



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

Geoffroy Couteau^{1(✉)} and Pierre Meyer^{2(✉)}

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

² École Normale Supérieure de Lyon and IDC Herzliya, Herzliya, Israel
pierre.meyer@ens-lyon.fr

Abstract. In this work we introduce a new (circuit-dependent) *homomorphic secret sharing* (HSS) scheme for all \log / \log -local circuits, with communication proportional only to the width of the circuit, and polynomial computation, 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, from short seeds, pseudorandom instances of an arbitrary \log / \log -local additive correlation.

Our main application, and the main 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) -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 (additive) 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 a circularly secure variant of DCR [FGJS17], 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. 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 might be 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 .¹ [FGJS17] later followed this blueprint and switched out the DDH assumption for the circular security of the Paillier encryption scheme. 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 A depth- d circuit is layered if it can be divided into d layers such that any wire connects adjacent layers.

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 $T = T(\lambda)$,

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

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, some 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 all log log-depth circuits. More precisely,

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 $\epsilon < 1/4$). 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 $O(m \cdot \text{poly}(s) \cdot (\log |\mathbb{F}|)^2)$.*

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 [BGH16a], 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 $s^3 \cdot \text{polylogs} \cdot (\log |\mathbb{F}|)^2$ 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 $2^{O(\lambda)}$.² 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).

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;

² 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.

– 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.*

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

In the above, $\vec{a}^{\otimes b}$ denotes a vector \vec{a} tensored with itself b times. In order to build shares of $(\vec{r}, \vec{r}^{\otimes 2})$ for some (pseudo)random $\vec{r} \in \mathbb{F}^n$ (the bilinear correlation), the PCG of [BCG+19b] uses a multi-point function secret sharing scheme (MPFSS) (defined in Sect. 3.1) to give the parties small seeds which can be expanded locally to shares of $(\vec{e}, \vec{e}^{\otimes 2})$ for some random sparse vector $\vec{e} \in \mathbb{F}^n$. Thence, if H is some suitable public matrix the parties can get shares of $\vec{r} := H \cdot \vec{e}$, which is pseudorandom under LPN, and of $\vec{r}^{\otimes 2} = H^{\otimes 2} \cdot \vec{e}^{\otimes 2}$ by locally multiplying their shares of \vec{e} and $\vec{e}^{\otimes 2}$ by H and $H^{\otimes 2}$ respectively. The main issue in using this approach directly is that performing the expanding $\vec{r}^{\otimes \text{tpp}} = H[S_i]^{\otimes \text{tpp}} \cdot \vec{e}^{\otimes \text{tpp}}$ (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 \vec{r} be a (pseudo)random mask, we construct \vec{r} as a sum of $n \cdot \log n$ vectors \vec{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 \vec{r} (and in particular S_1, \dots, S_{n_s}) will be expressible as a sum of ‘not too many’ of the \vec{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 \vec{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 \vec{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 [BCGI18, 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].

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 (\vec{x}, \vec{y}). By default, vectors are assumed to be column vectors. If \vec{x} and \vec{y} are two (column) vectors, we use $\vec{x}||\vec{y}$ to denote the (column) vector obtained by their concatenation. We write $\vec{x} \otimes \vec{y}$ to denote the tensor product between \vec{x} and \vec{y} , i.e., the vector of length $n_x n_y$ with coordinates $x_i y_j$ (where n_x is the length of \vec{x} and n_y is the length of \vec{y}). We write $\vec{x}^{\otimes 2}$ for $\vec{x} \otimes \vec{x}$, and more generally, $\vec{x}^{\otimes n}$ for the n -th tensor power of \vec{x} , $\vec{x} \otimes \vec{x} \otimes \dots \otimes \vec{x}$. Given a vector \vec{x} of length $|\vec{x}| = n$, the notation $\text{HW}(x)$ denotes the Hamming weight \vec{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 a 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 $\vec{x} \in \mathbb{F}^n$ is done by assigning the coordinates of \vec{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 \vec{x} , denoted $C(\vec{x})$, 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 $\text{tpp} = K$, 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:

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

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 $(\vec{r}, \vec{r}^{\otimes 2}, \dots, \vec{r}^{\otimes d})$ for some constant d , where \vec{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 $(\vec{r}, Q(\vec{X} - \vec{r}))$ 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 \vec{y}_i into w/β subvectors, each containing β consecutive coordinates of \vec{y}_i . Then, the core observation is that a simple variant of the PCG of [BCG+19b] allows in fact to generate shares of $(\vec{r}, \vec{r}^{\otimes 2}, \dots, \vec{r}^{\otimes 2^k})$ for some pseudorandom r , where $\vec{r}^{\otimes j}$ denotes the tensor product of \vec{r} with itself j times (which we call from now on the *j-th tensor power of \vec{r}*): this correlation is enough to generate shares of all degree- 2^k polynomial in \vec{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 \vec{y}_{i+1} , which we denote \vec{v} . Observe that the entire subvector \vec{v} can depend on at most $\beta \cdot 2^k$ coordinates of \vec{y}_i , since each coordinate of \vec{v} depends on at most 2^k coordinates of \vec{y}_i . Therefore, we can now see the computation of \vec{v} from \vec{y}_i as evaluating β multivariate polynomials $(Q_1 \dots, Q_\beta)$, where all multivariate polynomials take as input the same size- $(\beta 2^k)$ subset of coordinates of \vec{y}_i . To securely compute shares of \vec{v} from shares of \vec{y}_i , the parties can use the following

type of correlated randomness: they will have shares of $(\vec{r}, \vec{r}^{\otimes 2}, \dots, \vec{r}^{\otimes 2^k})$, where \vec{r} is a random mask of length $\beta \cdot 2^k$. Consider the following polynomials:

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

Each coefficient of each \hat{Q} can be computed as a degree- 2^k multivariate polynomial in the coordinates of \vec{r} – or, equivalently, as a linear combination of the coordinates of $(\vec{r}, \vec{r}^{\otimes 2}, \dots, \vec{r}^{\otimes 2^k})$. Hence, given additive shares of $(\vec{r}, \vec{r}^{\otimes 2}, \dots, \vec{r}^{\otimes 2^k})$, 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 $(\vec{r}, \vec{r}^{\otimes 2}, \dots, \vec{r}^{\otimes 2^k})$ 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 $(\vec{r}, \vec{r}^{\otimes 2}, \dots, \vec{r}^{\otimes 2^k})$ 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 $\vec{y} + \vec{r}$, where \vec{y} contains values of some layer, and \vec{r} is some mask. Recall that with the method which we just outlined, the parties must generate such a length- $(\beta 2^k)$

mask \vec{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 $\vec{y} + \vec{r}$ 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 \vec{r}_i to mask a full layer \vec{y}_i , and define the mask for the subset $S_{i,j}$ of ancestors of a target value $y_{i+1,j}$ to be $\vec{r}_i[S_{i,j}]$. This implies that the parties must now broadcast a single masked vector $\vec{y}_i + \vec{r}_i$ 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 $(\vec{r}, \dots, \vec{r}^{\otimes 2^k})$ where \vec{r} is of length w , the PCG will first generate a *very sparse* vector \vec{r}^t , with some number t of nonzero coordinates. Then, each $(\vec{r}^t)^{\otimes n}$ 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 $(\vec{r}^t)^{\otimes n}$ to length- $t^n \cdot \log w$ seeds. Then, the final pseudorandom correlation is obtained by letting the parties locally compress \vec{r}^t by multiplying it with a large public matrix H , giving a vector $\vec{r} = H \cdot \vec{r}^t$. Similarly, $\vec{r}^{\otimes n}$ can be reconstructed by computing $H^{\otimes n} \cdot (\vec{r}^t)^{\otimes n} = (H \cdot \vec{r}^t)^{\otimes n} = \vec{r}^{\otimes n}$, 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 \vec{r}^t is indistinguishable from random, under the dual LPN assumption (which is equivalent to the standard LPN assumption). Concretely, one can think of \vec{r}^t as being of length $2w$, and of H as being a matrix from $\mathbb{F}^{w \times 2w}$ which compresses \vec{r}^t 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 \vec{r} , the computation for expanding the seed to these pseudorandom tensor powers will grow super-polynomially with the length of *entire* vector w . Indeed, consider generating the 2^k -th tensor power of a subvector $\vec{r}[S]$ of \vec{r} , for some size- $\beta \cdot 2^k$ subset S of $[w]$. Then with the PCG of [BCG+19b], this requires computing $(H[S])^{\otimes 2^k} \cdot (\vec{r}[S])^{\otimes 2^k}$, where the share of $(\vec{r}[S])^{\otimes 2^k}$ 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 \vec{r} as the sum of a certain number m of shorter masks $\vec{r}^1, \vec{r}^2, \dots$ which each only cover θ values (note that they may – and will – overlap). This way the 2^k -th tensor power of a subvector \vec{v} can be obtained from appropriate linear combinations of coordinates of the 2^k -th tensor power of the concatenation of *only* the \vec{r}^j which overlap with \vec{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 $\vec{r}[S_i]$, for some random $\vec{r} \in \mathbb{F}^w$. 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 $\vec{r}^j \in \mathbb{F}^K$: together they define a sparse subvector of \mathbb{F}^w . If we let \vec{r} be the sum of these sparse vectors, it is clear that for any $i \in [w/\beta]$, each element of $(1_{\mathbb{F}} \parallel \vec{r}[S_i])^{\otimes 2^k}$ 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 \vec{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 Sect. 2.1) allows the parties to evaluate the circuit in a block-by-block 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 . For the formal definition of FSS, we refer the reader to the full version of this paper. 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 3 (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 4 (Multi-Point Function [BCGI18]). *An (n,t) -multi-point function over an abelian group $(\mathbb{G}, +)$ is a function $f_{S,\vec{y}} : [n] \rightarrow \mathbb{G}$, where $S = (s_1, \dots, s_t)$ is an ordered subset of $[n]$ of size t and $\vec{y} = (y_1, \dots, y_t) \in \mathbb{G}^t$, defined by $f_{S,\vec{y}}(s_i) = y_i$ for any $i \in [t]$, and $f_{S,\vec{y}}(x) = 0$ for any $x \in [n] \setminus S$.*

We assume that the description of S includes the input domain $[n]$ so that $f_{S,\vec{y}}$ 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 $f_{S,\vec{y}}$ 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 [BFLK93] (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 5 (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*

$$\begin{aligned} & \{(A, \vec{b}) \mid A \xleftarrow{s} \mathbb{F}^{q \times k}, \vec{e} \xleftarrow{s} \text{Ber}_r(\mathbb{F})^q, \vec{s} \xleftarrow{s} \mathbb{F}^k, \vec{b} \leftarrow A \cdot \vec{s} + \vec{e}\} \\ & \approx \{(A, \vec{b}) \mid A \xleftarrow{s} \mathbb{F}^{q \times k}, \vec{b} \xleftarrow{s} \mathbb{F}^q\} \end{aligned}$$

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 [BFLK93, 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 \vec{e} is not sampled from $\text{Ber}_r(\mathbb{F})^q$, but it is sampled uniformly from the set $\text{HW}_{rq}(\mathbb{F}^q)$ 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 6 (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

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

Solving the dual LPN assumption is easily seen to be at least as hard as solving LPN: given a sample (A, \vec{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 $(H, H \cdot \vec{b})$ to the dual LPN solver. Note that the parity check matrix of a random matrix is distributed as a random matrix. Furthermore, when $\vec{b} = A \cdot \vec{s} + \vec{e}$, we have $H \cdot \vec{b} = H \cdot (A \cdot \vec{s} + \vec{e}) = H \cdot \vec{e}$. 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) (k_0, k_1) such that each key k_σ can be expanded to a long string $R_\sigma = \text{Expand}(\sigma, k_\sigma)$, with the following guarantees: given the key $k_{1-\sigma}$, the string R_σ is indistinguishable from a random string sampled conditioned on satisfying the target correlation with the string $R_{1-\sigma} = \text{Expand}(1 - \sigma, k_{1-\sigma})$. The formal definition of PCGs is given in the full version of this paper

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

Lemma 7 (Block-Decomposition of Layered Circuits). *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 $\vec{x} \in \mathbb{F}^n$, $A(\vec{x}) = C(\vec{x})$ (that is, A computes C);

function $A(\vec{x})$


```

 $\vec{x}_0 \leftarrow \vec{x}$ 
for  $i = 0$  to  $t - 1$  do  $(\vec{x}_{i+1}, \vec{y}_i) \leftarrow f_i(\vec{x}_i)$ 
 $\vec{y} \leftarrow \vec{y}_0 || \dots || \vec{y}_{t-1}$ 
return  $\vec{y}$ 

```

- For any $i \in \llbracket 0, t - 1 \rrbracket$, $j \leq s_{i+1} + m_i$, the j -th output³ 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 $\sum_{i=0}^{t-1} m_i = m$.

4.2 Securely Computing C in the Correlated Randomness Model

We represent in Fig. 1 the ideal functionality for securely evaluating the layered arithmetic circuit C .

Ideal Functionality \mathcal{F}_C

- **Parameters.** The functionality is parametrised with an arithmetic circuit C with n inputs over a finite field \mathbb{F} .
- **Parties.** An adversary \mathcal{A} and N parties P_1, \dots, P_N . Each party P_ℓ has $p_\ell \in [0, n]$ inputs over \mathbb{F} , with $\sum_{\ell \leq N} p_\ell = n$.

The functionality aborts if it receives any incorrectly formatted message.

1. On input a message (input, \vec{x}_ℓ) from each party P_ℓ where $\vec{x}_\ell \in \mathbb{F}^{p_\ell}$, set

$$\vec{x} \leftarrow \vec{x}_1 || \dots || \vec{x}_N \in \mathbb{F}^n.$$
2. Compute $\vec{y} \leftarrow C(\vec{x})$. Output \vec{y} to all parties, and terminate.

Fig. 1. Ideal functionality \mathcal{F}_C for securely evaluating an arithmetic circuit C among N parties.

We represent on Fig. 2 an ideal functionality for distributing (function-dependent) correlated randomness between the parties.

Theorem 8. *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 securely generate with low communication overhead. The protocol and its security analysis are given in the full version.

³ i.e. the j^{th} coordinate of the image by f_i , seen as $f_i : \mathbb{F}^n \times \mathbb{F}^{s_i} \rightarrow \mathbb{F}^{s_{i+1} + m_i}$.

Ideal Functionality $\mathcal{F}_{\text{corr}}$

- **Parameters.** For every $i = 0, \dots, \lceil d/k \rceil - 1$, functionality is parameterised with subsets $(U_{i,j}^{\text{in}}, U_{i,j})_{1 \leq j \leq \lceil s_{i+1}/\beta \rceil}$ and $(V_{i,j}^{\text{in}}, V_{i,j})_{1 \leq j \leq \lceil m_i/\beta \rceil}$.
- **Parties.** An adversary \mathcal{A} and N parties P_1, \dots, P_N .

The functionality aborts if it receives any incorrectly formatted message.

1. On input a message $(\text{corrupt}, D)$ with $D \subsetneq [N]$ from \mathcal{A} , set $H \leftarrow [N] \setminus D$ and store (H, D) .
2. On input a message input with from each party P_ℓ , send ready to \mathcal{A} .
3. *Setup input masks:* On input a message $(\text{setinputshare}, (\vec{r}_{\text{in},\ell})_{\ell \in D})$ from \mathcal{A} with $\forall \ell \in D, \vec{r}_{\text{in},\ell} \in \mathbb{F}^n$, sample $(\vec{r}_{\text{in},\ell})_{\ell \in H} \xleftarrow{\$} (\mathbb{F}^n)^{|H|}$, and set $\vec{r}_{\text{in}} \leftarrow \sum_{\ell \in [N]} \vec{r}_{\text{in},\ell}$.
4. For $i = 1$ to $\lceil d/k \rceil - 1$:
 - (a) *Setup masks for the computation gates of the first layer of the i^{th} chunk:* On input a message $(\text{setblockshare}, i, (\vec{r}_{i,\ell})_{\ell \in D})$ from \mathcal{A} with $\forall \ell \in D, \vec{r}_{i,\ell} \in \mathbb{F}^{s_i}$, sample $(\vec{r}_{i,\ell})_{\ell \in H} \xleftarrow{\$} (\mathbb{F}^{s_i})^{|H|}$, and set $\vec{r}_i \leftarrow \sum_{\ell \in [N]} \vec{r}_{i,\ell}$.
 - (b) *Setup evaluation of the computation gates on the final layer of the i^{th} chunk:*
 - For $j = 1$ to $\lceil s^{i+1}/\beta \rceil$, set:

$$\vec{\pi}^{(i,j)} \leftarrow \left(1 \parallel \vec{r}_{\text{in}}[U_{i,j}^{\text{in}}] \parallel \vec{r}_i[U_{i,j}] \right)^{\otimes 2^k}.$$
 - Wait for a message $(\text{setshare}, (i, j), (\vec{\pi}_\ell^{(i,j)})_{\ell \in D})$ from \mathcal{A} with $\vec{\pi}_\ell^{(i,j)} \in \mathbb{F}^\delta$;
 - Compute uniformly random shares $(\vec{\pi}_\ell^{(i,j)})_{\ell \in [H]}$ of $\vec{\pi}^{(i,j)} - \sum_{\ell \in D} \vec{\pi}_\ell^{(i,j)}$.
 - (c) *Setup evaluation of the output gates in the i^{th} chunk:*
 - For $j = 1$ to $\lceil m_i/\beta \rceil$, set:

$$\vec{\pi}^{(i,j)} \leftarrow \left(1 \parallel \vec{r}_{\text{in}}[V_{i,j}^{\text{in}}] \parallel \vec{r}_i[V_{i,j}] \right)^{\otimes 2^k}.$$
 - Wait for a message $(\text{setoutputshare}, (i, j), (\vec{\pi}_\ell^{(i,j)})_{\ell \in D})$ from \mathcal{A} with $\vec{\pi}_\ell^{(i,j)} \in \mathbb{F}^\delta$;
 - Compute uniformly random shares $(\vec{\pi}_\ell^{(i,j)})_{\ell \in [H]}$ of $\vec{\pi}^{(i,j)} - \sum_{\ell \in D} \vec{\pi}_\ell^{(i,j)}$.
5. Output $(\vec{r}_{\text{in},\ell}, (\vec{r}_{i,\ell}, (\vec{\pi}_\ell^{(i,j)})_{1 \leq j \leq \lceil s_{i+1}/\beta \rceil}, (\vec{\pi}_{\text{out},\ell}^{(i,j)})_{1 \leq j \leq \lceil m_i/\beta \rceil})_{0 \leq i < \lceil d/k \rceil})$ to each party P_ℓ .

Fig. 2. Ideal corruptible functionality $\mathcal{F}_{\text{corr}}$ to deal out correlated randomness to the parties.

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 pseudorandom 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:

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

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, since it is an additive correlation: given any fixed share share_0 , a matching share can be reverse-sampled by sampling \vec{r} and setting $\text{share}_1 \leftarrow (\vec{r}, ((1_{\mathbb{F}} \parallel \vec{r}[S_i])^{\otimes \text{tpp}})_{1 \leq i \leq n_s}) - \text{share}_0$. 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 \vec{v} of size n , a family $(S_i)_{i \in [t]} \in \mathcal{P}([n])^t$ of t (*short*) subsets of coordinates of \vec{v} , and a (*small*) bound $B > 0$, the problem is to find a family $(\vec{v}_j)_{j \in [M]}$ of some number m of size- K subvectors of \vec{v} such that:

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

We call such a family a *B-Good Cover* of $(\vec{v}, (S_i)_{i \in [t]})$. 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 9 (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 $A = (\vec{\alpha}^j)_{j \in [M]} \in \binom{[n]}{K}^M$ is a *B-Good Cover* of $(S_i)_{i \in [t]}$ if:*

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

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 10 (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 $A = (\vec{\alpha}^j)_{j \in [M]} \in \binom{[n]}{K}^M$ 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 $A = (\vec{\alpha}^j)_{j \in [M]}$ 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 is given in the full version.

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 $\vec{r}^{\otimes \text{tpp}}$, for a parameter tpp and a pseudorandom \vec{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 Sect. 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 $H_{n',n} \xleftarrow{\$} \mathbb{F}^{(n'-n) \times n'}$ (Fig. 3).

PCG for Degree-tpp Polynomial Correlations

PCG.Gen: On input 1^λ :

1. Pick a random λ -sparse vector $\vec{e} \xleftarrow{\$} \text{HW}_\lambda(\mathbb{F}^{n'})$. Note that $\vec{e}^{\otimes \text{tpp}} \in \text{HW}_{\lambda \text{tpp}}(\mathbb{F}^{(n')^{\text{tpp}}})$. 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 $\vec{e}^{\otimes \text{tpp}}$.
2. Compute $(K_0^{\text{fss}}, K_1^{\text{fss}}) \xleftarrow{\$} \text{MPFSS.Gen}(1^\lambda, f)$. Output $k_0 \leftarrow (n, K_0^{\text{fss}})$ and $k_1 \leftarrow (n, K_1^{\text{fss}})$.

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

Fig. 3. PCG for low-degree polynomials from [BCG+19b].

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

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

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 $(\mathbf{k}_\sigma, \vec{r}_{1-\sigma})$ as in the first distribution above. Then, generate $(\mathbf{k}_\sigma, \vec{r}_{1-\sigma})$ as before, and generate an alternative key \mathbf{k}'_σ solely from the parameters $(1^\lambda, \mathbb{F}, n, n', t, \text{tpp})$, using the simulator of the MPFSS. Output $(\mathbf{k}'_\sigma, \vec{r}_{1-\sigma})$; under the security of the MPFSS, this distribution is indistinguishable from the previous one. Note that \mathbf{k}'_σ does not depend anymore on the noise vector \vec{e} . In the next hybrid, generate $\vec{r} \stackrel{s}{\leftarrow} H_{n',n} \cdot \vec{e}$ and set $\vec{r}_{1-\sigma} \leftarrow \vec{r}^{\otimes \text{tpp}} - \text{Expand}(\sigma, \mathbf{k}_\sigma)$; this game is perfectly indistinguishable from the previous one. Finally, replace $\vec{r} \stackrel{s}{\leftarrow} H_{n',n} \cdot \vec{e}$ by $\vec{r} \stackrel{s}{\leftarrow} \mathbb{F}^n$; 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 generality⁴, that $\bigcup_{i=1}^{n_s} S_i = [w]$. Our goal is for the parties to obtain shares of some pseudorandom vector $\vec{r} \in \mathbb{F}^w$ as well as shares of $(1 \parallel \vec{r}[S_i])^{\otimes \text{tpp}} \in \mathbb{F}^{w \cdot \text{tpp}}$ 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 $(\alpha_j, \vec{r}_j)_{j \in [m]} \in ((\frac{[w]}{\theta}) \times \mathbb{F}^\theta)^m$ where each \vec{r}_j is pseudorandom. We generate each of the m pseudorandom masks \vec{r}_j using a different instance of xLPN, *i.e.* $\vec{r}_j \leftarrow H_j \cdot \vec{e}_j$, where $\vec{e}_j \in \mathbb{F}^{\theta'}$ is λ -sparse and $H_j \stackrel{s}{\leftarrow} \mathbb{F}^{\theta \times \theta'}$ 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 $\vec{r} := \sum_{j=1}^m f_{\alpha_j, \vec{r}_j} \in \mathbb{F}^w$, where $f_{\alpha_j, \vec{r}_j} \in \mathbb{F}^w$ is the sparse vector defined by $(f_{\alpha_j, \vec{r}_j})_{\alpha_j} = \vec{r}_j$ (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 \vec{r}_j is pseudorandom, \vec{r} is also pseudorandom.

Note that for any given $i \in [n_s]$, $(1_{\mathbb{F}} \parallel \vec{r}[S_i])$ is a subvector of the vector \tilde{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 $\vec{e}'_i = (1_{\mathbb{F}} \parallel e_{j_1} \parallel \dots \parallel e_{j_{|I_i|}})$. Therefore for any tensor power tpp (*i.e.* the degree of the polynomial correlation), $\tilde{r}_i^{\otimes \text{tpp}} = (H'_i \cdot \vec{e}'_i)^{\otimes \text{tpp}} = (H'_i)^{\otimes \text{tpp}} \cdot (\vec{e}'_i)^{\otimes \text{tpp}}$. If the parties use an MPFSS scheme to generate small seeds which expand to $(\vec{e}'_i)^{\otimes \text{tpp}}$, they can then locally obtain shares of $\tilde{r}_i^{\otimes \text{tpp}}$ (since $(H'_i)^{\otimes \text{tpp}}$ is public), and therefore of $(1_{\mathbb{F}} \parallel \vec{r}[S_i])^{\otimes \text{tpp}}$. From all these shares of all the $(1_{\mathbb{F}} \parallel \vec{r}[S_i])^{\otimes \text{tpp}}, i \in [n_s]$ the parties can locally extract shares of all the $\vec{r}[S_i]$ and thence shares of \vec{r} (since $\bigcup_{i=1}^{n_s} S_i = [w]$). The protocol is given in Fig. 4.

⁴ If $\bigcup_{i=1}^{n_s} S_i \neq \emptyset$, and with the notations of the rest of the section, the vector \vec{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 $\vec{r}' = H' \cdot \vec{e}'$, pseudorandom under xLPN, to cover $[w] \setminus \bigcup_{i=1}^{n_s} S_i$. Since the parties do not need shares of $(\vec{r}')^{\otimes \text{tpp}}$, the communication complexity of generating the λ -sparse \vec{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 $\vec{e}_j \xleftarrow{\$} \text{HW}_\lambda(\mathbb{F}^{\theta'})$, $j \in [m]$ and define:

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

$$\tilde{\vec{r}}_i \leftarrow (\mathbb{1}_{\mathbb{F}} \parallel H_{j_1} \cdot \vec{e}_{j_1}^\top \parallel \dots \parallel H_{j_{m_i}} \cdot \vec{e}_{j_{m_i}}^\top)^\top.$$
 - (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 $f_i(x) = (\mathbb{1}_{\mathbb{F}} \parallel \vec{e}_{j_1} \parallel \dots \parallel \vec{e}_{j_{m_i}})^{\otimes \text{tpp}}[x]$. 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 $H'_i \leftarrow \text{Diag}(\mathbb{1}_{\mathbb{F}}, H_{j_1}, \dots, H_{j_{m_i}})$, compute

$$\vec{v}_{i,\sigma} \leftarrow \text{MPFSS.FullEval}(\sigma, K_{i,\sigma}^{\text{fss}}) \in \mathbb{F}^{(1+m_i \cdot \lambda)^{\text{tpp}}}$$
 and set $\vec{y}_\sigma \leftarrow ((H'_i)^{\otimes \text{tpp}} \cdot \vec{v}_\sigma)_{1 \leq i \leq n_s}$.
2. Extract from \vec{y}_σ the appropriate linear combinations of its elements corresponding to a share of $(\vec{r}, ((\mathbb{1}_{\mathbb{F}} \parallel \vec{r}[S_i])^{\otimes \text{tpp}})_{i \in [n_s]})$. // 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. 4. Pseudorandom correlation generator PCG_{stp} for generating pseudorandom instances of the subsets tensor powers correlation over a field \mathbb{F} .

Theorem 11. *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} - $\text{xLPN}(\theta, \theta', \lambda)$ 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 $(\vec{r}, ((\mathbb{1}_{\mathbb{F}} \parallel \vec{r}[S_i])^{\otimes \text{tpp}})_{1 \leq i \leq n_s})$ where $\vec{r} \in \mathbb{F}^w$.*

- 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 $\text{PCG}_{\text{stp.Gen}}$ 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 $\text{PCG}_{\text{stp.Expand}}$ 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} .

The proof of the above theorem is omitted in this version of the paper.

5.4 Instantiating the MPFSS

Theorem 11 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 [BG116b, BCG118] 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 recall in the full version of the paper.

5.5 Securely Distributing MPFSS.Gen an Π_{stp}

The seeds of the MPFSS scheme of [BCG118] 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 12. *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 $t' \cdot \nu \cdot \lambda$.*

We prove the above corollary by exhibiting Π_{MPFSS} in the full version. As a direct corollary of Corollary 12, since the seeds of PCG_{stp} contain exactly n_s independent MPFSS seeds, we have:

Corollary 13. *There exists a semi-honest secure two-party protocol Π_{stp} which distributes the seeds of the pseudorandom correlation generator PCG_{stp} represented on Fig. 4, 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 $m \cdot (\theta \cdot \theta' \cdot \log |\mathbb{F}| + \log \binom{w}{\theta'})$ (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, omitted from this version. Wrapping up, using Π_{stp} with an appropriate good cover suffices to construct a protocol Π_{corr} for securely implementing the functionality $\mathcal{F}_{\text{corr}}$. The detailed choice of parameters is deferred to the full version. Below, we describe a specific choice of parameters for the full construction which suffices to arrive at the claimed result.

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, and the other for computation (and incidentally for the strength of the security assumption). The full discussion is deferred to the full version.

Combining Theorem 8—which provides a secure protocol in the $\mathcal{F}_{\text{corr}}$ -hybrid model—and the instantiation of the $\mathcal{F}_{\text{corr}}$ as provided in the full version, with an

appropriate choice of parameters, also made explicit in the full version, 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 $s^3 \cdot \text{polylogs} \cdot (\log |\mathbb{F}|)^2$ bits of computation (where s , n , and m are respectively the number of gates, inputs, and outputs of the circuit).

Instantiating the protocol with an alternative choice of parameters, also detailed in the full version, 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. Applebaum, B., Haramaty-Krasne, N., Ishai, Y., Kushilevitz, E., Vaikuntanathan, V.: Low-complexity cryptographic hash functions, pp. 7:1–7:31 (2017)
- AIK07. Applebaum, B., Ishai, Y., Kushilevitz, E.: Cryptography with constant input locality. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 92–110. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_6
- AIK09. Applebaum, B., Ishai, Y., Kushilevitz, E.: J. Cryptol. Cryptography with constant input locality. **22**(4), 429–469 (2009)
- AJL+12. Asharov, G., Jain, A., López-Alt, A., Tromer, E., Vaikuntanathan, V., Wichs, D.: Multiparty computation with low communication, computation and interaction via threshold FHE. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 483–501. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_29

- BCG+17. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Orrù, M.: Homomorphic secret sharing: Optimizations and applications, pp. 2105–2122 (2017)
- BCG+19a. Boyle, E., et al.: Efficient two-round OT extension and silent non-interactive secure computation, pp. 291–308 (2019)
- BCG+19b. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Scholl, P.: Efficient pseudorandom correlation generators: silent OT extension and more. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11694, pp. 489–518. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_16
- BCG+20a. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Scholl, P.: Correlated pseudorandom functions from variable-density LPN, pp. 1069–1080 (2020)
- BCG+20b. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Scholl, P.: Efficient pseudorandom correlation generators from ring-LPN. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12171, pp. 387–416. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56880-1_14
- BCGI18. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y.: Compressing vector OLE, pp. 896–912 (2018)
- BFKL93. Blum, A., Furst, M., Kearns, M., Lipton, R.J.: Cryptographic primitives based on hard learning problems. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 278–291. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_24
- BFKL94. Blum, A., Furst, M.L., Kearns, M.J., Lipton, R.J.: Cryptographic primitives based on hard learning problems, pp. 278–291 (1994)
- BFKR91. Beaver, D., Feigenbaum, J., Kilian, J., Rogaway, P.: Security with low communication overhead. In: Menezes, A.J., Vanstone, S.A. (eds.) CRYPTO 1990. LNCS, vol. 537, pp. 62–76. Springer, Heidelberg (1991). https://doi.org/10.1007/3-540-38424-3_5
- BGI15. Boyle, E., Gilboa, N., Ishai, Y.: Function secret sharing. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 337–367. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_12
- BGI16a. Boyle, E., Gilboa, N., Ishai, Y.: Breaking the circuit size barrier for secure computation under DDH. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9814, pp. 509–539. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53018-4_19
- BGI16b. Boyle, E., Gilboa, N., Ishai, Y.: Function secret sharing: improvements and extensions, pp. 1292–1303 (2016)
- BI05. Barkol, O., Ishai, Y.: Secure computation of constant-depth circuits with applications to database search problems. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 395–411. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_24
- BKS19. Boyle, E., Kohl, L., Scholl, P.: Homomorphic secret sharing from lattices without FHE. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. Homomorphic secret sharing from lattices without FHE, vol. 11477, pp. 3–33. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17656-3_1
- BKW00. Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model, pp. 435–440 (2000)
- BKW03. Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. J. ACM (JACM) **50**(4), 506–519 (2003)

- BLVW19. Brakerski, Z., Lyubashevsky, V., Vaikuntanathan, V., Wichs, D.: Worst-case hardness for LPN and cryptographic hashing via code smoothing. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11478, pp. 619–635. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17659-4_21
- CG97. Chor, B., Gilboa, N.: Computationally private information retrieval (extended abstract), pp. 304–313 (1997)
- CGKS95. Chor, B., Goldreich, O., Kushilevitz, E., Sudan, M.: Private information retrieval, pp. 41–50 (1995)
- Cou19. Couteau, G.: A note on the communication complexity of multiparty computation in the correlated randomness model. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11477, pp. 473–503. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17656-3_17
- DFH12. Damgård, I., Faust, S., Hazay, C.: Secure two-party computation with low communication. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 54–74. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28914-9_4
- DGH+20. Döttling, N., Garg, S., Hajiabadi, M., Masny, D., Wichs, D.: Two-round oblivious transfer from CDH or LPN. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12106, pp. 768–797. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_26
- FGJS17. Fazio, N., Gennaro, R., Jafarikhah, T., Skeith, W.E.: Homomorphic secret sharing from Paillier encryption. In: Okamoto, T., Yu, Y., Au, M.H., Li, Y. (eds.) ProvSec 2017. LNCS, vol. 10592, pp. 381–399. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68637-0_23
- Gen09. Gentry, C.: Fully homomorphic encryption using ideal lattices, pp. 169–178 (2009)
- GI14. Gilboa, N., Ishai, Y.: Distributed point functions and their applications. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 640–658. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_35
- GMW87a. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or a completeness theorem for protocols with honest majority, pp. 218–229 (1987)
- GMW87b. Goldreich, O., Micali, S., Wigderson, A.: How to prove all NP statements in zero-knowledge and a methodology of cryptographic protocol design (Extended Abstract). In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 171–185. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-47721-7_11
- IKM+13. Ishai, Y., Kushilevitz, E., Meldgaard, S., Orlandi, C., Paskin-Cherniavsky, A.: On the power of correlated randomness in secure computation. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 600–620. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36594-2_34
- JKPT12. Jain, A., Krenn, S., Pietrzak, K., Tentes, A.: Commitments and efficient zero-knowledge proofs from learning parity with noise. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 663–680. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34961-4_40
- KO97. Kushilevitz, E., Ostrovsky, R.: Replication is NOT needed: SINGLE database, computationally-private information retrieval, pp. 364–373 (1997)

- Lyu05. Lyubashevsky, V.: The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem. In: Chekuri, C., Jansen, K., Rolim, J.D.P., Trevisan, L. (eds.) APPROX/RANDOM -2005. LNCS, vol. 3624, pp. 378–389. Springer, Heidelberg (2005). https://doi.org/10.1007/11538462_32
- NN01. Naor, M., Nissim, K.: Communication preserving protocols for secure function evaluation, pp. 590–599 (2001)
- Pra62. Prange, E.: The use of information sets in decoding cyclic codes. IRE Trans. Inf. Theory **8**(5), 5–9 (1962)
- SGRR19. Schoppmann, P., Gascón, A., Reichert, L., Raykova, M.: Distributed vector-OLE: improved constructions and implementation, pp. 1055–1072 (2019)
- Wak68. Waksman, A.: A permutation network. J. ACM (JACM) **15**(1), 159–163 (1968)
- Yao86. Yao, A.C.C.: How to generate and exchange secrets (extended abstract), pp. 162–167 (1986)
- YZW+19. Yu, Yu., Zhang, J., Weng, J., Guo, C., Li, X.: Collision resistant hashing from sub-exponential learning parity with noise. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11922, pp. 3–24. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34621-8_1