets $A \in \binom{[n]}{K}^M$ and where the good cover is a family of sparse vectors—given by a set of coordinates and a short vector of values— $A \in (\binom{[n]}{K} \times \mathbb{F}^K)^M$.

Lemma 14 (Random Covers are Good Covers.). *Let $n, \kappa, \kappa' \in \mathbb{N} \setminus \{0, 1\}$, and $(S_i)_{i \in [t]} \in \binom{[n]}{\leq q}^t$ a family of t subsets of $[n]$ of size at most q each. Let $A = (\alpha^j)_{j \in [M]} \in \binom{[n]}{K}^M$ be a sequence of M i.i.d. uniform random size- K subsets of $[n]$, with $M = \kappa \cdot n \ln n / K$. Let $B \leftarrow \kappa' \kappa \cdot q \cdot \ln n$.*

*It holds that $A = (\alpha^j)_{j \in [M]}$ is a *B-Good Cover* of $(S_i)_{i \in [t]}$ with probability at least:*

$$1 - \frac{1}{n^{\kappa-1}} - \frac{t}{n^{(\kappa'-2)\kappa \cdot q/2}}.$$

The proof, which is obtained in a straightforward fashion by combining the union and Chernoff bounds, is given in

Proof.

1. The probability that $i \notin A$ (for any $i \in [n]$) is equal to $(1 - \binom{n-1}{K-1} / \binom{n}{K})^M = (1 - \frac{K}{n})^M = e^{M \cdot \ln(1-K/n)} \leq e^{-MK/n} = n^{-\kappa}$ (using $\forall x \geq -1, \ln(1+x) \leq x$), so by union bound, the probability that A does not cover $[n]$ is at most $n \cdot n^{-\kappa}$.

2. For each $i \in [q]$, we arbitrarily extend S_i to an \tilde{S}_i of size exactly q ; in particular $\forall i \in [q], S_i \subseteq \tilde{S}_i$ so $|\{j: \alpha^j \cap S_i \neq \emptyset\}| \leq |\{j: \alpha^j \cap \tilde{S}_i \neq \emptyset\}|$. For any $j \in [M]$, the indicator random variable of the event $\{\alpha^j \cap \tilde{S}_i \neq \emptyset\}$ follows a Bernoulli law with parameters $(M, \frac{q}{p})$, where $p = 1 - \binom{n-q}{K} / \binom{n}{K}$. Its expectancy is $\mu := \mathbb{E}(X) = M \cdot (1 - \binom{n-q}{K} / \binom{n}{K})$. Note that, using Bernoulli's inequality ($\forall r \in \mathbb{N}^*, \forall x \geq -1, (1+x)^r > (1+rx)$):

$$\begin{aligned} \mu &= M \cdot \left(1 - \frac{\prod_{j=0}^{K-1} (n-q-j)}{\prod_{j=0}^{K-1} (n-j)} \right) \leq m \cdot \left(1 - \frac{(n-2q)^K}{n^K} \right) \\ &= M \cdot (1 - (1 - 2q/n)^K) \leq M \frac{2qK}{n} = 2\kappa \cdot q \cdot \ln n. \end{aligned}$$

Therefore the probability that $X := |\{j: \alpha_j \in \tilde{S}_i\}| \leq B$ is, by (a looser version of the multiplicative form of) Chernoff's bound, at most:

$$\forall \delta \geq 0, \Pr[X > (1+\delta)\mu] \leq \exp\left(-\frac{\delta^2\mu}{2+\delta}\right).$$

Setting $\delta \leftarrow B/\mu - 1$ yields:

$$\begin{aligned} \Pr[X > B] &\leq \exp\left(-\frac{\delta^2\mu}{2+\delta}\right) \leq \exp(-\delta\mu/2) = \exp\left(-\frac{B-\mu}{2}\right) \\ &\leq n^{-(\kappa'-2)\kappa \cdot q/2}. \end{aligned}$$

3. The desired bound is then obtained by union bound. □

5.3 PCG for Subsets Tensor Powers (PCG_{stp})

We now proceed with the description of a pseudorandom correlation generator for subsets tensor powers.

PCG for Low-Degree Polynomials from [BCG⁺19b]. We start by recalling a natural variant of pseudorandom correlation generator of [BCG⁺19b, Section 6], which generates shares of $\mathbf{r}^{\otimes \text{tpp}}$, for a parameter tpp and a pseudorandom \mathbf{r} . It relies on the xLPN assumption with dimension n , number of samples $n' > n$, and a number λ of noisy coordinates. In our instantiation, we will typically consider $n' = O(n)$, e.g. $n' = 12n$; this corresponds to a particularly conservative variant of LPN with a very limited number of samples, and is equivalent to the hardness of decoding a random constant-rate linear code (which is known as the *syndrome decoding* problem). As discussed in Section 3, all known attacks on the syndrome decoding problem for constant-rate codes have complexity $2^{O(\lambda)}$. The PCG of [BCG⁺19b] is parametrised by integers $1^\lambda, n, n', \lambda, \text{tpp} \in \mathbb{N}$ (where $n' > n$), a field \mathbb{F} , and a random parity-check matrix $H_{n',n} \stackrel{\$}{\leftarrow} \mathbb{F}^{(n'-n) \times n'}$.

PCG for Degree-tpp Polynomial Correlations

PCG.Gen: On input 1^λ :

1. Pick a random λ -sparse vector $\mathbf{e} \stackrel{\$}{\leftarrow} \text{HW}_\lambda(\mathbb{F}^{n'})$. Note that $\mathbf{e}^{\otimes \text{tpp}} \in \text{HW}_{\lambda \text{tpp}}(\mathbb{F}^{(n')^{\text{tpp}}})$. Let $f: [(n')^{\text{tpp}}] \mapsto \mathbb{F}$ be the multi-point function with λ^{tpp} points, such that $f(i)$ returns the i -th coordinate of $\mathbf{e}^{\otimes \text{tpp}}$.
2. Compute $(K_0^{\text{fss}}, K_1^{\text{fss}}) \stackrel{\$}{\leftarrow} \text{MPFSS.Gen}(1^\lambda, f)$. Output $\mathbf{k}_0 \leftarrow (n, K_0^{\text{fss}})$ and $\mathbf{k}_1 \leftarrow (n, K_1^{\text{fss}})$.

PCG.Expand: On input $(\sigma, \mathbf{k}_\sigma)$, compute $\mathbf{v}_\sigma \leftarrow \text{MPFSS.FullEval}(\sigma, K_\sigma^{\text{fss}})$ in $\mathbb{F}^{(n')^{\text{tpp}}}$ and set $\mathbf{r}_\sigma \leftarrow H_{n',n}^{\otimes \text{tpp}} \cdot \mathbf{v}_\sigma$. Output \mathbf{r}_σ .

Fig. 6: PCG for Low-Degree Polynomials from [BCG⁺19b].

Correctness follows from the fact that $\mathbf{v}_0 + \mathbf{v}_1 = \mathbf{e}^{\otimes \text{tpp}}$ by the correctness of MPFSS, and $H_{n',n}^{\otimes \text{tpp}} \cdot \mathbf{e}^{\otimes \text{tpp}} = (H_{n',n} \cdot \mathbf{e})^{\otimes \text{tpp}}$ by multilinearity of the tensor product. Hence, denoting $\mathbf{r} = H_{n',n} \cdot \mathbf{e}$, it holds that $\mathbf{r}_0 + \mathbf{r}_1 = \mathbf{r}^{\otimes \text{tpp}}$. For security, we must show that the following distributions are indistinguishable for any $\sigma = 0, 1$:

$$\begin{aligned} & \{(k_\sigma, \mathbf{r}_{1-\sigma}) : (k_0, k_1) \xleftarrow{\$} \text{Gen}(1^\lambda), \mathbf{r}_{1-\sigma} \leftarrow \text{Expand}(1-\sigma, k_{1-\sigma})\} \\ & \stackrel{c}{\approx} \{(k_\sigma, \mathbf{r}_{1-\sigma}) : (k_0, k_1) \xleftarrow{\$} \text{Gen}(1^\lambda), \mathbf{r}_\sigma \leftarrow \text{Expand}(\sigma, k_\sigma), \mathbf{r} \xleftarrow{\$} \mathbb{F}^n, \\ & \quad \mathbf{r}_{1-\sigma} \leftarrow \mathbf{r}^{\otimes \text{tpp}} - \mathbf{r}_\sigma\} \end{aligned}$$

Proof. We sketch the analysis for the sake of completeness; the full proof is given in [BCG⁺19b]. Security is shown with the following sequence of hybrids: first generate $(k_\sigma, \mathbf{r}_{1-\sigma})$ as in the first distribution above. Then, generate $(k_\sigma, \mathbf{r}_{1-\sigma})$ as before, and generate an alternative key k'_σ solely from the parameters $(1^\lambda, \mathbb{F}, n, n', t, \text{tpp})$, using the simulator of the MPFSS. Output $(k'_\sigma, \mathbf{r}_{1-\sigma})$; under the security of the MPFSS, this distribution is indistinguishable from the previous one. Note that k'_σ does not depend anymore on the noise vector \mathbf{e} . In the next hybrid, generate $\mathbf{r} \xleftarrow{\$} H_{n',n} \cdot \mathbf{e}$ and set $\mathbf{r}_{1-\sigma} \leftarrow \mathbf{r}^{\otimes \text{tpp}} - \text{Expand}(\sigma, k_\sigma)$; this game is perfectly indistinguishable from the previous one. Finally, replace $\mathbf{r} \xleftarrow{\$} H_{n',n} \cdot \mathbf{e}$ by $\mathbf{r} \xleftarrow{\$} \mathbb{F}^n$; under the LPN assumption, this last game (which correspond exactly to the second distribution) is computationally indistinguishable from the previous one, and security follows. \square

Our New PCG. We now describe a variant of the above PCG, tailored to computing the tensor powers of many short subsets. The PCG is parametrised by $(S_i)_{i \in [K]} \in \binom{[w]}{\leq K}^{n_s}$, n_s subsets of at most K indices taken from $[w]$. We assume for simplicity, but morally without loss of generality⁶, that $\bigcup_{i=1}^{n_s} S_i = [w]$. Our goal is for the parties to obtain shares of some pseudorandom vector $\mathbf{r} \in \mathbb{F}^w$ as well as shares of $(1 \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}} \in \mathbb{F}^{w \cdot \text{tpp}}$ for each $i \in [n_s]$.

We start by generating a B -good cover (for some integer B) of the $(S_i)_i$ of the form $(\alpha_j, \mathbf{r}_j)_{j \in [m]} \in \left(\binom{[w]}{\theta} \times \mathbb{F}^\theta\right)^m$ where each \mathbf{r}_j is pseudorandom. We generate each of the m pseudorandom masks \mathbf{r}_j using a different instance of xLPN, *i.e.* $\mathbf{r}_j \leftarrow H_j \cdot \mathbf{e}_j$, where $\mathbf{e}_j \in \mathbb{F}^{\theta'}$ is λ -sparse and $H_j \xleftarrow{\$} \mathbb{F}^{\theta \times \theta'}$ for some $\theta' = O(\theta)$. For each S_i , we denote $I_i := \{j \in [m] : \alpha_j \cap S_i \neq \emptyset\} = \{j_1, \dots, j_{|I_i|}\}$ the set of the indices of the masks which ‘intersect’ with S_i . Note that $\forall i \in [n_s], |I_i| \leq B$ by definition of a B -good cover. We can now proceed with our main goal: generating shares of a subsets tensor powers correlation.

We define $\mathbf{r} := \sum_{j=1}^m f_{\alpha_j, \mathbf{r}_j} \in \mathbb{F}^w$, where $f_{\alpha_j, \mathbf{r}_j} \in \mathbb{F}^w$ is the sparse vector defined by $(f_{\alpha_j, \mathbf{r}_j})_{\alpha_j} = \mathbf{r}_j$ (and which is equal to $0_{\mathbb{F}}$ on $[w] \setminus \alpha_j$). Since $\bigcup_{i=1}^{n_s} S_i = [w]$ and each of the \mathbf{r}_j is pseudorandom, \mathbf{r} is also pseudorandom.

Note that for any given $i \in [n_s]$, $(1_{\mathbb{F}} \parallel \mathbf{r}[S_i])$ is a subvector of the vector $\tilde{\mathbf{r}}_i$ obtained by multiplying the block-diagonal matrix $H'_i = \text{Diag}((1_{\mathbb{F}}), H_{j_1}, \dots, H_{j_{|I_i|}})$ with the vector $\mathbf{e}'_i = (1_{\mathbb{F}} \parallel \mathbf{e}_{j_1} \parallel \dots \parallel \mathbf{e}_{j_{|I_i|}})$. Therefore for any tensor power tpp (*i.e.* the degree of the polynomial correlation), $\tilde{\mathbf{r}}_i^{\otimes \text{tpp}} = (H'_i \cdot \mathbf{e}'_i)^{\otimes \text{tpp}} = (H'_i)^{\otimes \text{tpp}} \cdot (\mathbf{e}'_i)^{\otimes \text{tpp}}$. If the parties use an MPFSS scheme to generate small seeds which expand to $(\mathbf{e}'_i)^{\otimes \text{tpp}}$, they can then locally obtain shares of $\tilde{\mathbf{r}}_i^{\otimes \text{tpp}}$ (since $(H'_i)^{\otimes \text{tpp}}$ is public), and therefore of $(1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}}$. From all these shares of all the $(1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}}, i \in [n_s]$ the parties can locally extract shares of all the $\mathbf{r}[S_i]$ and thence shares of \mathbf{r} (since $\bigcup_{i=1}^{n_s} S_i = [w]$). The protocol is given in Figure 7.

Theorem 15. *Let $w > 0$, and $(S_i)_{i \in [n_s]}$ a list of n_s subsets of $[w]$. Let B, θ' such that there exists a B -good cover of $(S_i)_{i \in [n_s]}$ comprised of size- θ' vectors, and let $\theta < \theta'$. Assume that the \mathbb{F} -xLPN(θ, θ', λ) assumption holds, and that MPFSS is a secure multi-point function secret-sharing scheme for the family of $(1 + \mu \cdot \lambda)^{\text{tpp}}$ -point functions from $[(1 + \mu \cdot \theta')^{\text{tpp}}]$ to \mathbb{F} for all $\mu \in [B]$. Then PCG_{stp} is a secure pseudorandom correlation generator, which generates pseudorandom shares of a subsets tensor powers correlation $(\mathbf{r}, ((1_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}})_{1 \leq i \leq n_s})$ where $\mathbf{r} \in \mathbb{F}^w$.*

⁶ If $\bigcup_{i=1}^{n_s} S_i \neq \emptyset$, and with the notations of the rest of the section, the vector \mathbf{r} we generate is equal to $0_{\mathbb{F}}$ on $[w] \setminus \bigcup_{i=1}^{n_s} S_i$, hence not pseudorandom. However, we can simply have the parties generate another mask $\mathbf{r}' = H' \cdot \mathbf{e}'$, pseudorandom under xLPN, to cover $[w] \setminus \bigcup_{i=1}^{n_s} S_i$. Since the parties do not need shares of $(\mathbf{r}')^{\otimes \text{tpp}}$, the communication complexity of generating the λ -sparse \mathbf{e}' using an MPFSS is not an issue.

Pseudorandom Correlation Generator PCG_{stp}

Parameters: $w, \text{tpp}, \lambda \in \mathbb{N}$ and $(S_i)_{1 \leq i \leq n_s} \subseteq [w]^{n_s}$.

Gen: On input 1^λ :

1. Generate a family of subsets $(\alpha_j)_{1 \leq j \leq m} \in \binom{[m]}{\theta'}$ which form a B -good cover of the $(S_i)_{i \in [n_s]}$ (when the α_j are paired with length- θ' vectors in $\mathbb{F}^{\theta'}$), and contracting matrices^a $(H_j)_{j \in [m]} \in (\mathbb{F}^{\theta \times \theta'})^m$.
2. Pick m random λ -sparse vectors $\mathbf{e}_j \xleftarrow{\$} \text{HW}_\lambda(\mathbb{F}^{\theta'})$, $j \in [m]$ and define:

$$\mathbf{r}_j \leftarrow H_j \cdot \mathbf{e}_j^\top, \text{ for all } j \in [m].$$

3. For each $i = 1 \dots n_s$:

(a) Denoting $I_i := \{j \in [m] : \alpha_j \cap S_i \neq \emptyset\} = \{j_1, \dots, j_{m_i}\}$ (with $m_i \leq B$), set:

$$\tilde{\mathbf{r}}_i \leftarrow (\mathbf{1}_{\mathbb{F}} \parallel H_{j_1} \cdot \mathbf{e}_{j_1}^\top \parallel \dots \parallel H_{j_{m_i}} \cdot \mathbf{e}_{j_{m_i}}^\top)^\top.$$

(b) Let $f_i : [(1 + m_i \cdot \theta')^{\text{tpp}}] \rightarrow \mathbb{F}$ be the multi-point function with $(1 + m_i \cdot \lambda)^{\text{tpp}}$ points, such that $f_i(x) = (\mathbf{1}_{\mathbb{F}} \parallel \mathbf{e}_{j_1} \parallel \dots \parallel \mathbf{e}_{j_{m_i}})^{\otimes \text{tpp}}[x]$. Compute $(K_{i,0}^{\text{fss}}, K_{i,1}^{\text{fss}}) \xleftarrow{\$} \text{MPFSS.Gen}(1^\lambda, f_i)$.

4. Output $\mathbf{k}_0 \leftarrow (w, (K_{i,0}^{\text{fss}})_{i \leq n_s})$ and $\mathbf{k}_1 \leftarrow (w, (K_{i,1}^{\text{fss}})_{i \leq n_s})$.

Expand: On input $(\sigma, \mathbf{k}_\sigma)$, parse \mathbf{k}_σ as $(w, (K_{i,\sigma}^{\text{fss}})_{i \leq n_s})$.

1. For each $i = 1 \dots n_s$:

Set $H'_i \leftarrow \text{Diag}(\mathbf{1}_{\mathbb{F}}, H_{j_1}, \dots, H_{j_{m_i}})$, compute

$$\mathbf{v}_{i,\sigma} \leftarrow \text{MPFSS.FullEval}(\sigma, K_{i,\sigma}^{\text{fss}}) \in \mathbb{F}^{(1+m_i \cdot \lambda)^{\text{tpp}}}$$

and set $\mathbf{y}_\sigma \leftarrow ((H'_i)^{\otimes \text{tpp}} \cdot \mathbf{v}_\sigma)_{1 \leq i \leq n_s}$.

2. Extract from \mathbf{y}_σ the appropriate linear combinations of its elements corresponding to a share of $(\mathbf{r}, ((\mathbf{1}_{\mathbb{F}} \parallel \mathbf{r}[S_i])^{\otimes \text{tpp}})_{i \in [n_s]})$. // If there are several ways to do so, it must be consistent across $\sigma \in \{0, 1\}$.

^a Implicitly, the H_j are supposed to be ‘suitably chosen’ for xLPN to be presumed hard, e.g. that they were randomly and independently sampled.

Fig. 7: Pseudorandom correlation generator PCG_{stp} for generating pseudorandom instances of the subsets tensor powers correlation over a field \mathbb{F} .

- Communication: If the MPFSS seeds have size $O[\lambda \cdot (1+B\lambda)^{\text{tpp}} \cdot \log((1+B\theta')^{\text{tpp}})]$ and MPFSS.FullEval can be computed with $O((1+B\lambda)^{\text{tpp}} \cdot (1+B\theta')^{\text{tpp}} \cdot \frac{\log|\mathbb{F}|}{\lambda})$ invocations of a pseudorandom generator $\text{PRG} : \{0,1\}^\lambda \mapsto \{0,1\}^{2\lambda+2}$, then $\text{PCG}_{\text{stp}}.\text{Gen}$ outputs seeds of size:

$$|k_\sigma| = O(n_s \cdot \lambda \cdot (1+B\lambda)^{\text{tpp}} \cdot \log((1+B\theta')^{\text{tpp}})).$$

- Computation: The computational complexity of $\text{PCG}_{\text{stp}}.\text{Expand}$ is predominantly that of $O(n_s \cdot (1+B\lambda)^{\text{tpp}} \cdot (1+B\theta') \cdot \frac{\log|\mathbb{F}|}{\lambda})$ invocations of a PRG, plus n_s matrix-vector products with a matrix of dimensions $(1+B\theta)^{\text{tpp}} \times (1+B\theta')^{\text{tpp}}$ which requires at most $O(n_s \cdot (B\theta)^{\text{tpp}} \cdot (B\theta')^{\text{tpp}}) \subseteq O(n_s \cdot (B\theta')^{2 \cdot \text{tpp}})$ arithmetic operations over \mathbb{F} .

Proof. Let us first show correctness of our candidate PCG. By correctness of the MPFSS, $\mathbf{v}_{i,0} + \mathbf{v}_{i,1} = (1_{\mathbb{F}} \| (\mathbf{e}_\ell)_{\ell \in I_i})^\top$ for every $i \in [n_s]$. Therefore,

$$\begin{aligned} \mathbf{y}_0 + \mathbf{y}_1 &= ((H'_i)^{\otimes \text{tpp}} \cdot \mathbf{v}_0 + (H'_i)^{\otimes \text{tpp}} \cdot \mathbf{v}_1)_{1 \leq i \leq n_s} \\ &= ((H'_i)^{\otimes \text{tpp}} \cdot ((1_{\mathbb{F}} \| (\mathbf{e}_\ell)_{\ell \in I_i})^\top))_{1 \leq i \leq n_s} = (\tilde{\mathbf{r}}_i^{\otimes \text{tpp}})_{1 \leq i \leq n_s}. \end{aligned}$$

$\mathbf{r}[S_i] = \left(\sum_{j=1}^m f_{\alpha_j, \mathbf{r}_j} \right) [S_i] = \left(\sum_{j \in I_i} f_{\alpha_j, \mathbf{r}_j} \right) [S_i]$, where $f_{\alpha_j, \mathbf{r}_j}$ is the sparse vector obtained by spreading the vector \mathbf{r}_j over the coordinates in the ordered set α_j . It follows that $(1_{\mathbb{F}} \| \mathbf{r}[S_i])^{\otimes \text{tpp}}$ can be extracted from $\tilde{\mathbf{r}}_j$ (since any degree- $(\leq \text{tpp})$ polynomial in $\mathbf{a} + \mathbf{b}$ is also a degree- $(\leq \text{tpp})$ polynomial in $\mathbf{a} \| \mathbf{b}$).

Security follows exactly by the same sequence of hybrids as in the previous analysis. We show the following indistinguishability:

$$\begin{aligned} &\{ (k_\sigma, \mathbf{r}_{1-\sigma}) : (k_0, k_1) \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda), \mathbf{r}_{1-\sigma} \leftarrow \text{Expand}(1-\sigma, k_{1-\sigma}) \} \\ &\stackrel{c}{\approx} \{ (k_\sigma, \mathbf{r}_{1-\sigma}) : (k_0, k_1) \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda), \mathbf{r}_\sigma \leftarrow \text{Expand}(\sigma, k_\sigma), \\ &\quad \mathbf{r}' \stackrel{\$}{\leftarrow} \mathbb{F}^m, \mathbf{r}_{1-\sigma} \leftarrow ((1_{\mathbb{F}} \| \mathbf{r}')^{\otimes \text{tpp}})_{i \leq n_s} - \mathbf{r}_\sigma \}. \end{aligned}$$

We first switch all the $K_{i,\sigma}^{\text{fss}}$ to simulated keys using the security of the MPFSS, then replace the $H_i \cdot \mathbf{e}_i^\top$ by random vectors, applying m times the security of LPN, once for each replacement. The sum of sparse random vectors which form a good cover is itself a random vector (by item 1 in definition 13), therefore the resulting distribution is exactly the second distribution above, hence security follows.

Finally, the efficiency claims can be read directly from the construction, and follows from the fact that Gen consists of n_s MPFSS seeds, and the cost of Expand is dominated by n_s calls to MPFSS.FullEval and n_s matrix-vector products. \square

5.4 Instantiating the MPFSS

Theorem 15 assumes the existence of an MPFSS scheme MPFSS for the family of all $(1+\mu \cdot \lambda)^{\text{tpp}}$ -point functions from $[(1+\mu \cdot \theta')^{\text{tpp}}]$ to \mathbb{F} for some $\mu \in [B]$ (or, equivalently, an MPFSS for each μ which can then all be combined into one scheme), with the following efficiency guarantees: $\text{MPFSS.Gen}(1^\lambda)$ outputs seeds of size $O((1+B\lambda)^{\text{tpp}} \cdot \lambda \cdot \log((1+B\theta')^{\text{tpp}}))$, and MPFSS.FullEval can be computed with $O((1+B\lambda)^{\text{tpp}} \cdot (1+B\theta')^{\text{tpp}} \cdot \frac{\log|\mathbb{F}|}{\lambda})$ invocations of a pseudorandom generator $\text{PRG} : \{0,1\}^\lambda \mapsto \{0,1\}^{2\lambda+2}$. The works of [BGI16b, BCGI18] provides exactly such a construction, which makes a black box use of any pseudorandom generator $\text{PRG} : \{0,1\}^\lambda \mapsto \{0,1\}^{2\lambda+2}$.

We instantiate the PRG using the LPN-based construction of [BKW03], which we now recall. Fix some (constant) noise rate ε , a random matrix $A \stackrel{\$}{\leftarrow} \mathbb{F}_2^{n_1 \times n_2}$. Given a random bitstring $r \in \{0,1\}^{n_2+n_1 \cdot h(\varepsilon)}$, where h is the binary entropy function, define $\mathbf{s} = \mathbf{s}(r) \in \mathbb{F}_2^{n_2}$ to be the n_2 first bits of r , and use the remaining $n_1 \cdot h(\varepsilon)$ bits to sample a random vector $\mathbf{e}(r)$ from $\text{Ber}_\varepsilon(\mathbb{F}_2)^{n_1}$ (it is well known that this distribution can be sampled using roughly $n_1 \cdot h(\varepsilon)$ bits of randomness). Define the pseudorandom generator $\text{PRG} : \{0,1\}^\lambda \mapsto \{0,1\}^{2\lambda+2}$ as $\text{PRG}(r) = A \cdot \mathbf{s}(r) + \mathbf{e}(r) \in \mathbb{F}_2^{n_1}$. Security follows from the \mathbb{F}_2 -LPN(n_2, n_1, ε) assumption, and this PRG stretches $\lambda = n_2 + n_1 \cdot h(\varepsilon)$ bits to $n_1 = 2\lambda + 2$ bits when $n_2 = n_1 \cdot (1/2 - h(\varepsilon)) - 1$. Hence, given the security parameter λ , security follows from the \mathbb{F}_2 -LPN($\lambda \cdot (1 - 2h(\varepsilon)) - 2h(\varepsilon), 2\lambda + 2, \varepsilon$) assumption for any constant ε ; for example, setting $\varepsilon = 1/8$, this assumption is implied by the \mathbb{F}_2 -LPN($\lambda/4, 3\lambda, 1/8$) assumption. The cost of evaluating PRG is dominated by the matrix-vector product $A \cdot \mathbf{s}$, which requires at most $n_1 n_2 = O(\lambda^2)$ arithmetic operations.

5.5 Securely Distributing MPFSS.Gen and Π_{stp}

The seeds of the MPFSS scheme of [BCGI18] can be securely generated by using parallel instances of a generic secure computation protocols to securely evaluate the above PRG. Using GMW to instantiate the generic protocol, we have:

Corollary 16. *There exists a semi-honest secure two-party protocol Π_{MPFSS} which distributes the seeds of a multi-point function secret-sharing scheme MPFSS for the family of t' -point functions from $[(1 + B\theta')^{\text{tpp}}]$ to \mathbb{F} , using $O(t' \cdot \nu \cdot \lambda^2)$ calls to an ideal oblivious transfer functionality, where $\nu = \log((1 + B\theta')^{\text{tpp}})$ and $t' = (1 + B\lambda)^{\text{tpp}}$, with an additional communication of $O(t' \cdot \nu \cdot \lambda^2)$ bits, and total computation polynomial in $t' \cdot \nu \cdot \lambda$.*

Proof. Let $t' \leftarrow (1 + B\lambda)^{\text{tpp}}$. The MPFSS scheme MPFSS for the family of t' -point functions from $[(1 + B\theta')^{\text{tpp}}]$ to \mathbb{F} of [BCGI18] is constructed using t' independent instances of a single-point function secret sharing scheme (also called *distributed point function* [GI14]): any t' -point function $f : [(1 + B\theta')^{\text{tpp}}] \mapsto \mathbb{F}$ can be written as the sum $\sum_{i=1}^{t'} f_i$ of point functions $f_i : [(1 + B\theta')^{\text{tpp}}] \mapsto \mathbb{F}$ which evaluate to $0_{\mathbb{F}}$ everywhere, except on a single entry j_i where they take the value $f(j_i)$ (the j_i being the entries on which f does not evaluate to $0_{\mathbb{F}}$). Given a distributed point function (Gen, Eval), the construction and its analysis are straightforward:

- MPFSS.Gen : On input $(1^\lambda, f)$, decompose f as $\sum_{i=1}^{t'} f_i$ as above, and output $(\text{Gen}(1^\lambda, f_i))_{1 \leq i \leq t'}$.
- MPFSS.Eval : On input $(\sigma, (K_{i,\sigma})_{i \in [t']}, x)$, output $y_\sigma \leftarrow \sum_{i=1}^{t'} \text{Eval}(K_{i,\sigma}, f_i, x)$.

Therefore, securely distributing MPFSS seeds reduces to t' invocations of a secure protocol for distributing the seeds of a distributed point function (DPF). The DPF.Gen construction of [BGI16b] for point functions over the domain $[(1 + B\theta')^{\text{tpp}}]$ works as follows: the two output keys are $K_0 = (s_0^{(0)}, cw_1, \dots, cw_{\nu+1})$ and $K_1 = (s_1^{(0)}, cw_1, \dots, cw_{\nu+1})$ where $s_0^{(0)}, s_1^{(0)}$ are two random seeds for the PRG and $\nu = \log((1 + B\theta')^{\text{tpp}})$. Gen proceeds in $\nu + 1$ steps. In the i -th step it expands $s_0^{(i-1)}$ and $s_1^{(i-1)}$ by using one PRG invocation for each seed and obtains $s_0^{(i)}, s_1^{(i)}$, and cw_i . In the final step the algorithm computes $cw_{\nu+1}$ as a function of the expanded seeds and the target value.

Securely Generating DPF Seeds. We recall that the well known GMW protocol [GMW87b] allows two parties to securely evaluate any circuit of size s (in the semi-honest model) using $O(s)$ calls to an oblivious transfer functionality, and $O(s)$ additional bits of communication. Using GMW for distributing the Gen procedure of the DPF over a field \mathbb{F} , the communication and computation of the protocol are dominated by two factors: $O(\lambda)$ oblivious transfers for a seed and location of the designated point and by $O(\nu)$ secure evaluations of the PRG. Since evaluating the PRG can be done using $O(\lambda^2)$ arithmetic operations over \mathbb{F}_2 , it can be generically computed using $O(\lambda^2)$ calls to an oblivious transfer functionality, and $O(\lambda^2)$ additional bits of communication. Hence the following lemma:

Lemma 17. *There exists a semi-honest secure two-party protocol Π_{DPF} which distributes the seeds of a distributed point function DPF for the family of point functions from $[(1 + B\theta')^{\text{tpp}}]$ to \mathbb{F} , using $O(\nu \cdot \lambda^2)$ calls to an ideal oblivious transfer functionality, where $\nu = \log((1 + B\theta')^{\text{tpp}})$, with an additional communication of $O(\nu \cdot \lambda^2)$ bits, and total computation polynomial in $\nu \cdot \lambda$.*

□

As a direct corollary of Corollary 16, since the seeds of PCG_{stp} contain exactly n_s independent MPFSS seeds, we have:

Corollary 18. *There exists a semi-honest secure two-party protocol Π_{stp} which distributes the seeds of the pseudorandom correlation generator PCG_{stp} represented on Figure 7, using $O(n_s \cdot t' \cdot \nu \cdot \lambda^2)$ calls to an ideal oblivious transfer functionality, where $\nu = \log((B\theta' + 1)^{\text{tpp}})$ and $t' = (1 + B\lambda)^{\text{tpp}}$, with an additional communication of $O(n_s \cdot t' \cdot \nu \cdot \lambda^2)$ bits, and total computation $O(n_s \cdot \text{poly}(t' \cdot \nu \cdot \lambda))$.*

Instantiating the oblivious transfer. To execute the GMW protocol, we need an oblivious transfer. Under the \mathbb{F}_2 -LPN($\lambda, O(\lambda), 1/\lambda^\delta$) assumption (δ is any small constant), there exists oblivious transfers (with simulation security) with $\text{poly}(\lambda)$ communication and computation; see for example [DGH⁺20].

Constructing Π_{corr} . The work of [BCG⁺19b] shows that any corruptible functionality distributing the output of a correlation generator \mathcal{C} can be secure instantiated using any semi-honest secure two-party protocol Π for distributing the Gen procedure of a PCG for \mathcal{C} , with the same communication as Π , and with computational complexity dominated by the computational complexity of Π plus the computational complexity for computing the PCG.Expand procedure. Therefore, using their result together with our protocol Π_{stp} for generating the seeds of a PCG for subsets tensor powers correlation allows to securely instantiate $\mathcal{F}_{\text{corr}}$ (with $N = 2$).

Recall that the computation of $\text{PCG}_{\text{stp}}.\text{Expand}$ is dominated by $O(n_s \cdot (1+B\lambda)^{\text{tpp}} \cdot (1+B\theta')^{\text{tpp}} \cdot \frac{\log|\mathbb{F}|}{\lambda})$ invocations of a PRG – which requires at most $O(\lambda^2 \cdot n_s \cdot (1+B\lambda)^{\text{tpp}} \cdot (1+B\theta')^{\text{tpp}} \cdot \frac{\log|\mathbb{F}|}{\lambda})$ operations over \mathbb{F}_2 using the simple LPN-based PRG from [BKW03] –, plus an additional $O(n_s \cdot (1+B\theta)^{\text{tpp}} \cdot (1+B\theta')^{\text{tpp}})$ arithmetic operations over \mathbb{F} . Since each operation over \mathbb{F} can be computed with $O(\log|\mathbb{F}|)^2$ boolean operations, combining the two, we get computation $O(\lambda \cdot n_s \cdot (1+B\theta)^{\text{tpp}} \cdot (1+B\theta')^{\text{tpp}} \cdot (\log|\mathbb{F}|)^2)$.

All that remains is for the parties to generate the necessary material for PCG_{stp} : m random $\mathbb{F}^{\theta \times \theta'}$ matrices and m size- θ' subsets of $[w]$. At its core, this is just a matter for the parties to generate and hold the same $m \cdot (\theta \cdot \theta' \cdot \log|\mathbb{F}| + \log(\frac{w}{\theta'}))$ (pseudo)-random bits. This can be achieved by having one party sample a seed of size λ , send it to the other, and both parties can expand it locally by calling the length-doubling PRG from [BKW03] (and used above) $m \cdot \theta' \cdot (\theta \cdot \log|\mathbb{F}| + \log w)/\lambda$ times (in a GGM tree-like approach). This requires λ bits of communication and $O(m \cdot \theta' \cdot (\theta \cdot \log|\mathbb{F}| + \log w) \cdot \lambda)$ bits of local computation. This is summarised in an intermediate theorem, Theorem 19 below.

Theorem 19. *Assume the $\mathbb{F}\text{-xLPN}(\theta, \theta' - \theta, \lambda/\theta')$ and $\mathbb{F}_2\text{-LPN}(\lambda, O(\lambda), 1/\lambda^\delta)$ – where δ is any small enough constant – assumptions hold. Then there exists a semi-honest secure two-party protocol Π_{stp} which securely generates a subsets tensor powers correlation for subsets $(S_i)_{i \in [n_s]}$ of $[w]$ and for which there exists a B -good cover comprised of m size- θ' masks, using the following resources:*

– Communication:

$$O\left(n_s \cdot \text{poly}(\lambda) \cdot (1+B\lambda)^{\text{tpp}} \cdot \log(1+B\theta')^{\text{tpp}}\right).$$

– Computation:

$$O\left(\lambda \cdot n_s \cdot (1+B\theta)^{\text{tpp}} \cdot (1+B\theta')^{\text{tpp}} \cdot (\log|\mathbb{F}|)^2 + n_s \cdot \text{poly}((1+B\lambda)^{\text{tpp}} \cdot \log(1+B\theta')^{\text{tpp}} \cdot \lambda) + m \cdot \theta' \cdot (\theta \cdot \log|\mathbb{F}| + \log w) \cdot \lambda\right).$$

Wrapping up, using Π_{stp} with an appropriate good cover suffices to construct a protocol Π_{corr} for securely implementing the functionality $\mathcal{F}_{\text{corr}}$. For each $i = 1 \dots \lceil d/k \rceil - 1$, the parties need to generate a B -good cover of the $((U_{i,j}^{\text{in}})_{j \in \lceil [s^{i+1}/\beta] \rceil}, (U_{i,j})_{j \in \lceil [s^{i+1}/\beta] \rceil}, (V_{i,j}^{\text{in}})_{j \in \lceil [m^i/\beta] \rceil}, (V_{i,j})_{j \in \lceil [m^i/\beta] \rceil})$ seen as subsets of $[n + s_i]$. A way to do so is to generate:

- a $B/2$ -good cover A_{in} of $((U_{i,j}^{\text{in}})_{j \in \lceil [s^{i+1}/\beta] \rceil}, (V_{i,j}^{\text{in}})_{j \in \lceil [m^i/\beta] \rceil})_{1 \leq i < \lceil d/k \rceil}$ seen as subsets of $[n]$ comprised of M_{in} size- θ' masks;
- for each $i = 1 \dots \lceil d/k \rceil - 1$, a $B/2$ -good cover A_i of $((U_{i,j})_{j \in \lceil [s^{i+1}/\beta] \rceil}, (V_{i,j})_{j \in \lceil [m^i/\beta] \rceil})$ seen as subsets of $[s_i]$ comprised of M_i size- θ' masks.

Let $\kappa, \kappa', \kappa_{\text{in}}$ be correctness parameters. We set $M_{\text{in}} \leftarrow \kappa_{\text{in}} \cdot n \cdot \log n$, $M_i \leftarrow \kappa \cdot s_i \cdot \log s_i$ (which is upper-bounded by $M := \kappa \cdot s \cdot \log s$), and $B \leftarrow 2\kappa' \cdot \kappa \cdot \ln s$. The probability $p = p(\kappa_{\text{in}}, \kappa, \kappa')$ all the above conditions are satisfied is then, by union-bound (and with $s_i/\beta \leq s/\beta$), at least:

$$1 - \left(\frac{1}{n^{\kappa_{\text{in}}-1}} + \frac{(s/k)/\beta}{n^{(\kappa'-2) \cdot \kappa_{\text{in}} \cdot \theta/2}} \right) - \lceil d/k \rceil \cdot \left(\frac{1}{(s/k)^{\kappa-1}} + \frac{s/\beta}{(s/k)^{(\kappa'-2) \cdot \kappa \cdot \theta/2}} \right).$$

Wrapping-up, we get the following parametrised theorem.

Theorem 20. *Assume the $\mathbb{F}\text{-xLPN}(\theta, \theta' - \theta, \lambda/\theta')$ and $\mathbb{F}_2\text{-LPN}(\lambda, O(\lambda), 1/\lambda^\delta)$ – where δ is any small enough constant – assumptions hold. Then there exists a probabilistic semi-honest secure two-party protocol Π_{corr} which securely implements the functionality $\mathcal{F}_{\text{corr}}$ given on Figure 3 with success probability:*

$$p^* = 1 - \left(\frac{1}{n^{\kappa_{\text{in}}-1}} + \frac{(s/k)/\beta}{n^{(\kappa'-2) \cdot \kappa_{\text{in}} \cdot \theta/2}} \right) - \lceil d/k \rceil \cdot \left(\frac{1}{(s/k)^{\kappa-1}} + \frac{s/\beta}{(s/k)^{(\kappa'-2) \cdot \kappa \cdot \theta/2}} \right).$$

Furthermore, it uses the following resources, where $B = 2\kappa\kappa' \cdot \ln s$:

– Communication:

$$O\left(\sum_{i=1}^{\lceil d/k \rceil - 1} \frac{s_{i+1} + m_i}{\beta} \cdot \text{poly}(\lambda) \cdot (1 + B\lambda)^{2^k} \cdot \log(1 + B\theta')^{2^k}\right).$$

– Computation:

$$O\left(\sum_{i=1}^{\lceil d/k \rceil - 1} \frac{s_{i+1} + m_i}{\beta} \left[\lambda \cdot (1 + B\theta)^{\text{tpp}} \cdot (1 + B\theta')^{\text{tpp}} \cdot (\log |\mathbb{F}|)^2 \right. \right. \\ \left. \left. + \cdot \text{poly}((1 + B\lambda)^{\text{tpp}} \cdot \log(1 + B\theta')^{\text{tpp}} \cdot \lambda) \right] \right. \\ \left. + (\kappa_{in} \cdot n \log n \cdot \theta' \cdot (\theta \cdot \log |\mathbb{F}| + \log n) \cdot \lambda) \right. \\ \left. + (2\kappa \cdot s \log s \cdot \theta' \cdot (\theta \cdot \log |\mathbb{F}| + \log s) \cdot \lambda) \right).$$

6 Choice of Parameters

In this section, we tune the parameters of our protocol. We want to ensure the scheme is correct with all but negligible probability, that it is secure, that the communication is sublinear, and that the computation is polynomial. We make two sets of choices for the parameters: the first optimising for communication (*i.e.* maximising the sublinearity factor), and the other for computation (and incidentally for the strength of the security assumption). Since the latter entails minimising the sublinearity factor k , we can see any other choice as a tradeoff between the two, as a smooth continuum. Before we proceed, let us recall what the parameters are and how they are constrained.

6.1 Parameters.

$k = k(s)$ is the depth of the blocks in the *layer separation* phase; it is also the sublinearity factor of the final protocol. β is the number of outputs batched together for evaluation in a block, in the *output separation* phase. For the *input separation* phase, we need one B -good cover comprised of M_{in} size- θ' masks and $(\lceil d/k \rceil - 1)$ B -good covers comprised of M size- θ' masks. For convenience we prefer to consider $\kappa_{in} := M_{in}/(n \log n)$, $\kappa := M/(n \log n)$, and $\kappa' := B/(2\kappa \ln s)$. λ is the security parameter in the $\mathbb{F}\text{-xLPN}(\theta, \theta' - \theta, \lambda/\theta')$ assumption (which additionally introduces the parameter $\theta \leq \theta'$) for the PCG and in the $\mathbb{F}_2\text{-LPN}(\lambda, O(\lambda), 1/\lambda^\delta)$ – where δ is any small enough constant – assumption for the length-doubling PRG and the oblivious transfer.

6.2 Constraints.

1. *Correctness*: We have several choices on how to tune the parameters $\kappa, \kappa', \kappa_{in}$:
 - $p > 0$: If we choose the parameters so that p is non-zero, by a probabilistic argument, there must exist a choice of randomness which satisfies all the requirements. Then, if we so choose, we can preprocess circuits to find such a choice, or be satisfied with the non-constructive proof of existence of a sublinear protocol.
 - $p = \Omega(1)$: If we now choose them so that $1 \geq p \geq c$ for some positive constant c , then the parties can sample fresh randomness until a combination of good covers is obtained, which occurs in an expected constant number of tries.
 - $p = 1 - o(1)$: Finally if p is taken to be close to 1, the protocol succeeds on the first try of sampling randomness with almost certain probability (or even with all but negligible probability if $p = 1 - \text{negl}$).

We focus only on the last one, which is the most natural requirement. While it is the same

$$\left(\frac{1}{n^{\kappa_{in}-1}} + \frac{(s/k)/\beta}{n^{(\kappa'-2)\cdot\kappa_{in}\cdot\theta/2}} \right) + \lceil d/k \rceil \cdot \left(\frac{1}{(s/k)^{\kappa-1}} + \frac{s/\theta}{(s/k)^{(\kappa'-2)\cdot\kappa\cdot\theta/2}} \right) \\ = \frac{1}{s^{\omega(1)}} = 2^{-\omega(\log s)}.$$

Since $\theta = \omega(1)$, κ, κ' are unconstrained (other than $\kappa' > 2$ and $\kappa > 1$) and can be taken to be constant. Therefore the constraint can be simplified to κ_{in} being subject to $\kappa_{\text{in}} = \omega\left(\frac{\log s}{\log n}\right)$. Hence the following correctness constraints:

$$\begin{cases} M_{\text{in}} = \omega(n \cdot \log s) \\ M = \Omega(s \cdot \log s) \\ B = \Omega(\log s) \end{cases} \quad (1)$$

2. *Security*: Security depends on two security assumptions \mathbb{F} -xLPN($\theta, \theta' - \theta, \lambda/\theta'$) (for the PCG) and \mathbb{F}_2 -LPN($\lambda, O(\lambda), 1/\lambda^\delta$) – where δ is any small enough constant – (for the underlying PRG and OT). As we define the parameters θ, θ', λ , we will refine the exact flavour of assumption required.
3. *Communication*: The total communication of our protocol is the following:

$$O\left(\underbrace{n + \frac{s}{k} + m}_{\substack{\text{Cost of} \\ \text{reconstructing} \\ \text{masked values (} n \\ \text{inputs, } \leq s/k \\ \text{intermediary gates,} \\ \text{and } m \text{ outputs)}}} + \underbrace{\frac{s/k+m}{\beta}}_{\substack{\text{Number of} \\ \text{batches of} \\ \beta \text{ nodes} \\ \text{(i.e. number of} \\ \text{calls to PCG)}}} \times \underbrace{\text{poly}(\lambda)}_{\substack{\text{Overhead} \\ \text{from} \\ \text{the PRG} \\ \text{(including} \\ \text{the OTs)}}} \times \underbrace{(1 + B\lambda)^{2^k}}_{\substack{\text{Number of DPF} \\ \text{calls in MPFSS}}} \times \underbrace{\log\left((1 + B\theta')^{2^k}\right)}_{\substack{\text{Seed for one DPF}}} \Bigg).$$

We want this quantity to be $O(n + s/k + m)$, which translates to the following condition:

$$\beta = \Omega\left(\text{poly}(\lambda) \cdot (1 + B\lambda)^{2^k} \cdot \log\left((1 + B\theta')^{2^k}\right)\right). \quad (2)$$

4. *Computation*: The total computation is the following:

$$O\left(\frac{s/k+m}{\beta} \left[\lambda \cdot (1 + B\theta)^{\text{tpp}} \cdot (1 + B\theta')^{\text{tpp}} \cdot \log|\mathbb{F}| \right. \right. \\ \left. \left. + \cdot \text{poly}\left((1 + B\lambda)^{\text{tpp}} \cdot \log(1 + B\theta')^{\text{tpp}} \cdot \lambda\right) \right] \right. \\ \left. + M_{\text{in}} \cdot \theta' \cdot (\theta \cdot \log|\mathbb{F}| + \log n) \cdot \lambda + M \cdot \theta' \cdot (\theta \cdot \log|\mathbb{F}| + \log s) \cdot \lambda \right).$$

We want it to remain polynomial in s .

6.3 Choice of Parameters – Optimising for Communication.

Choosing k . Computation involves a term of the form B^{2^k} by eq. (2) with $B = \Omega(\log s)$ by eq. (1), which must remain polynomial in s .

$$(\log s)^{2^k} \leq s^{O(1)} \iff k \leq \log \log s - \log \log \log s + O(1).$$

Since k is the sublinearity factor, we also need it to be in $\omega(1)$, thus yielding the following constraint:

$$\omega(1) \leq k \leq \log \log s - \log \log \log s + O(1). \quad (3)$$

When we try and maximise k , it becomes more convenient to introduce the quantity $\epsilon := k/\log \log s$. Since we are only interested in asymptotic sublinearity, we can strengthen the condition $\epsilon \leq 1 - \frac{\log \log \log s - O(1)}{\log \log s}$ to $\epsilon < \frac{1}{4}$ (this particular choice of constant will become apparent as we choose the other parameters as a function of k and hence of ϵ) without ultimately weakening the result. While it makes most sense in the vicinity of $O(\log \log s)$, it is well defined for the full spectrum of values of k :

$$\omega\left(\frac{1}{\log \log s}\right) \leq \epsilon < \frac{1}{4} \iff \omega(1) \leq k < \frac{\log \log s}{4}.$$

Choosing B , M_{in} , M . It makes sense so minimise these parameters, i.e. make them match the lower bounds of eq. (1). Therefore, (if we choose to set M_{in} to the explicit $n \cdot (\log s)^2$ rather than an arbitrary small $\omega(n \cdot \log s)$ for simplicity), with $\kappa \leftarrow 2$ and $\kappa' \leftarrow 3$:

$$\begin{cases} M_{\text{in}} = n \cdot (\log s)^2 \\ M = 2 \cdot s \cdot \log s \\ B = 12 \cdot \log s \end{cases} \quad (4)$$

Choosing β . We choose to push β as high as it will go, i.e. we just need to ensure β^{2^k} remains polynomial in s (indeed in the general case, a batch of β nodes can have up to $2^k \cdot \beta$ ancestors in a depth- 2^k circuit, hence we need $\beta \leq B\theta'$). This is the optimal choice if we want to prioritise communication complexity and security, but as we shall show, it will not hinder optimising for computation complexity.

$$\beta \leftarrow s^{O(1/2^k)} = 2^{O(\log^{1-\epsilon} s)}.$$

Choosing λ . Since λ determines the security of the LPN scheme it makes sense to maximise it, especially as it has no bearing on the asymptotic computation complexity. The only upper bound on λ is $\beta = \Omega((B\lambda)^{2^k+2} \cdot \log s)$ i.e. $(12 \cdot \log s \cdot \lambda)^{2^k+2} \cdot \log s = O(\beta) = O(s^{1/2^k})$. This translates to the following for λ :

$$\begin{aligned} \lambda &\leq \left(\frac{s^{O(1/2^k)}}{12 \cdot \log s} \right)^{O(1/(2^k+2))} / \log s \\ &= 2^{O(\log^{1-2\epsilon} s - \log \log s \cdot \log^\epsilon s - \log \log s)} \\ &= 2^{O(\log^{1-2\epsilon} s)}, \text{ since } \epsilon < \frac{1}{4}. \end{aligned}$$

If we set $\lambda \leftarrow 2^{O(\log^{1-2\epsilon} s)}$ then $s = 2^{O(\log^{1-2\epsilon} \lambda)} = \lambda^{O(\log^{\epsilon'} \lambda)}$, where $\epsilon' := \frac{2\epsilon}{1-2\epsilon} \in]0, 1[$. In other terms, if $\lambda \leftarrow (s/\log s)^{O(1/2^k)}$ then $s = O(2^{2^k} \cdot (\lambda)^{O(2^k)} \cdot \log \lambda) = (\lambda)^{O(2^k)}$. Summarising:

$$\begin{cases} s = \lambda^{O(\log^{\epsilon'} \lambda)} \\ s = \lambda^{O(2^k)} \end{cases}$$

If we make sure $\theta' = \Theta(\theta)$, then this means that we can now choose the flavour of security assumption we need for \mathbb{F} -xLPN($\theta, \theta' - \theta, \lambda/\theta'$). If the adversary runs in time $\text{poly}(s)$ then we want their advantage to be $1/s^{\omega(1)} = 1/2^{\text{poly}(\log(\lambda))}$ which is the case if we assume the super-polynomial security of \mathbb{F} -xLPN(θ, θ', λ). Note however that if k (and therefore 2^{2^k}) is taken to be an arbitrarily small $\omega(1)$ function, then we only need to assume its quasi-polynomial security (which smoothly transitions to polynomial as k tends to $\Theta(1)$).

Choosing θ and θ' . We assumed in the previous paragraph that $\theta' = \Theta(\theta)$, and we choose the constant θ'/θ so that \mathbb{F}_2 -LPN($\lambda, O(\lambda), 1/\lambda^\delta$) – where δ is any small enough constant – is implied by \mathbb{F}_2 -xLPN($\theta, \theta' - \theta, \lambda/\theta'$). The computation (and to a lesser extent the communication) grows in $B\theta'$ rather than in θ' alone, so the exact value of θ' isn't all that important (and can be compensated by B). Since $B\theta'$ must be at least $\beta 2^k$ in the general case (since all the ancestors of a batch of β outputs must be covered by the concatenation of B size- θ' masks) and it makes sense for B to be a positive integer, we want $\theta' \geq \beta 2^k$. We therefore set $\theta \leftarrow \beta \cdot 2^k = \log^\epsilon s \cdot 2^{O(\log^{1-\epsilon} s)}$.

Summary. We now summarise the choice of parameters which optimises for communication in fig. 8.

| Parameter | Description | Choice for Optimal Comm. |
|----------------------|----------------------------|---|
| k | Subl. Factor / Block Depth | $\epsilon \cdot (\log \log s)$, $\epsilon \in]0, 1/4[$ |
| β | Output batch size | $2^{O(\log^{1-\epsilon} s)}$ |
| λ | Security Parameter | $2^{O(\log^{1-2\epsilon} s)}$ |
| θ | – | $\beta \cdot \log^\epsilon s$ |
| θ' | Mask Length in Good Cover | $O(\theta)$ |
| κ_{in} | Defines M_{in} | $(\log s)^2 / \log n$ |
| κ | Defines M | 2 |
| κ' | Defines B | 3 |
| B | "Goodness" of Cover | $12 \cdot \log s$ |
| Total Communication: | | $\Theta([n + s/\log \log s + m] \cdot \log \mathbb{F})$ |
| Total Computation: | | $s^3 \cdot \text{polylog}(s) \cdot (\log \mathbb{F})^2$ |

Fig. 8: Choice of parameters for Main Theorem 1. The constants hidden in the exponents of β and λ are correlated, but omitted here for simplicity.

6.4 Choice of Parameters – Optimising for Computation and Security Assumption.

If we fix the sublinearity factor $k = \omega(1)$ to be some very slow-growing function (e.g. the inverse Ackermann function), we can set:

$$\beta \leftarrow s^{2^{-2k}}; \lambda \leftarrow s^{2^{-3k}}; \theta \leftarrow \beta \cdot 2^k; \theta' \leftarrow \text{cst} \cdot \theta; \kappa \leftarrow 2; \kappa' \leftarrow 3; \kappa_{\text{in}} \leftarrow \frac{(\log s)^2}{\log n}.$$

With this choice, we have $s = \lambda^{\omega(1)}$ where s need only be larger than any polynomial in λ , so the assumed security of LPN need only be *quasi-polynomial*. Using these parameters for our sublinear protocol, we get a quasi-linear amount of computation $s^{1+o(1)}$ (where the exact nature of the $o(1)$ function depends on how we tune the parameters), with a $(\log |\mathbb{F}|)^2$ overhead.

6.5 Wrapping-Up: Main Results with Optimal Choices of Parameters

Combining Theorems 12 and 20 (the former providing a secure protocol in the $\mathcal{F}_{\text{corr}}$ -hybrid model, and the latter instantiating the $\mathcal{F}_{\text{corr}}$ functionality) with the parameter choice of Section 6.3 (and summarised in fig. 8) we get our main theorem, Main Theorem 1 below.

Main Theorem 1 (Sublinear Computation of Layered Circuits – Optimised for Communication).
Assuming the super-polynomial security of

- \mathbb{F} -LPN with super-polynomial dimension ℓ , $O(\ell)$ samples, and inverse super-polynomial rate,
- \mathbb{F}_2 -LPN with super-polynomial dimension $\ell' = s^{O(1)}$, $O(\ell')$ samples, and inverse polynomial rate (which is implied by the above if $\mathbb{F} = \mathbb{F}_2$),

there exists a probabilistic semi-honest two-party protocol which securely evaluates any layered arithmetic circuit over \mathbb{F} with success probability $1 - \text{negl}(s)$ and which uses $O([n + s/\log \log s + m] \cdot \log |\mathbb{F}|)$ bits of communication and $s^3 \cdot \text{polylog} s \cdot (\log |\mathbb{F}|)^2$ bits of computation (where s , n , and m are respectively the number of gates, inputs, and outputs of the circuit).

Instantiating the protocol with the parameters in Section 6.4 instead yields the following.

Main Theorem 2 (Sublinear Computation of Layered Circuits – Optimised for Computation).
Assuming the quasi-polynomial security of

- \mathbb{F} -LPN with quasi-polynomial dimension ℓ , $O(\ell)$ samples, and inverse quasi-polynomial rate,
- \mathbb{F}_2 -LPN with quasi-polynomial dimension ℓ' , $O(\ell')$ samples, and inverse polynomial rate (which is implied by the above if $\mathbb{F} = \mathbb{F}_2$),

there exists a probabilistic semi-honest two-party protocol which securely evaluates any layered arithmetic circuit over \mathbb{F} with success probability $1 - \text{negl}(s)$ and which uses $O([n + o(s) + m] \cdot \log |\mathbb{F}|)$ bits of communication and $s^{1+o(1)} \cdot (\log |\mathbb{F}|)^2$ bits of computation (where s , n , and m are respectively the number of gates, inputs, and outputs of the circuit).

References

- AHI⁺17. B. Applebaum, N. Haramaty, Y. Ishai, E. Kushilevitz, and V. Vaikuntanathan. Low-complexity cryptographic hash functions. pages 7:1–7:31, 2017.
- AIK07. B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography with constant input locality. pages 92–110, 2007.
- AIK09. B. Applebaum, Y. Ishai, and E. Kushilevitz. Cryptography with constant input locality. 22(4):429–469, October 2009.
- AJL⁺12. G. Asharov, A. Jain, A. López-Alt, E. Tromer, V. Vaikuntanathan, and D. Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. pages 483–501, 2012.
- BCG⁺17. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, and M. Orrù. Homomorphic secret sharing: Optimizations and applications. pages 2105–2122, 2017.
- BCG⁺19a. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, P. Rindal, and P. Scholl. Efficient two-round OT extension and silent non-interactive secure computation. pages 291–308, 2019.
- BCG⁺19b. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, and P. Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. pages 489–518, 2019.
- BCG⁺20a. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, and P. Scholl. Correlated pseudorandom functions from variable-density LPN. pages 1069–1080, 2020.
- BCG⁺20b. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, and P. Scholl. Efficient pseudorandom correlation generators from ring-LPN. pages 387–416, 2020.
- BCGI18. E. Boyle, G. Couteau, N. Gilboa, and Y. Ishai. Compressing vector OLE. pages 896–912, 2018.
- BFKL93. A. Blum, M. L. Furst, M. J. Kearns, and R. J. Lipton. Cryptographic primitives based on hard learning problems. In *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, pages 278–291, 1993.
- BFKL94. A. Blum, M. L. Furst, M. J. Kearns, and R. J. Lipton. Cryptographic primitives based on hard learning problems. pages 278–291, 1994.
- BFKR91. D. Beaver, J. Feigenbaum, J. Kilian, and P. Rogaway. Security with low communication overhead. pages 62–76, 1991.
- BGI15. E. Boyle, N. Gilboa, and Y. Ishai. Function secret sharing. pages 337–367, 2015.
- BGI16a. E. Boyle, N. Gilboa, and Y. Ishai. Breaking the circuit size barrier for secure computation under DDH. pages 509–539, 2016.
- BGI16b. E. Boyle, N. Gilboa, and Y. Ishai. Function secret sharing: Improvements and extensions. pages 1292–1303, 2016.
- BI05. O. Barkol and Y. Ishai. Secure computation of constant-depth circuits with applications to database search problems. pages 395–411, 2005.
- BKS19. E. Boyle, L. Kohl, and P. Scholl. Homomorphic secret sharing from lattices without FHE. pages 3–33, 2019.
- BKW00. A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. pages 435–440, 2000.
- BKW03. A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM (JACM)*, 50(4):506–519, 2003.
- BLVW19. Z. Brakerski, V. Lyubashevsky, V. Vaikuntanathan, and D. Wichs. Worst-case hardness for LPN and cryptographic hashing via code smoothing. pages 619–635, 2019.
- CG97. B. Chor and N. Gilboa. Computationally private information retrieval (extended abstract). pages 304–313, 1997.
- CGKS95. B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. pages 41–50, 1995.
- Cou19. G. Couteau. A note on the communication complexity of multiparty computation in the correlated randomness model. pages 473–503, 2019.
- DFH12. I. Damgård, S. Faust, and C. Hazay. Secure two-party computation with low communication. pages 54–74, 2012.
- DGH⁺20. N. Döttling, S. Garg, M. Hajiabadi, D. Masny, and D. Wichs. Two-round oblivious transfer from cdh or lpn. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 768–797. Springer, 2020.
- FGJS17. N. Fazio, R. Gennaro, T. Jafarikhah, and W. E. Skeith III. Homomorphic secret sharing from paillier encryption. pages 381–399, 2017.
- Gen09. C. Gentry. Fully homomorphic encryption using ideal lattices. pages 169–178, 2009.
- GI14. N. Gilboa and Y. Ishai. Distributed point functions and their applications. pages 640–658, 2014.
- GMW87a. O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. pages 218–229, 1987.

- GMW87b. O. Goldreich, S. Micali, and A. Wigderson. How to prove all NP-statements in zero-knowledge, and a methodology of cryptographic protocol design. pages 171–185, 1987.
- IKM⁺13. Y. Ishai, E. Kushilevitz, S. Meldgaard, C. Orlandi, and A. Paskin-Cherniavsky. On the power of correlated randomness in secure computation. pages 600–620, 2013.
- JKPT12. A. Jain, S. Krenn, K. Pietrzak, and A. Tentes. Commitments and efficient zero-knowledge proofs from learning parity with noise. pages 663–680, 2012.
- KO97. E. Kushilevitz and R. Ostrovsky. Replication is NOT needed: SINGLE database, computationally-private information retrieval. pages 364–373, 1997.
- Lyu05. V. Lyubashevsky. The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In *Approximation, randomization and combinatorial optimization. Algorithms and techniques*, pages 378–389. Springer, 2005.
- NN01. M. Naor and K. Nissim. Communication preserving protocols for secure function evaluation. pages 590–599, 2001.
- OSY21. C. Orlandi, P. Scholl, and S. Yakoubov. The rise of paillier: Homomorphic secret sharing and public-key silent ot. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 678–708. Springer, 2021.
- Pra62. E. Prange. The use of information sets in decoding cyclic codes. *IRE Transactions on Information Theory*, 8(5):5–9, 1962.
- RS21. L. Roy and J. Singh. Large message homomorphic secret sharing from dcr and applications. 2021.
- SGRR19. P. Schoppmann, A. Gascón, L. Reichert, and M. Raykova. Distributed vector-OLE: Improved constructions and implementation. pages 1055–1072, 2019.
- Wak68. A. Waksman. A permutation network. *Journal of the ACM (JACM)*, 15(1):159–163, 1968.
- Yao86. A. C.-C. Yao. How to generate and exchange secrets (extended abstract). pages 162–167, 1986.
- YZW⁺19. Y. Yu, J. Zhang, J. Weng, C. Guo, and X. Li. Collision resistant hashing from sub-exponential learning parity with noise. pages 3–24, 2019.