

# Half-Tree: Halving the Cost of Tree Expansion in COT and DPF

Xiaojie Guo<sup>1,2</sup>[0000-0001-5295-2781], Kang Yang<sup>1</sup>[0000-0002-7453-4043],  
Xiao Wang<sup>3</sup>[0000-0002-5991-7417], Wenhao Zhang<sup>3</sup>[0000-0001-6031-8684],  
Xiang Xie<sup>4,5</sup>[0000-0001-5044-9576], Jiang Zhang<sup>1</sup>[0000-0002-4787-0316], and  
Zheli Liu<sup>2</sup>[0000-0002-2984-2661]

<sup>1</sup> State Key Laboratory of Cryptology, Beijing, China  
yangk@sklc.org, jiangzhang09@gmail.com

<sup>2</sup> Nankai University, Tianjin, China

xiaojie.guo@mail.nankai.edu.cn, liuzheli@nankai.edu.cn

<sup>3</sup> Northwestern University, Evanston, USA

wangxiao@cs.northwestern.edu, wenhao.zhang@northwestern.edu

<sup>4</sup> Shanghai Qi Zhi Institute, Shanghai, China

xiexiangiscas@gmail.com

<sup>5</sup> PADO Labs, Hong Kong, China

**Abstract.** GGM tree is widely used in the design of correlated oblivious transfer (COT), subfield vector oblivious linear evaluation (sVOLE), distributed point function (DPF), and distributed comparison function (DCF). Often, the cost associated with GGM tree dominates the computation and communication of these protocols. In this paper, we propose a suite of optimizations that can reduce this cost by half.

– **Halving the cost of COT and sVOLE.** Our COT protocol introduces extra correlation to each level of a GGM tree used by the state-of-the-art COT protocol. As a result, it reduces both the number of AES calls and the communication by half. Extending this idea to sVOLE, we are able to achieve similar improvement with either halved computation or halved communication.

– **Halving the cost of DPF and DCF.** We propose improved two-party protocols for the distributed generation of DPF/DCF keys. Our tree structures behind these protocols lead to more efficient full-domain evaluation and halve the communication and the round complexity of the state-of-the-art DPF/DCF protocols.

All protocols are provably secure in the random-permutation model and can be accelerated based on fixed-key AES-NI. We also improve the state-of-the-art schemes of puncturable pseudorandom function (PPRF), DPF, and DCF, which are of independent interest in dealer-available scenarios.

## 1 Introduction

The construction of Goldreich-Goldwasser-Micali (GGM) tree [26] yields a pseudorandom function (PRF) family from any length-doubling pseudorandom generator (PRG). In this construction, a PRF key serves as a root and is expanded into a full binary tree, where each non-leaf node defines two child nodes from

Protocol	Computation	Communication	# Rounds
COT (§ 4.1)	2×	2×	—
sVOLE (§ 4.1)	2×	1 ~ 2×	—
sVOLE (§ 4.2)	1.33×	2×	—
DPF (§ 5.2)	1.33×	3×	2×
DCF (§ 5.3)	1.6×	2 ~ 3×	2×

Table 1: **Improvements of our protocols in the random-permutation model.** Computation is measured as the number of fixed-key AES calls. In sVOLE, communication varies as per two field sizes  $|\mathbb{F}|$  and  $|\mathbb{K}|$ . In DCF protocol, communication varies as per the range size  $|\mathcal{R}|$  of comparison functions.

its PRG output. The PRF output for an input bit-string is defined as the leaf node labeled by this bit-string. GGM tree has been adapted widely for various cryptographic applications, especially in recent years.

A recent appealing application of GGM tree is to build efficient pseudorandom correlation generators (PCGs) [8,42,10,46,12,43], e.g., correlated oblivious transfer (COT), subfield vector oblivious linear evaluation (sVOLE), etc. In this context, a GGM tree essentially serves as a puncturable pseudorandom function (PPRF). PCGs serve as essential building blocks for secure multi-party computation (MPC) (e.g., [33,27]), zero-knowledge proofs (e.g., [43,21,2]), private set intersection (e.g., [23,40]), etc. Another related application of GGM tree is to build function secret sharing (FSS). In an FSS scheme, a dealer produces two keys, each defining an additive secret sharing of the full-domain evaluation result of some function  $f$  without revealing the parameters of  $f$ . FSS is very useful even for a simple  $f$ , and the dealer can be emulated using an MPC protocol. A distributed point function (DPF) [25] is an FSS scheme for the family of point functions  $f_{\alpha,\beta}^\bullet(x)$  that output  $\beta$  if  $x = \alpha$  and 0 otherwise. DPF has found various applications, including RAM-based secure computation [22], two-server PIR [25,13], private heavy hitters [6], oblivious linear evaluation (OLE) [12], etc. One important variation of DPF is distributed comparison function (DCF), which is an FSS scheme for the family of comparison functions  $f_{\alpha,\beta}^<(x)$  that output  $\beta$  if  $x < \alpha$  and 0 otherwise. DCF has been applied to design mixed-mode MPC [14,7], secure machine-learning inference [30], etc.

In all applications above, the cost associated with GGM tree can often be significant. For example, in the most recent silent OT protocol [18], distributing GGM-tree-related correlations takes more than 70% of the computation and essentially all communication. Similar bottlenecks have also been observed in DPF. For example, in the DPF-based secure RAM computation [22], local expansion of DPF keys takes a majority of the time as well.

### 1.1 Our Contribution

We propose a suite of *half-trees* as tailored alternatives for several GGM-tree-based protocols, leading to halved computation/communication/round complexity (Table 1, detailed complexity is compared in the sections). Our constructions work in the random-permutation model (RPM) [41,4], which can be efficiently instantiated via, e.g., fixed-key AES-NI.

**Correlated GGM trees for half-cost COT and sVOLE.** We introduce correlated GGM (cGGM), a tree structure leading to both improved computation and communication in COT. It has an invariant that all same-level nodes sum up to the same global offset. We keep this invariant by setting a left child as the hash of its parent and the associated right child as the parent minus the left child. By plugging this tree into the state-of-the-art COT protocols [46,18], we can prove the security of the whole protocol in the random-permutation model by carefully choosing the hash function. Compared to the optimized GGM tree [28], this tree reduces the number of random-permutation calls and the communication by half.

Using cGGM tree, we can realize sVOLE for any large field and its subfield. This protocol reduces the computation of the prior protocols [10,43] by  $2\times$  using a field-based random permutation. However, it only halves the communication when the subfield size is significantly smaller than the field size. Then, we modify our cGGM tree to obtain a pseudorandom correlated GGM (pcGGM) tree, which is similar to a cGGM tree but has pseudorandom leaves. In contrast, pcGGM tree leads to a  $2\times$  saving in communication and a  $1.33\times$  saving in computation.

**Halved communication and round complexity in distributed key generation of DPF and DCF.** We introduce another binary tree structure, which adapts our pcGGM tree into a secretly shared form. This tree leads to a new DPF scheme with an improved distributed key generation protocol. This DPF protocol reduces the computation, communication, and round complexity of the prior work roughly by  $1.33\times$ ,  $3\times$ , and  $2\times$ , respectively. When the range of point functions is a general ring, this shared tree allows simpler secure computation than the prior works in terms of the last correction word.

We also use an extended version of this shared pcGGM tree to design a new DCF scheme also with an improved distributed key generation protocol. The tree expansion in our DCF is much simpler than the prior work [7], where each parent node has to quadruple in length to produce additional correction words. In our extended shared pcGGM tree, this expansion factor in length is two or three, and the resulting additional correction words are more 2PC-friendly. When used in our DCF protocol with typical parameters, this extended tree leads to about  $1.6\times$ ,  $2 \sim 3\times$ , and  $2\times$  savings in terms of computation, communication, and round complexity in contrast to the prior work.

## 1.2 Concurrent Work

Recently, Boyle et al. [9] propose two unpredictable punctured functions (UPFs) that can be converted to PPRF with additional  $0.5N$  RO calls for  $N$ -sized domain. Their first UPF construction needs  $N$  RO calls and is provably secure while the second UPF construction needs  $N$  RP calls but relies on an ad-hoc conjecture. For  $m$ -sized sVOLE tuples, the sVOLE extension protocols based on their proposal either needs  $1.5m$  RO calls, or needs  $m$  RP calls plus  $0.5m$  RO calls. They also propose an sVOLE extension protocol that is based on a stronger variation of UPF and requires  $m$  RO calls in total.

In contrast, our protocol is secure in the random-permutation model without any conjecture. Our COT protocol, as a special case of sVOLE protocol, only requires  $m$  RP calls and can reduce communication by half; our two sVOLE

Assump.		Corr.	Computation	Communication (bits)	
				$P_0 \rightarrow P_1$	$P_1 \rightarrow P_0$
[9]	ROM	sVOLE	$m$ RO calls	$2t(\log \frac{m}{t} - 1)\lambda$ $+ 3t \log  \mathbb{K} $	$t \log  \mathbb{F} $
	Ad-hoc <sup>1</sup>	sVOLE	$m$ RP calls $+ 0.5m$ RO calls		
This work		COT	$m$ RP calls	$t(\log \frac{m}{t} - 1)\lambda + \lambda$	—
	RPM	sVOLE	$m$ RP calls	$t(\log \frac{m}{t} - 1) \log  \mathbb{K} $ $+ \lambda$	$t(\log \frac{m}{t} + 1) \log  \mathbb{F} $
		sVOLE	$1.5m$ RP calls	$t(\log \frac{m}{t} - 2)\lambda$ $+ 3t \log  \mathbb{K}  + \lambda$	$t \log  \mathbb{F} $

<sup>1</sup> Security relies on the conjecture that the punctured result of the RPM-based UPF is unpredictable. This UPF uses GGM-style tree expansion  $G(x) := H_0(x) \parallel H_1(x)$  for  $H_0(x) := H(x) \oplus x$  and  $H_1(x) := H(x) + x \bmod 2^\lambda$ .

Table 2: **Comparison with the concurrent work.** “RO/ROM” (resp., “RP/RPM”) is for random oracle (resp., permutation) and the model.  $P_0$  is the sender with a global key, and  $P_1$  is the receiver. Assume weight- $t$  regular LPN noises in sVOLE extension with output length  $m$ , field  $\mathbb{F}$ , and extension field  $\mathbb{K}$ . Computation is measured by the amount of symmetric-key operations, and there is also LPN-related computation in practice. Communication is measured by assuming  $P_0$  and  $P_1$  have access to random precomputed tuples: (i) [9]:  $t \log \frac{m}{t}$  COTs ( $+ t$  sVOLEs, for general sVOLE extension), (ii) our COT extension:  $t \log \frac{m}{t}$  COTs, (iii) our first sVOLE extension:  $t(\log \frac{m}{t} + 1)$  sVOLEs, and (iv) our second sVOLE extension:  $t \log \frac{m}{t}$  COTs  $+ t$  sVOLEs.

protocols need  $m$  or  $1.5m$  RP calls with different levels of communication reduction. More importantly, we also demonstrate how the idea can be applied to DPF/DCF protocols as well.

In Table 2, we compare the cost of sVOLE extension in the two works. The sVOLE extension in both works can be easily turned into the extension of random OTs via the standard transformation [34, 3, 10]. If we regard one (length-preserving) RO call as two RP calls according to the XOR-based construction of [5], our work also beats the concurrent one in terms of concrete efficiency.

## 2 Preliminaries

### 2.1 Notation

Let  $\lambda$  denote the computational security parameter.  $n = n(\lambda)$  means that  $n \in \mathbb{N}$  is polynomial in  $\lambda$ . Let  $\text{negl}(\cdot)$  denote an unspecified negligible function and  $\log(\cdot)$  denote the logarithm in base 2. Let  $x \leftarrow S$  denote sampling  $x$  uniformly at random from a finite set  $S$ . Let  $[a, b) := \{a, \dots, b-1\}$  and  $[a, b] := \{a, \dots, b\}$ . Let  $\mathbb{G}$  (resp.,  $\mathcal{R}$ ) denote finite group (resp., ring). We use bold lowercase letters (e.g.,  $\mathbf{a}$ ) for vectors. For  $i \geq 0$ , let  $\mathbf{a}^{(i)}$  denote the  $i$ -th entry of vector  $\mathbf{a}$ . Let  $\text{unit}_{\mathbb{G}}(n, \alpha, \beta) \in \mathbb{G}^n$  denote the vector whose  $\alpha$ -th entry is  $\beta$  and others are 0. For some field  $\mathbb{F}$  and irreducible polynomial  $f(X) \in \mathbb{F}[X]$ , let  $\mathbb{K} = \mathbb{F}[X]/f(X)$  denote an extension field. For some  $n \in \mathbb{N}$ , we interchangeably use  $\mathbb{F}_{2^n}$ ,  $\mathbb{F}_2^n$ , and  $\{0, 1\}^n$ , where  $\oplus$  is for bitwise-XOR. For some bit-string  $x \in \{0, 1\}^n$ , let  $\text{lsb}(x)$

denote its least significant bit (LSB),  $\text{hb}(x)$  denote its high  $n - 1$  bits, and  $x_i$  denote its  $i$ -th bit such that  $x_1$  is the most significant one. We use  $\parallel$  for bit-string concatenation and  $\circ$  for function composition. Let  $\text{Convert}_{\mathbb{G}} : \{0, 1\}^* \rightarrow \mathbb{G}$  denote a function that maps random strings to pseudorandom  $\mathbb{G}$  elements (see Appendix F.1 of the full version [29] for its implementation).

**Binary trees.** In an  $n$ -level tree, let  $X_i^j$  denote the  $j$ -th node on its  $i$ -level for  $i \in [1, n]$  and  $j \in [0, 2^i)$ . We can write the superscript  $j$  into  $i$ -bit decomposition, i.e.,  $X_i^{j_1 \dots j_i} := X_i^j$ . When a node  $X_i^j \in \{0, 1\}^n$ , we can decompose it into a seed  $s_i^j := \text{hb}(X_i^j) \in \{0, 1\}^{n-1}$  and a control bit  $t_i^j := \text{lsb}(X_i^j) \in \{0, 1\}$  such that  $X_i^j = (s_i^j \parallel t_i^j)$ . We usually omit the superscript  $j$  if it is the  $i$ -bit prefix of a path  $\alpha \in \{0, 1\}^n$  of particular interest in a given context. For completeness, let  $X_0$  denote the root. For some  $i \in [1, n]$  and  $b \in \{0, 1\}$ , let  $K_i^b$  denote the sum of the  $2^{i-1}$   $b$ -side (i.e., left or right) nodes on the  $i$ -th level.

**Secret sharings.** For some additive Abelian group  $\mathbb{G}$  and  $x \in \mathbb{G}$ , we use  $\langle x \rangle^{\mathbf{A}}$  to mean that  $x$  is additively shared between two parties and call it a secret for short. For some secret  $\langle x \rangle^{\mathbf{A}}$  for  $x \in \mathbb{G}$  and party  $b \in \{0, 1\}$ , let  $\langle x \rangle_b^{\mathbf{A}} \in \mathbb{G}$  denote the secret share of the party  $b$  such that  $x = \langle x \rangle_0^{\mathbf{A}} + \langle x \rangle_1^{\mathbf{A}}$ . We abbreviate  $\langle x \rangle^{\mathbf{A}}$  to  $\langle x \rangle$  and  $\langle x \rangle_b^{\mathbf{A}}$  to  $\langle x \rangle_b$  if  $\mathbb{G} = \{0, 1\}^n$ . For some secret  $\langle x \rangle$  for  $x \in \{0, 1\}^n$  and efficiently computable (possibly non-linear) Boolean circuit  $H : \{0, 1\}^n \rightarrow \{0, 1\}^*$ , let  $H(\langle x \rangle)$  denote such a *linear evaluation* that returns a secret  $\langle y \rangle$  with share  $\langle y \rangle_b := H(\langle x \rangle_b)$  for each  $b \in \{0, 1\}$ .

## 2.2 Security Model and Functionalities

We use the *universal composability* (UC) framework [15] to prove security in the presence of a semi-honest, static adversary. We say that a protocol  $\Pi$  *UC-realizes* an ideal functionality  $\mathcal{F}$  if for any probabilistic polynomial-time (PPT) adversary  $\mathcal{A}$ , there exists a PPT adversary (simulator)  $\mathcal{S}$  such that for any PPT environment  $\mathcal{Z}$  with arbitrary auxiliary input  $z$ , the output distribution of  $\mathcal{Z}$  in the *real-world* execution where the parties interact with  $\mathcal{A}$  and execute  $\Pi$  is computationally indistinguishable from the output distribution of  $\mathcal{Z}$  in the *ideal-world* execution where the parties interact with  $\mathcal{S}$  and  $\mathcal{F}$ .

Our protocols use the functionality  $\mathcal{F}_{\text{sVOLE}}$  (Figure 1) of subfield vector oblivious linear evaluation. If  $\mathbb{K} = \mathbb{F}_{2^\lambda}$  and  $\mathbb{F} = \mathbb{F}_2$ ,  $\mathcal{F}_{\text{sVOLE}}$  degenerates to the COT functionality  $\mathcal{F}_{\text{COT}}$  in [46]. If  $\mathbb{K} = \mathbb{F}$ ,  $\mathcal{F}_{\text{sVOLE}}$  serves as the VOLE functionality in [8, 42, 40]. We omit the session IDs and sub-session IDs in the functionalities for simplicity. By convention, we can write sVOLE tuples as two-party *information-theoretic message authentication codes* (IT-MACs) [38, 20]. Let  $\Delta_b \in \mathbb{K}$  denote the global key of one party  $P_b$ .  $P_b$  authenticates a value  $x \in \mathbb{F}$  of the other party  $P_{1-b}$  by sampling a uniform one-time key  $K_b[x] \leftarrow \mathbb{K}$  and giving to  $P_{1-b}$  the MAC  $M_{1-b}[x] := K_b[x] + x \cdot \Delta_b \in \mathbb{K}$ . If identity  $b \in \{0, 1\}$  is clear in a given context, we write  $\Delta$ ,  $K[x]$ , and  $M[x]$  for  $\Delta_b$ ,  $K_b[x]$ , and  $M_{1-b}[x]$ , respectively.

## 2.3 Circular Correlation Robustness

Circular correlation robustness (CCR) [17, 28] is the security notion first introduced for the circuit garbling with Free-XOR optimization [37], where there exists a global key  $\Delta$  offsetting the inputs and outputs of some function  $H$ . [28]

**Functionality  $\mathcal{F}_{\text{subVOLE}}$**

**Parameters:** Field  $\mathbb{F}$  and its extension field  $\mathbb{K}$ .

**Initialize:** Upon receiving (init) from  $P_0$  and  $P_1$ , sample  $\Delta \leftarrow \mathbb{K}$  if  $P_0$  is honest; otherwise, receive  $\Delta \in \mathbb{K}$  from the adversary. Store  $\Delta$  and send it to  $P_0$ . Ignore all subsequent (init) commands.

**Extend:** This functionality allows polynomially many (extend) commands. Upon receiving (extend,  $m$ ) from  $P_0$  and  $P_1$ :

1. If  $P_0$  is honest, sample  $\mathbf{v} \leftarrow \mathbb{K}^m$ ; otherwise, receive  $\mathbf{v} \in \mathbb{K}^m$  from the adversary.
2. If  $P_1$  is honest, sample  $\mathbf{u} \leftarrow \mathbb{F}^m$ , and compute  $\mathbf{w} := \mathbf{v} + \mathbf{u} \cdot \Delta \in \mathbb{K}^m$ ; otherwise, receive  $(\mathbf{u}, \mathbf{w}) \in \mathbb{F}^m \times \mathbb{K}^m$  from the adversary, and recompute  $\mathbf{v} := \mathbf{w} - \mathbf{u} \cdot \Delta \in \mathbb{K}^m$ .
3. Send  $\mathbf{v}$  to  $P_0$  and  $(\mathbf{u}, \mathbf{w})$  to  $P_1$ .

**Global-key queries:** If  $P_1$  is corrupted, upon receiving (guess,  $\Delta'$ ), where  $\Delta' \in \mathbb{K}$ , from the adversary, send (success) to the adversary if  $\Delta = \Delta'$ ; send (fail) to the adversary otherwise.

Fig. 1: Functionality for subfield VOLE.

showed that a CCR function  $H$  can be constructed from a fixed-key block cipher (e.g., AES) modeled as random permutation and a *linear orthomorphism*<sup>6</sup>. In this construction, it takes one block-cipher call to invoke a CCR function.

**Definition 1 (Circular Correlation Robustness, [28]).** Let  $H : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ ,  $\chi$  be a distribution on  $\{0, 1\}^\lambda$ , and  $\mathcal{O}_{H, \Delta}^{\text{CCR}}(x, b) := H(x \oplus \Delta) \oplus b \cdot \Delta$  be an oracle for  $x, \Delta \in \{0, 1\}^\lambda$  and  $b \in \{0, 1\}$ .  $H$  is  $(t, q, \rho, \epsilon)$ -CCR if, for any distinguisher  $\mathcal{D}$  running in time at most  $t$  and making at most  $q$  queries to  $\mathcal{O}_{H, \Delta}^{\text{CCR}}(\cdot, \cdot)$ , and any  $\chi$  with min-entropy at least  $\rho$ , it holds that

$$\left| \Pr_{\Delta \leftarrow \chi} [\mathcal{D}^{\mathcal{O}_{H, \Delta}^{\text{CCR}}(\cdot, \cdot)}(1^\lambda) = 1] - \Pr_{f \leftarrow \mathcal{F}_{\lambda+1, \lambda}} [\mathcal{D}^f(\cdot, \cdot)(1^\lambda) = 1] \right| \leq \epsilon,$$

where  $\mathcal{D}$  cannot query both  $(x, 0)$  and  $(x, 1)$  for any  $x \in \{0, 1\}^\lambda$ .

In this work,  $\mathcal{D}$  can only make CCR queries with *restricted* forms, which are reminiscent of those in the Half-Gate garbling scheme [47]. We defer the formal definition of these restricted queries to Appendix A of the full version [29].

## 2.4 Function Secret Sharing

A function secret sharing (FSS) is a secret sharing scheme where a dealer distributes the shares of a function  $f$  to multiple parties, and each party can use

<sup>6</sup> A mapping  $\sigma : \mathbb{G} \rightarrow \mathbb{G}$  for an additive Abelian group  $\mathbb{G}$  is a linear orthomorphism if (i)  $\sigma$  is a permutation, (ii)  $\sigma(x + y) = \sigma(x) + \sigma(y)$  for any  $x, y \in \mathbb{G}$ , and (iii)  $\sigma'(x) := \sigma(x) - x$  is also a permutation. [28] presents two efficient instantiations of  $\sigma$  (with well-defined efficient  $\sigma^{-1}$ ,  $\sigma'$ , and  $\sigma'^{-1}$ ): (i) if  $\mathbb{G}$  is a field,  $\sigma(x) := c \cdot x$  for some  $c \neq 0, 1 \in \mathbb{G}$ , and (ii) if  $\mathbb{G} = \{0, 1\}^n$ ,  $\sigma(x) = \sigma(x_L \| x_R) := (x_L \oplus x_R) \| x_L$  where  $x_L$  and  $x_R$  are the left and right halves of  $x$ .

its share to *locally* compute the share of  $f(x)$  for any *public*  $x$  in the domain of  $f$ . In this work, we focus on two-party FSS schemes.

**Definition 2 (Function Secret Sharing, [13,7]).** For a family  $\mathcal{F}_{\mathcal{X},\mathbb{G}}$  of functions with domain  $\mathcal{X}$  and range  $\mathbb{G}$ , where  $\mathbb{G}$  is an Abelian group, a two-party FSS scheme with key space  $\mathcal{K}_0 \times \mathcal{K}_1$  has the following syntax:

- $(k_0, k_1) \leftarrow \text{Gen}(1^\lambda, \hat{f})$ . On input  $1^\lambda$  and the description  $\hat{f} \in \{0,1\}^*$  of a function  $f \in \mathcal{F}_{\mathcal{X},\mathbb{G}}$ , output a key pair  $(k_0, k_1) \in \mathcal{K}_0 \times \mathcal{K}_1$ .
- $f_b(x) \leftarrow \text{Eval}(b, k_b, x)$ . On input the party identifier  $b \in \{0,1\}$ , the party's key  $k_b \in \mathcal{K}_b$ , and a point  $x \in \mathcal{X}$ , output the share  $f_b(x) \in \mathbb{G}$ .

A two-party FSS scheme  $(\text{Gen}, \text{Eval})$  is secure for the function family  $\mathcal{F}_{\mathcal{X},\mathbb{G}}$  with leakage  $\text{Leak} : \{0,1\}^* \rightarrow \{0,1\}^*$  if the following properties hold.

- **Correctness.** For any function  $f \in \mathcal{F}_{\mathcal{X},\mathbb{G}}$  with description  $\hat{f}$ , and any  $x \in \mathcal{X}$ ,

$$\Pr \left[ (k_0, k_1) \leftarrow \text{Gen}(1^\lambda, \hat{f}) : \sum_{b \in \{0,1\}} \text{Eval}(b, k_b, x) = f(x) \right] = 1.$$

- **Security.** There exists a PPT simulator  $\text{Sim}$  such that, for any function  $f \in \mathcal{F}_{\mathcal{X},\mathbb{G}}$  with the description  $\hat{f}$ , any  $b \in \{0,1\}$ , and any PPT adversary  $\mathcal{A}$ ,

$$\left| \Pr \left[ (k_0, k_1) \leftarrow \text{Gen}(1^\lambda, \hat{f}) : \mathcal{A}(1^\lambda, k_b) = 1 \right] - \Pr \left[ k_b \leftarrow \text{Sim}(1^\lambda, b, \text{Leak}(\hat{f})) : \mathcal{A}(1^\lambda, k_b) = 1 \right] \right| \leq \text{negl}(\lambda).$$

By default, the leakage  $\text{Leak}(\hat{f})$  only involves the domain and the range of  $f$ . The following two special FSS schemes have been proposed in [13,7].

**Distributed Point Functions (DPFs).** A two-party *distributed point function* (DPF.Gen, DPF.Eval) with domain  $\mathcal{X}$  and range  $\mathbb{G}$  is a two-party FSS scheme for the function family  $\mathcal{F}_{\mathcal{X},\mathbb{G}} = \{f_{\alpha,\beta}^\bullet\}_{\alpha \in \mathcal{X}, \beta \in \mathbb{G}}$  where  $f_{\alpha,\beta}^\bullet$  is a *point function* such that  $f_{\alpha,\beta}^\bullet(\alpha) = \beta$ , and  $f_{\alpha,\beta}^\bullet(x) = 0$  for  $x \neq \alpha \in \mathcal{X}$ .

**Distributed Comparison Functions (DCFs).** A two-party *distributed comparison function* (DCF.Gen, DCF.Eval) with domain  $\mathcal{X}$  and range  $\mathbb{G}$  is a two-party FSS scheme for the function family  $\mathcal{F}_{\mathcal{X},\mathbb{G}} = \{f_{\alpha,\beta}^<\}_{\alpha \in \mathcal{X}, \beta \in \mathbb{G}}$  where  $f_{\alpha,\beta}^<$  is a *comparison function* such that  $f_{\alpha,\beta}^<(x) = \beta$  if  $x < \alpha \in \mathcal{X}$ , and  $f_{\alpha,\beta}^<(x) = 0$  otherwise.

### 3 Technical Overview

#### 3.1 Improved COT/sVOLE from Correlated GGM Trees

Since COT/sVOLE can be constructed from its “single-point” version using an appropriate LPN assumption, we focus on single-point COT/sVOLE, where the vector  $\mathbf{u}$  in a COT/sVOLE tuple  $\mathbf{w} = \mathbf{v} + \mathbf{u} \cdot \Delta$  has exactly one non-zero entry.

**Correlated OT from correlated GGM.** The core idea behind our single-point COT protocol is that, instead of using a GGM tree with pseudorandom nodes as the state-of-the-art works, our protocol uses a *correlated* GGM (cGGM) tree where *the sum of all same-level nodes equals a global offset  $\Delta$* . This invariant



can be maintained by using a generalized Davies-Meyer construction with a hash function  $H$ : every parent  $x$  has left child  $H(x)$  and right child  $x - H(x)$ . **cGGM** tree leads to two improvements: (i) no additional hash computation is needed for every right child so that the computation is halved, and (ii) if the global offset  $\Delta$  (i.e., the difference between two first-level nodes) is set up by precomputed random COT tuples, the single-point COT protocol sends only  $\lambda$  bits per level, in contrast to  $2\lambda$  bits from a standard OT per level in the state-of-the-art works.

To explain our second improvement in detail, we first recall the prior construction from the perspective of GGM tree. In this construction, the sender holds an  $n$ -level GGM tree, whose  $2^n$  leaves in  $\mathbb{F}_{2^\lambda}$  forms a vector  $\mathbf{v} \in \mathbb{F}_{2^\lambda}^{2^n}$ . The receiver with a punctured point  $\alpha = \alpha_1 \dots \alpha_n \in \{0, 1\}^n$  uses, for each  $i \in [1, n]$ , a standard OT to select the XOR of all  $\bar{\alpha}_i$ -side nodes on the  $i$ -th level. From these  $n$  XORs, the receiver recovers the  $n$  off-path GGM-tree nodes just leaving the path  $\alpha$  and use these  $n$  nodes to recover all leaves except the  $\alpha$ -th one, corresponding to a vector  $\mathbf{w} \in \mathbb{F}_{2^\lambda}^{2^n}$  with the punctured entry  $\mathbf{w}^{(\alpha)}$ . The sender samples  $\Delta \leftarrow \mathbb{F}_{2^\lambda}$ , defines its output as  $(\Delta, \mathbf{v})$ , and sends  $\psi := \Delta \oplus (\oplus_{j \in [0, 2^n)} \mathbf{v}^{(j)}) \in \mathbb{F}_{2^\lambda}$  to the receiver. The receiver patches  $\mathbf{w}^{(\alpha)} := \psi \oplus (\oplus_{j \neq \alpha} \mathbf{w}^{(j)})$  and defines its output as  $(\mathbf{u}, \mathbf{w})$  for  $\mathbf{u} = \mathbf{unit}_{\mathbb{F}_2}(2^n, \alpha, 1)$ . The computation is dominated by the full GGM-tree expansion while the communication is from  $n$  parallel standard OTs, which need  $n$  precomputed COT tuples via the standard technique [34, 3].

In contrast, our **cGGM**-tree single-point COT, where the global offset in a **cGGM** tree coincides with the global key in the  $n$  precomputed COT tuples, can directly use these tuples. For each level  $i \in [1, n]$ , let  $M[r_i] = K[r_i] \oplus r_i \cdot \Delta$  be such a tuple where the sender has  $(\Delta, K[r_i]) \in \mathbb{F}_{2^\lambda} \times \mathbb{F}_{2^\lambda}$  and the receiver has  $(r_i, M[r_i]) \in \mathbb{F}_2 \times \mathbb{F}_{2^\lambda}$ , and  $K_i^b \in \mathbb{F}_{2^\lambda}$  be the XOR of all  $b$ -side nodes for  $b \in \{0, 1\}$ . To select  $K_i^{\bar{\alpha}_i}$  as in the prior construction, the receiver sends  $\bar{\alpha}_i \oplus r_i$  to the sender, receives back  $c_i := K_i^0 \oplus K[r_i] \oplus (\bar{\alpha}_i \oplus r_i) \cdot \Delta$ , and computes

$$c_i \oplus M[r_i] = K_i^0 \oplus K[r_i] \oplus (\bar{\alpha}_i \oplus r_i) \cdot \Delta \oplus M[r_i] = K_i^0 \oplus \bar{\alpha}_i \cdot \Delta = K_i^{\bar{\alpha}_i},$$

where the last equality holds since the **cGGM** tree uses  $\Delta$  as global offset. For each level, the sender sends  $\lambda$  bits to the receiver, only a half of the  $2\lambda$  bits in a standard OT. When the point  $\alpha$  is random, the message  $\bar{\alpha}_i \oplus r_i$  can be avoided as well. The single-point COT outputs are defined as in the prior construction, except that the receiver locally patches  $\mathbf{w}^{(\alpha)} := \oplus_{j \neq \alpha} \mathbf{w}^{(j)}$ .

The security against the semi-honest sender is straightforward. However, a subtle issue arises in proving the security against the semi-honest receiver. Note that the environment  $\mathcal{Z}$  can observe the global key  $\Delta$  from the honest sender's output and use it to distinguish the two worlds. Let  $\{X_i^{\alpha_1 \dots \alpha_{i-1} \bar{\alpha}_i}\}_{i \in [1, n]}$  be the **cGGM**-tree off-path nodes recovered by the receiver. In the real world, these off-path nodes satisfy the consistency with  $\Delta$ : for  $j \in [2, n]$ ,  $X_j^{\alpha_1 \dots \alpha_{j-1} \bar{\alpha}_j}$  equals

$$H\left(\Delta \oplus \bigoplus_{i \in [1, j-1]} X_i^{\alpha_1 \dots \alpha_{i-1} \bar{\alpha}_i}\right) \oplus \bar{\alpha}_j \cdot \left(\Delta \oplus \bigoplus_{i \in [1, j-1]} X_i^{\alpha_1 \dots \alpha_{i-1} \bar{\alpha}_i}\right). \quad (1)$$

However, this consistency does not hold in the ideal world where  $\{c_i\}_{i \in [1, n]}$  sent by the simulator are sampled at random so that the  $n$  off-path nodes will be independently uniform in the ideal world. Thus,  $\mathcal{Z}$  can trivially distinguish the



two worlds by using the known  $\Delta$  to check (1). Our security proof addresses this issue by carefully constructing  $H$  from a random permutation, allowing global-key queries in the single-point COT functionality, and programming the random permutation and its inverse to keep the consistency. The intuition is that, to distinguish the two worlds,  $\mathcal{Z}$  must query the random permutation or its inverse with  $\Delta$ -related transcripts. Thus, the simulator can observe these queries and extract every potential  $\Delta$  from them. Using global-key queries, the simulator checks whether an extracted  $\Delta$  matches that in the single-point COT functionality or not. If so, it immediately programs the two permutation oracles using this  $\Delta$  so that they are consistent with the simulated  $\{c_i\}_{i \in [1, n]}$ . Similar proof technique in the random-oracle model have been used in TinyOT [38, 32].

**Subfield VOLE from correlated GGM.** We further propose a cGGM-based blueprint of single-point sVOLE for field  $\mathbb{F}$  and its exponentially large extension  $\mathbb{K}$ . In this blueprint, we construct an  $n$ -level cGGM tree from a hash function  $H : \mathbb{K} \rightarrow \mathbb{K}$  so that all nodes are in  $\mathbb{K}$ , and extend the spirit of our single-point COT. The spirit is that the equality  $\mathbf{w}^{(\alpha)} = \mathbf{v}^{(\alpha)} \oplus \Delta$  at the punctured point  $\alpha$  automatically holds by embedding  $\Delta$  into a cGGM tree. For single-point sVOLE, we want to likewise keep  $\mathbf{w}^{(\alpha)} = \mathbf{v}^{(\alpha)} + \beta \cdot \Delta$  for some  $\beta \in \mathbb{F}^*$  and  $\Delta \in \mathbb{K}$  at the punctured point  $\alpha$ . However, we cannot use  $\beta \cdot \Delta$ , which is unknown to the sender, as the cGGM-tree global offset. Instead, we can define this offset as the sender's additive share of  $\beta \cdot \Delta$  so that the receiver can correct the automatically preserved result at the point  $\alpha$  by using its additive share of  $\beta \cdot \Delta$ .

In detail, the two parties use a random sVOLE tuple  $M[\beta] = K[\beta] + \beta \cdot \Delta$  for the  $\beta \cdot \Delta$  term, where the sender has  $(\Delta, K[\beta]) \in \mathbb{K} \times \mathbb{K}$  and the receiver has  $(\beta, M[\beta]) \in \mathbb{F}^* \times \mathbb{K}$ . The sender uses  $K[\beta]$  as the global offset of its cGGM tree, and the receiver selects, for each level  $i$ , the sum of all  $\bar{\alpha}_i$ -side nodes. For the  $i$ -th level, let  $K_i^b \in \mathbb{K}$  be the sum of all  $b$ -side nodes for  $b \in \{0, 1\}$ , and let the two parties have access to a special sVOLE tuple<sup>7</sup>  $M[r_i] = K[r_i] + r_i \cdot K[\beta]$ , where the sender has  $K[r_i] \in \mathbb{K}$  and the receiver has  $(r_i, M[r_i]) \in \mathbb{F}_2 \times \mathbb{K}$ . The sender sends  $c_i := K[r_i] + K_i^0 \in \mathbb{K}$  to the receiver, who defines  $\bar{\alpha}_i := r_i$  and can compute

$$(-1)^{r_i} \cdot (-M[r_i] + c_i) = (-1)^{\bar{\alpha}_i} \cdot (K_i^0 - \bar{\alpha}_i \cdot K[\beta]) = K_i^{\bar{\alpha}_i},$$

where the last equality holds due to the cGGM invariant. The  $n$  selected sums allow the receiver to recover, in a top-down manner, the  $n$  off-nodes with respect to  $\alpha$  and the  $2^n$  cGGM leaves except the  $\alpha$ -th one. The sender defines  $\mathbf{v} \in \mathbb{K}^{2^n}$  from its  $2^n$  cGGM-tree leaves, while the receiver defines  $\mathbf{w} \in \mathbb{K}^{2^n}$  from the  $\alpha$ -exclusive  $2^n - 1$  leaves and the locally patched punctured leaf  $\mathbf{w}^{(\alpha)} := M[\beta] - \sum_{j \neq \alpha} \mathbf{w}^{(j)} = M[\beta] - (\sum_{j \neq \alpha} \mathbf{w}^{(j)} + \mathbf{v}^{(\alpha)}) + \mathbf{v}^{(\alpha)} = \mathbf{v}^{(\alpha)} + \beta \cdot \Delta$ . If the sender defines its output as  $(\Delta, \mathbf{v})$  and the receiver defines its output as  $(\mathbf{u}, \mathbf{w})$  for  $\mathbf{u} := \text{unit}_{\mathbb{F}}(2^n, \alpha, \beta)$ , the two parties share a single-point sVOLE correlation.

Our cGGM-based single-point sVOLE protocol also has the issue in proving the security against the semi-honest receiver as the environment  $\mathcal{Z}$  sees  $\Delta$  from the honest sender's output.  $\mathcal{Z}$  can compute the cGGM offset  $K[\beta] = M[\beta] - \beta \cdot \Delta$

<sup>7</sup> The special sVOLE tuples for selecting  $n$  sums can be obtained from  $n$  precomputed random sVOLE tuples by the receiver sending  $n \cdot \log |\mathbb{F}|$  bits.

and, to distinguish the two worlds, check if the consistency (1) holds for  $K[\beta]$  or not. As in our cGGM-based single-point COT, our simulator addresses this issue by extracting every possible  $K[\beta]$  and the associated  $\Delta = \beta^{-1} \cdot (M[\beta] - K[\beta])$ , querying the single-point sVOLE functionality with  $\Delta$ , and programming the random permutation and its inverse if the global-key query succeeds.

**Subfield VOLE from pseudorandom correlated GGM.** There is another single-point sVOLE blueprint [10,43] basing its security on the pseudorandomness of GGM-tree nodes: for some path  $\alpha \in \{0,1\}^n$ , the  $n$  off-path nodes and the  $\alpha$ -th leaf are pseudorandom. Our cGGM tree cannot be used in this blueprint since its same-level nodes are correlated under the global offset. However, we observe that a cGGM tree can be modified into a *pseudorandom* cGGM (pcGGM) tree with the required pseudorandomness.

In an  $n$ -level pcGGM tree, we preserve the cGGM invariant for the  $\mathbb{F}_{2^\lambda}$  nodes on the first  $n-1$  levels, i.e., using a hash function  $H' : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$  and Davies-Meyer construction to keep that all same-level nodes are XORed to a global offset  $\Delta \in \mathbb{F}_{2^\lambda}$ . Nevertheless, we break the last-level correlation in the pcGGM tree: every parent  $x \in \mathbb{F}_{2^\lambda}$  on the  $(i-1)$ -th level has left child  $H'(x)$  and right child  $H'(x \oplus 1)$ . In sVOLE protocols for  $\mathbb{K} \neq \mathbb{F}_{2^\lambda}$ , the pcGGM leaves will be further converted by the function  $\text{Convert}_{\mathbb{K}} : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{K}$ .

Our core observation for arguing the pseudorandomness of the  $n+1$  pcGGM nodes is that the inputs of the hash function  $H'$  are of CCR forms. More specifically, a global  $\Delta \in \mathbb{F}_{2^\lambda}$  offsets the two first-level nodes of the pcGGM tree and induces the first  $n-1$  off-path nodes  $\{X_i^{\alpha_1 \dots \alpha_{i-1} \bar{\alpha}_i}\}_{i \in [1, n-1]}$  according to (1) for  $H'$ . Meanwhile, the last off-path node  $X_n^{\alpha_1 \dots \alpha_{n-1} \bar{\alpha}_n} \in \mathbb{F}_{2^\lambda}$  and the  $\alpha$ -th pcGGM leaf  $X_n \in \mathbb{F}_{2^\lambda}$  come from two hash calls of the following form: for  $b \in \{0,1\}$ ,

$$X_n^{\alpha_1 \dots \alpha_{n-1} b} = H' \left( \Delta \oplus \left( \bigoplus_{i \in [1, n-1]} X_i^{\alpha_1 \dots \alpha_{i-1} \bar{\alpha}_i} \right) \oplus b \right).$$

Intuitively, we can use a CCR hash function  $H'$  to argue the pseudorandomness of the  $n$  off-path nodes and the  $\alpha$ -th leaf, which is sufficient for the single-point sVOLE blueprint. The challenge in this security reduction is to show that the CCR queries to  $H'$  are legal (i.e., no  $(x,0)$  and  $(x,1)$  for the same  $x$ ) with overwhelming probability. We address this challenge by resorting to the observation that these inputs are *restricted* so that they are well-structured and are not arbitrarily chosen by the corrupted receiver (the only case where we need the pseudorandomness). Such restricted inputs are reminiscent of the “naturally derived keys” [47,28] in the Half-Gate garbling scheme so that we can bound the probability similarly. We defer the details to Appendix A of the full version [29]. Note that even if one uses  $\text{Convert}_{\mathbb{K}}$  to map the leaves into  $\mathbb{K}$ , the pseudorandomness of these nodes still holds due to the pseudorandomness of  $\text{Convert}_{\mathbb{K}}$ .

By plugging our pcGGM tree into the prior single-point sVOLE blueprint, we obtain a more efficient protocol. The improvement owes to the cGGM invariant in its first  $n-1$  levels. In terms of communication, the receiver can use  $n-1$  precomputed random COTs to select the XORs on these levels and recover the first  $n-1$  levels of the sender’s pcGGM tree; in contrast, the prior protocols use a standard OT per level due to the two pseudorandom XORs. For the last

level in our protocol, the two parties also need a standard OT due to the broken correlation of the two sums. Given the random-permutation-based CCR hash functions in [28], our pcGGM-based single-point sVOLE protocol is secure in the random-permutation model. In particular, this protocol can implement the single-point sVOLE functionality *without* global-key queries since  $\Delta \in \mathbb{F}_{2^\lambda}$  is only used in the pcGGM tree and is not included in the sender's output.

### 3.2 DPF/DCF from Shared Pseudorandom Correlated GGM Trees

**DPF scheme and protocol.** Using a pcGGM-like trick, we present a new DPF scheme, followed by a more efficient distributed protocol. Recall that, in the prior DPF scheme [13], there are two parties sharing an  $n$ -level GGM-style tree where the  $n$  nodes on some path  $\alpha \in \{0, 1\}^n$  are pseudorandom with LSB one, and others are zero. Then, the two-party shares of the  $\alpha$ -th leaf mask the DPF payload  $\beta \in \mathbb{G}$ . Our core observation is that we need the pseudorandom  $\alpha$ -th leaf to hide  $\beta$ , but the internal pseudorandom on-path nodes are not mandatory. Instead, the two parties can share an  $n$ -level pcGGM-style tree (say, spcGGM tree) where (i) the root  $X_0$  and the first  $n - 1$  on-path nodes equal a global offset  $\Delta \in \mathbb{F}_{2^\lambda}$  with  $\text{lsb}(\Delta) = 1$ , (ii) the last on-path node (i.e., the  $\alpha$ -th leaf) is pseudorandom with LSB one, and (iii) other nodes are zero. As in the prior scheme, the per-party share of this tree is compressed as a key including an XOR share of the root and  $n + 1$  public *pseudorandom* correction words.

We explain our construction of these correction words in detail. To keep the invariant (i), the spcGGM tree uses a correction procedure different from the prior one. For each level  $i \in [1, n - 1]$  with a public correction word  $\text{CW}_i \in \mathbb{F}_{2^\lambda}$ , and  $b \in \{0, 1\}$ , the  $b$ -side secret child of the  $(i - 1)$ -th on-path secret node  $\langle X_{i-1} \rangle = \langle s_{i-1} \parallel t_{i-1} \rangle$  is defined as follows:

$$\langle X_i^{\alpha_1 \dots \alpha_{i-1} b} \rangle := H'(\langle X_{i-1} \rangle) \oplus b \cdot \langle X_{i-1} \rangle \oplus \langle t_{i-1} \rangle \cdot \text{CW}_i.$$

Solving this *linear* equation for the *public*  $\text{CW}_i$  under the constraint (i), we have

$$\text{CW}_i = H'(\langle X_{i-1} \rangle_0) \oplus H'(\langle X_{i-1} \rangle_1) \oplus \bar{\alpha}_i \cdot \Delta.$$

As for (ii), we use a public correction word  $\text{CW}_n = (\text{HCW}, \text{LCW}^0, \text{LCW}^1) \in \mathbb{F}_{2^{\lambda-1}} \times \mathbb{F}_2 \times \mathbb{F}_2$  to follow the same last-level correction as the prior work. For  $b \in \{0, 1\}$ , define a function  $H'_b(\cdot) := H'(\cdot \oplus b)$  and the  $b$ -side secret child of the  $(n - 1)$ -th on-path secret node  $\langle X_{n-1} \rangle = \langle s_{n-1} \parallel t_{n-1} \rangle$  as follows:

$$\langle X_n^{\alpha_1 \dots \alpha_{n-1} b} \rangle := H'_b(\langle X_{n-1} \rangle) \oplus \langle t_{n-1} \rangle \cdot (\text{HCW} \parallel \text{LCW}^b).$$

Solving this *linear* equation for the *public*  $\text{CW}_n$  under the constraint (i) and (iii),

$$\begin{aligned} \text{HCW} &= \text{hb}\left(H'_{\bar{\alpha}_n}(\langle X_{n-1} \rangle_0) \oplus H'_{\bar{\alpha}_n}(\langle X_{n-1} \rangle_1)\right), \\ \forall b \in \{0, 1\} : \text{LCW}^b &= \text{lsb}\left(H'_b(\langle X_{n-1} \rangle_0) \oplus H'_b(\langle X_{n-1} \rangle_1)\right) \oplus \alpha_n \oplus \bar{b}. \end{aligned} \tag{2}$$

Note that the  $n$  off-path secret nodes  $\{\langle X_i^{\alpha_1 \dots \alpha_{i-1} \bar{\alpha}_i} \rangle\}_{i \in [1, n]}$  are zero secrets according to the above correction procedures. As a result, the two parties hold identical shares of these  $n$  off-path nodes and their subtrees, given that the share of a subtree is fully determined by the share of its root (i.e., an off-path node) and

the public correction words. This implies the constraint (iii). Finally, the  $(n+1)$ -th public correction word is defined from the secret  $\alpha$ -th leaf  $\langle X_n \rangle = \langle s_n \parallel t_n \rangle$  and the function  $\text{Convert}_{\mathbb{G}} : \mathbb{F}_{2^{\lambda-1}} \rightarrow \mathbb{G}$  as follows:

$$\text{CW}_{n+1} = (\langle t_n \rangle_0 - \langle t_n \rangle_1) \cdot \left( \text{Convert}_{\mathbb{G}}(\langle s_n \rangle_1) - \text{Convert}_{\mathbb{G}}(\langle s_n \rangle_0) + \beta \right) \in \mathbb{G},$$

where the DPF payload  $\beta$  is masked by the XOR shares of the  $\alpha$ -th leaf.

The DPF security primarily follows from that the first  $n$  correction words are of CCR forms, i.e., for  $i \in [0, n-1]$ ,  $\langle X_i \rangle_0 \oplus \langle X_i \rangle_1 = X_i = \Delta$  according to the XOR secret sharing scheme and the invariant (i). The  $\Delta$ -circular correlation in  $\text{CW}_1, \dots, \text{CW}_{n-1}$  is obvious for either corrupted party. In  $\text{CW}_n$ , the honest party's  $H'$  inputs also differ from the corrupted party's  $H'$  inputs by  $\Delta$ . Intuitively, these  $n$  correction words use CCR responses as one-time pads, and the underlying CCR queries are as structured as those in the original **pcGGM** tree. By using a CCR  $H'$  and upper bounding the probability of illegal CCR queries, we can prove the pseudorandomness of the first  $n$  correction words and the high  $\lambda - 1$  bits (i.e.,  $s_n$ ) of the  $\alpha$ -th leaf. The pseudorandom  $s_n = \langle s_n \rangle_0 \oplus \langle s_n \rangle_1$  and  $\text{Convert}_{\mathbb{G}}$  ensure the pseudorandom  $\text{CW}_{n+1}$  for either corrupted party.

Our DPF scheme enables a more efficient distributed key generation protocol due to the construction of the first  $n - 1$  correction words. The insight is that the two parties, who share  $\langle \alpha \rangle$  and  $\langle \beta \rangle^A$ , can use their precomputed COT tuples to set up a secret  $\langle \Delta \rangle$  with  $\text{lsb}(\Delta) = 1$  and share  $\{\langle \bar{\alpha}_i \cdot \Delta \rangle\}_{i \in [1, n]}$  in two rounds, and use the black-box evaluation technique in [22] to locally share each secret  $H'(\langle X_{i-1} \rangle)$ . This technique relies on the invariant (iii) so that, for each  $i \in [1, n]$ , summing the shares of the  $2^i$  nodes on the  $i$ -th level returns the share of the  $i$ -th level on-path node. Given the two-party shares of  $\langle \bar{\alpha}_i \cdot \Delta \rangle$  and  $H'(\langle X_{i-1} \rangle)$ , the secure computation of each  $\text{CW}_i$  only needs one round for revealing  $\langle \text{CW}_i \rangle$ , leading to  $n - 1$  rounds for the first  $n - 1$  correction words in total. In contrast, the prior protocol [22] uses (2) for each correction word, and the  $i$ -th level HCW depends on  $\bar{\alpha}_i$  and should be computed level-by-level. Thus, it securely computes the first  $n - 1$  correction words in  $2(n - 1)$  rounds: for each level, one round is to share  $\langle \text{CW}_i \rangle$  from standard OTs, and another round is to reveal this secret.

We remark that our  $\text{CW}_{n+1}$  construction uses  $\langle t_n \rangle_0 - \langle t_n \rangle_1$  to replace the  $(-1)^{\langle t_n \rangle_1}$  term in the prior construction. The correctness is unaffected due to the non-zero LSB (i.e.,  $t_n$ ) of the  $\alpha$ -th leaf. However, when  $\mathbb{G}$  is a ring, our  $\text{CW}_{n+1}$  allows the two parties to locally share  $\langle t_n \rangle_0 - \langle t_n \rangle_1$  on the ring via the black-box evaluation technique [22]. Thus, the secure computation of  $\text{CW}_{n+1}$  uses only one secure multiplication of two locally shared ring operands.

**DCF scheme and protocol.** We further show that our **spcGGM** tree can be extended to realize more efficient DCF scheme and its distributed protocol. Note that comparison function  $f_{\alpha, \beta}^<(x)$  can be written as the sum of point function  $f_{\alpha, -\alpha_n \cdot \beta}^\bullet(x)$  and a prefix function  $V_{\alpha, \beta}(x)$ , which returns  $\alpha_{h+1} \cdot \beta \in \mathbb{G}$  such that  $\alpha_1 \dots \alpha_h = x_1 \dots x_h$  is the longest common prefix of  $\alpha$  and  $x$  (for completeness,  $\alpha_{n+1} := \alpha_n$ ). We have shown how to realize the DPF scheme for point function  $f_{\alpha, -\alpha_n \cdot \beta}^\bullet(x)$  from **spcGGM** tree. Then, we want to compute  $V_{\alpha, \beta}(x)$  by reusing the prefix information with respect to  $\alpha$  and  $x$  when traversing the **spcGGM**

tree to evaluate the point function. Following the GGM-style DCF scheme [7], we do this by introducing more nodes to the spcGGM tree and an additional correction procedure to ensure that the sum of the introduced nodes along the path  $x$  equals  $V_{\alpha,\beta}(x)$ . However, our extended spcGGM tree can use less nodes and simpler correction words to compute  $V_{\alpha,\beta}(x)$ .

To give more details, we first recall how [7] works. It extends a shared GGM tree by replacing its length-doubling PRG with a length-quadrupling PRG so that each secret parent spawns two more secret children in  $\mathbb{F}_{2^\lambda}$ . For each level  $i \in [1, n]$ , let  $\langle v_i^0 \rangle$  and  $\langle v_i^1 \rangle$  denote such two secret children of the  $(i-1)$ -th on-path secret parent  $\langle X_{i-1} \rangle = \langle s_{i-1} \parallel t_{i-1} \rangle$ , and the two parties correct their additive shares for  $V_{\alpha,\beta}(x)$  via the public correction word  $\mathbf{VCW}_i$ :

$$\begin{aligned} V_{i-1} &:= \sum_{b \in \{0,1\}} (-1)^{1-b} \cdot \left( \text{Convert}_{\mathbb{G}}(\langle v_{i-1}^{\bar{\alpha}_{i-1}} \rangle_b) - \text{Convert}_{\mathbb{G}}(\langle v_{i-1}^{\alpha_{i-1}} \rangle_b) \right) \in \mathbb{G}, \\ \mathbf{VCW}_i &:= (-1)^{\langle t_{i-1} \rangle_1} \cdot \left( \left( \text{Convert}_{\mathbb{G}}(\langle v_i^{\bar{\alpha}_i} \rangle_1) - \text{Convert}_{\mathbb{G}}(\langle v_i^{\bar{\alpha}_i} \rangle_0) \right) \right. \\ &\quad \left. - V_{i-1} + (\alpha_i - \alpha_{i-1}) \cdot \beta \right) \in \mathbb{G}. \quad (V_0 := 0 \in \mathbb{G}, \alpha_0 = 0) \end{aligned}$$

The DCF key per party includes its DPF key for  $f_{\alpha, -\alpha_n, \beta}^\bullet(x)$  and  $\{\mathbf{VCW}_i\}_{i \in [1, n]}$ . The DCF security also requires the pseudorandomness of the  $n$   $\mathbf{VCW}_i$ 's.

In contrast, our DCF scheme shows that it is overkill to introduce two more secret children to each secret parent for the DCF security. For each  $i \in [1, n]$ , one additional secret child  $\langle v_i \rangle = \langle v_i^0 \rangle = \langle v_i^1 \rangle$  of the secret parent  $\langle X_{i-1} \rangle$  suffices, and the pseudorandomness of  $\mathbf{VCW}_i$  relies on a random  $v_i = \langle v_i \rangle_0 \oplus \langle v_i \rangle_1 \in \mathbb{F}_{2^\lambda}$  as  $\text{Convert}_{\mathbb{G}}$  maps random strings to pseudorandom  $\mathbb{G}$  elements. We can argue the pseudorandomness of  $v_i$  based on the CCR induced by  $X_{i-1} = \Delta$ , if we use  $v_i := \mathbf{H}'(\langle X_{i-1} \rangle_0 \oplus 2) \oplus \mathbf{H}'(\langle X_{i-1} \rangle_1 \oplus 2)$ . Collecting all  $\mathbf{H}'$  inputs for the DPF part and  $v_i$ 's, we find that these inputs are as structured as those in the original pcGGM tree. The DCF security can follow from a similar hybrid argument.

Our DCF protocol is extended from our DPF protocol with the additional secure computation of  $\{\mathbf{VCW}_i\}_{i \in [1, n]}$ . Compared with the prior work, our DCF protocol achieves better efficiency due to not only its optimized DPF part but also the structure of each  $\mathbf{VCW}_i$ . This structure makes the  $\text{Convert}_{\mathbb{G}}$  difference term independent of  $\bar{\alpha}_i$ . This independence allows the two parties to locally share the  $\text{Convert}_{\mathbb{G}}$  difference via the black-box evaluation technique [22], in contrast to the technique plus OT-based 2PC in the prior protocol. Since there is only one more secret child for each secret parent, the local computation for sharing this difference is halved as well. We can also replace the  $(-1)^{\langle t_{i-1} \rangle_1}$  term in the prior  $\mathbf{VCW}_i$  construction by a linear term  $\langle t_{i-1} \rangle_0 - \langle t_{i-1} \rangle_1$ , which can be locally shared via the same black-box evaluation technique if  $\mathbb{G}$  is a ring. As a result, except the 2PC for sharing  $\{\langle \alpha_i \cdot \beta \rangle^A\}_{i \in [1, n]}$ , the secure computation of  $\{\mathbf{VCW}_i\}_{i \in [1, n]}$  requires  $n$  secure multiplications of two shared ring elements. These secure multiplications can run in parallel with that for  $\mathbf{CW}_{n+1}$ .

In our DCF protocol, each  $\langle \alpha_i \cdot \beta \rangle^A$  is secretly shared by carefully reusing the two precomputed COT tuples, which were used to share  $\langle \bar{\alpha}_i \cdot \Delta \rangle$ , to run a COT-based multiplication between the XOR shared  $\alpha_i$  and the additively shared

**Functionality  $\mathcal{F}_{\text{spsVOLE}}$**

**Parameters:** Field  $\mathbb{F}$  and its extension field  $\mathbb{K}$ .

**Initialize:** Upon receiving (init) from  $P_0$  and  $P_1$ , sample  $\Delta \leftarrow \mathbb{K}$  if  $P_0$  is honest; otherwise, receive  $\Delta \in \mathbb{K}$  from the adversary. Store  $\Delta$  and send it to  $P_0$ . Ignore all subsequent (init) commands.

**Extend:** This functionality allows polynomially many (extend) commands. Upon receiving (extend,  $N$ ) from  $P_0$  and  $P_1$ :

1. If  $P_0$  is honest, sample  $\mathbf{v} \leftarrow \mathbb{K}^N$ ; otherwise, receive  $\mathbf{v} \in \mathbb{K}^N$  from the adversary.
2. If  $P_1$  is honest, sample  $\mathbf{u} \leftarrow \mathbb{F}^N$  with exactly one nonzero entry, and compute  $\mathbf{w} := \mathbf{v} + \mathbf{u} \cdot \Delta \in \mathbb{K}^N$ ; otherwise, receive  $(\mathbf{u}, \mathbf{w}) \in \mathbb{F}^N \times \mathbb{K}^N$  from the adversary, where  $\mathbf{u}$  has at most one nonzero entry, and recompute  $\mathbf{v} := \mathbf{w} - \mathbf{u} \cdot \Delta \in \mathbb{K}^N$ .
3. Send  $\mathbf{v}$  to  $P_0$  and  $(\mathbf{u}, \mathbf{w})$  to  $P_1$ .

**Global-key queries:** If  $P_1$  is corrupted, upon receiving (guess,  $\Delta'$ ), where  $\Delta' \in \mathbb{K}$ , from the adversary, send (success) to the adversary if  $\Delta = \Delta'$ ; send (fail) to the adversary otherwise.

Fig. 2: Functionality for single-point subfield VOLE.

$\beta$  on the ring. This multiplication generalizes the binary case [1, 28] for an XOR shared bit and an XOR shared string by using the well-known arithmetic XOR on the ring:  $\langle \alpha_i \rangle_0 \oplus \langle \alpha_i \rangle_1 = \langle \alpha_i \rangle_0 + \langle \alpha_i \rangle_1 - 2 \cdot \langle \alpha_i \rangle_0 \cdot \langle \alpha_i \rangle_1$ .

## 4 Subfield VOLE Extension

Our sVOLE extension follows the blueprint of [10, 42, 46, 43], which uses LPN to locally convert  $t$  single-point sVOLE (spsVOLE) tuples output by functionality  $\mathcal{F}_{\text{spsVOLE}}$  (Figure 2) into an sVOLE tuple. We focus on the efficient spsVOLE protocol that UC-realizes  $\mathcal{F}_{\text{spsVOLE}}$ . Note that the spsVOLE protocol dominates the computation and contributes all communication in sVOLE extension.

$\mathcal{F}_{\text{spsVOLE}}$  is parameterized by a field  $\mathbb{F}$  and its extension  $\mathbb{K}$ , and covers the single-point COT functionality  $\mathcal{F}_{\text{spCOT}}$  if  $\mathbb{F} = \mathbb{F}_2$  and  $\mathbb{K} = \mathbb{F}_{2^\lambda}$ . This functionality is the same as that in [43], except that  $\mathcal{F}_{\text{spsVOLE}}$  will not abort for an incorrect global-key query. Allowing for global-key queries has been considered in [38, 32] and does not weaken the effective security. In the spsVOLE protocol based on pseudorandom correlated GGM, such global-key queries can be removed.

In essence, our spsVOLE protocols work as the PCG protocol [11, 10, 12, 18] of spsVOLE correlation, although we do not divide the correlation generation into two explicit PCG phases. In Appendix E.1 of the full version [29], we show how to modify one of our spsVOLE protocols to define such two phases, in order to satisfy the “silent property” that a long spsVOLE tuple can be stored as two sublinearly short correlated seeds.

### 4.1 Single-point COT and sVOLE from Correlated GGM

In Figure 3, we present the two evaluation algorithms for our correlated GGM tree, which is defined by two first-level nodes  $(k, \Delta - k) \in \mathbb{K}^2$ . For every non-leaf

**Parameters:** Tree depth  $n \in \mathbb{N}$ . Field  $\mathbb{K}$  with  $|\mathbb{K}| \geq 2^\lambda$ . Hash function  $\mathbf{H} : \mathbb{K} \rightarrow \mathbb{K}$ .

**cGGM.FullEval** $(\Delta, k)$ : Given  $(\Delta, k) \in \mathbb{K}^2$ ,

- 1:  $X_1^0 := k \in \mathbb{K}$ ,  $X_1^1 := \Delta - k \in \mathbb{K}$ .
- 2: **for**  $i \in [2, n]$ ,  $j \in [0, 2^{i-1})$  **do**
- 3:      $X_i^{2j} := \mathbf{H}(X_{i-1}^j) \in \mathbb{K}$ ,  $X_i^{2j+1} := X_{i-1}^j - X_i^{2j} \in \mathbb{K}$ .
- 4:  $\mathbf{v} := (X_n^0, \dots, X_n^{2^n-1}) \in \mathbb{K}^{2^n}$ .
- 5: **for**  $i \in [1, n]$  **do**  $K_i^0 := \sum_{j \in [0, 2^{i-1})} X_i^{2j} \in \mathbb{K}$ .
- 6: **return**  $(\mathbf{v}, \{K_i^0\}_{i \in [1, n]})$

**cGGM.PuncFullEval** $(\alpha, \{K_i^{\bar{\alpha}_i}\}_{i \in [1, n]})$ : Given  $(\alpha, \{K_i^{\bar{\alpha}_i}\}_{i \in [1, n]}) \in \{0, 1\}^n \times \mathbb{K}^n$ ,

- 1:  $X_1^{\bar{\alpha}_1} := K_1^{\bar{\alpha}_1} \in \mathbb{K}$ .
- 2: **for**  $i \in [2, n]$  **do**
- 3:     **for**  $j \in [0, 2^{i-1})$ ,  $j \neq \alpha_1 \dots \alpha_{i-1}$  **do**
- 4:          $X_i^{2j} := \mathbf{H}(X_{i-1}^j) \in \mathbb{K}$ ,  $X_i^{2j+1} := X_{i-1}^j - X_i^{2j} \in \mathbb{K}$ .
- 5:      $X_i^{\alpha_1 \dots \alpha_{i-1} \bar{\alpha}_i} := K_i^{\bar{\alpha}_i} - \sum_{j \in [0, 2^{i-1}), j \neq \alpha_1 \dots \alpha_{i-1}} X_i^{2j + \bar{\alpha}_i} \in \mathbb{K}$ .
- 6:  $X_n^\alpha := - \sum_{j \in [0, 2^n), j \neq \alpha} X_n^j \in \mathbb{K}$ ,  $\mathbf{w} := (X_n^0, \dots, X_n^{2^n-1}) \in \mathbb{K}^{2^n}$ .
- 7: **return**  $\mathbf{w}$

Fig. 3: Two full-evaluation algorithms for correlated GGM tree.

node  $x \in \mathbb{K}$ , its left child is defined as  $\mathbf{H}(x) \in \mathbb{K}$  while its right child is defined as  $x - \mathbf{H}(x) \in \mathbb{K}$ . The following claim is straightforward from an induction.

*Claim (Leveled correlation).* For any two first-level nodes  $(k, \Delta - k) \in \mathbb{K}^2$  and any  $i \in [1, n]$ , the offset  $\Delta \in \mathbb{K}$  equals the sum of all nodes on the  $i$ -th level of the correlated GGM tree expanded from  $(k, \Delta - k)$  as per **cGGM.FullEval**.

**Corollary 1.** For any  $\alpha \in [0, 2^n)$ , any  $(k, \Delta - k) \in \mathbb{K}^2$ , and

$$(\mathbf{v}, \{K_i^0\}_{i \in [1, n]}) := \text{cGGM.FullEval}(\Delta, k),$$

$$\mathbf{w} := \text{cGGM.PuncFullEval}(\alpha, \{K_i^{\bar{\alpha}_i}\}_{i \in [1, n]}),$$

where  $K_i^{\bar{\alpha}_i} := \bar{\alpha}_i \cdot \Delta + (-1)^{\bar{\alpha}_i} \cdot K_i^0$  for  $i \in [1, n]$ , we have  $\mathbf{w}^{(\alpha)} - \mathbf{v}^{(\alpha)} = -\Delta$ .

*Proof.* Claim 4.1 and the definition of **cGGM.FullEval** imply that  $K_i^{\bar{\alpha}_i} \in \mathbb{K}$  in this corollary defines the sum of all  $\bar{\alpha}_i$ -side nodes on the  $i$ -th level of the correlated GGM tree. Then, it follows from the definition of **cGGM.PuncFullEval** that  $\mathbf{v}^{(j)} = \mathbf{w}^{(j)}$  for any  $j \neq \alpha \in [0, 2^n)$ . Using Claim 4.1 for the last level, we have  $\mathbf{w}^{(\alpha)} - \mathbf{v}^{(\alpha)} = - \sum_{j \in [0, 2^n), j \neq \alpha} \mathbf{w}^{(j)} - \mathbf{v}^{(\alpha)} = - \sum_{j \in [0, 2^n), j \neq \alpha} \mathbf{v}^{(j)} - \mathbf{v}^{(\alpha)} = -\Delta$ .

**Single-point COT** Figure 4 describes our single-point COT protocol  $\Pi_{\text{spCOT}}$  that runs in the  $\mathcal{F}_{\text{COT}}$ -hybrid model and uses the cGGM expansion in Figure 3.

**The same  $\Delta$  in correlated GGM trees.** Note that  $\mathcal{F}_{\text{spCOT}}$  produces single-point COT tuples with the same global key  $\Delta \in \mathbb{F}_{2^\lambda}$  in a number of **Extend** executions. To realize  $\mathcal{F}_{\text{spCOT}}$ , our protocol  $\Pi_{\text{spCOT}}$  proceeds as sketched in Section 3.1 but uses the same  $\Delta$  for the cGGM trees of these executions, each of which samples a fresh  $k \leftarrow \mathbb{F}_{2^\lambda}$  for **cGGM.FullEval** $(\Delta, k)$ . A merit of using the



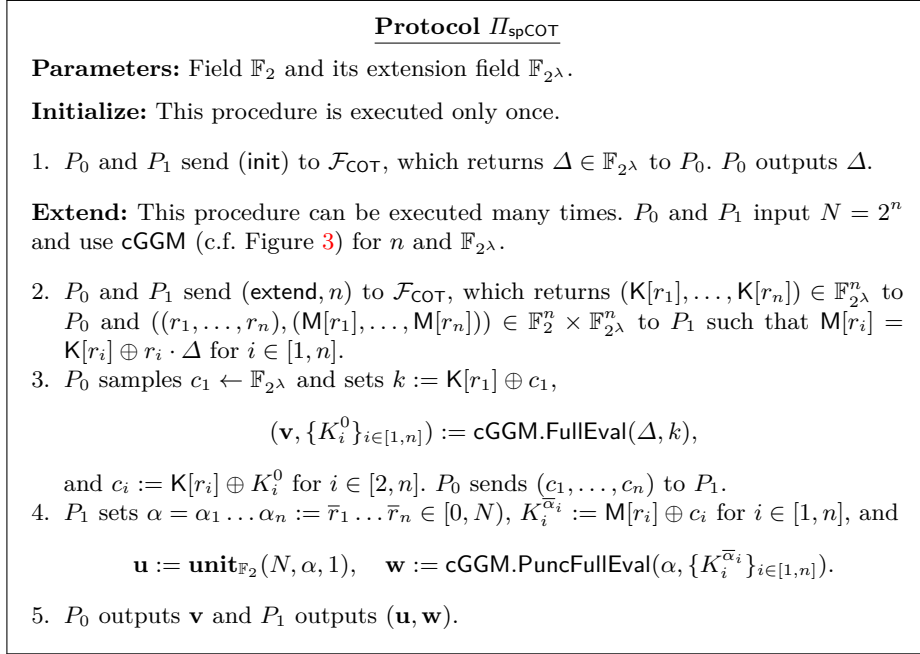


Fig. 4: cGGM-based single-point COT protocol in the  $\mathcal{F}_{\text{COT}}$ -hybrid model.

same  $\Delta$  in several tree instances is that  $\Pi_{\text{spCOT}}$  only invokes one  $\mathcal{F}_{\text{COT}}$  instance, and the amortized cost per precomputed COT tuple can be small.

**Security.** We prove Theorem 1 by following the sketched intuition in Section 3.1 and defer the proof to Appendix B.1 of the full version [29]. Our proof considers polynomially many concurrent **Extend** executions (strictly speaking, sub-sessions with unique sub-session IDs) that uses the one-time initialized  $\Delta$ .

**Theorem 1.** *Given random permutation  $\pi : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ , efficiently computable linear orthomorphism  $\sigma : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$  with efficiently computable  $\sigma^{-1}$ ,  $\sigma'(x) := \sigma(x) \oplus x$ , and  $\sigma'^{-1}$  (Footnote 6), and hash function  $H(x) := \pi(\sigma(x)) \oplus \sigma(x)$ , protocol  $\Pi_{\text{spCOT}}$  (Figure 4) UC-realizes functionality  $\mathcal{F}_{\text{spCOT}}$  (Figure 2) against any semi-honest adversary in the  $\mathcal{F}_{\text{COT}}$ -hybrid model and the RPM.*

**Communication optimization.** For  $t$  concurrent **Extend** executions (e.g., in COT extension), the random  $c_1$ 's in these executions can be compressed via a PRF  $F : \mathbb{F}_{2^\lambda} \times \{0, 1\}^* \rightarrow \mathbb{F}_{2^\lambda}$ . Concretely,  $P_0$  samples a PRF key  $k_{\text{prf}} \leftarrow \mathbb{F}_{2^\lambda}$  after receiving its COT outputs in all executions and sends this key to  $P_1$ . For each execution with sub-session ID  $\text{ssid}$ , the two parties locally defines the element  $c_1 := F(k_{\text{prf}}, \text{ssid})$ . This PRF key is only used for the  $t$  concurrent executions. The security of this optimization follows from the PRF security and the fact that, in the concurrent executions, the COT messages chosen by the corrupted receiver cannot depend on the PRF key to be sampled by the honest sender.

**Complexity analysis.** Consider the complexity per execution when the PRF-based optimization is used in  $t$  concurrent **Extend** executions.  $\Pi_{\text{spCOT}}$  needs  $n$

precomputed COT tuples.  $P_0$  sends  $(n - 1) \cdot \lambda + \frac{\lambda}{t}$  bits, and  $P_1$  sends nothing. The computation per party comes from the tree expansion with  $N$  RP calls.

In the  $\mathcal{F}_{\text{COT}}$ -hybrid model, the prior single-point COT protocol [46] consumes  $n$  precomputed COT tuples. However,  $P_0$  sends  $2n \cdot \lambda$  bits. Each party performs about  $N$  length-doubling PRG calls, which in turn result in  $2N$  RP calls. We can see that our protocol halves both the computation and communication in the prior work. When looking at the whole protocol, the improvement is still huge. For example, the micro benchmark in Silver [18] reported that 70% of the time is spent on GGM-tree-related computation, and thus our protocol will lead to at least 50% of end-to-end computational improvement in COT.

**Single-point sVOLE** We can also realize single-point sVOLE from our cGGM tree by using the high-level idea sketched in Section 3.1. This protocol extends  $\Pi_{\text{spCOT}}$  by using a cGGM tree whose nodes are in a general *exponentially large* extension field  $\mathbb{K}$ . The tree expansion therein uses a hash function constructed from a random permutation and a linear orthomorphism over  $\mathbb{K}$ . We defer the detailed protocol and its security proof to Appendix B.2 of the full version [29].

#### 4.2 Single-point sVOLE from Pseudorandom Correlated GGM

We can adapt our correlated GGM tree for a *pseudorandom* correlated one with the property that the leaf node at some punctured position  $\alpha$  is pseudorandom. This pseudorandom correlated GGM tree **pcGGM** is defined in Figure 5, where the first  $n - 1$  levels preserve the correlation in Claim 4.1 but all last-level nodes are processed by  $H_S$  to break this correlation. The keyed hash function  $H_S$  uses some key  $S \in \mathbb{F}_{2^\lambda}$ , which can be sampled by the receiver in single-point sVOLE and, for simplicity, is assumed to have been sent to the sender before protocol execution. The implementation of  $H_S$  is given in Theorem 2. In fact, this **pcGGM** tree yields PPRF, which is proved in Appendix C of the full version [29].

The pseudorandomness only at the cost of the last-level correlation allows us to follow the single-point sVOLE blueprint in [10, 43] but also take advantage of the correlation in the first  $n - 1$  levels. The protocol is presented in Figure 6. In this protocol, the sender  $P_0$  only sends  $\lambda$  bits to the receiver  $P_1$  for each of the first  $n - 1$  levels, given a precomputed COT tuple. For the last level, the two parties use a COT tuple and the standard technique [34, 3] to emulate the string OT as in the prior protocols. To amortize the cost per precomputed COT tuple, the **pcGGM** trees in many **Extend** executions also use the same  $\Delta$  set by  $\mathcal{F}_{\text{COT}}$ .

**Security.** The security against the semi-honest  $P_0$  resorts to the one-time pad  $s$  from  $\mathcal{F}_{\text{sVOLE}}$ . Meanwhile, the security against the semi-honest  $P_1$  relies on that (i) the **pcGGM** tree with a CCR structure has  $n$  pseudorandom off-path nodes and the punctured leaf, giving pseudorandom  $c_1, \dots, c_{n-1}$  and  $(c_n^r, \psi)$ , and (ii) the mask of the unselected message  $c_n^{\bar{r}}$  in the emulated last-level OT is computed by applying  $\text{Convert}_{\mathbb{K}}$  to a CCR response, which is for a legal CCR query with overwhelming probability due to the uniform  $\mu$ . The proof of Theorem 2 can be found in Appendix B.3 of the full version [29], where we consider polynomially many concurrent **Extend** executions, which use the one-time initialized  $\Delta$ .

**Parameters:** Tree depth  $n \in \mathbb{N}$ . Field  $\mathbb{K}$ . Keyed hash function  $H_S : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ . Function  $\text{Convert}_{\mathbb{K}} : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{K}$ .

**pcGGM.FullEval**( $\Delta, k$ ): Given  $(\Delta, k) \in \mathbb{F}_{2^\lambda}^2$ ,

- 1:  $X_1^0 := k \in \mathbb{F}_{2^\lambda}, X_1^1 := \Delta \oplus k \in \mathbb{F}_{2^\lambda}$ .
- 2: **for**  $i \in [2, n-1]$ ,  $j \in [0, 2^{i-1})$  **do**
- 3:  $X_i^{2j} := H_S(X_{i-1}^j) \in \mathbb{F}_{2^\lambda}, X_i^{2j+1} := X_{i-1}^j \oplus X_i^{2j} \in \mathbb{F}_{2^\lambda}$ .
- 4: **for**  $j \in [0, 2^{n-1})$ ,  $b \in \{0, 1\}$  **do**  $X_n^{2j+b} := \text{Convert}_{\mathbb{K}}(H_S(X_{n-1}^j \oplus b)) \in \mathbb{K}$ .
- 5:  $\mathbf{v} := (X_n^0, \dots, X_n^{2^n-1}) \in \mathbb{K}^{2^n}$ .
- 6: **for**  $i \in [1, n-1]$  **do**  $K_i^0 := \bigoplus_{j \in [0, 2^{i-1})} X_i^{2j} \in \mathbb{F}_{2^\lambda}$ .
- 7:  $(K_n^0, K_n^1) := (\sum_{j \in [0, 2^{n-1})} X_n^{2j}, \sum_{j \in [0, 2^{n-1})} X_n^{2j+1}) \in \mathbb{K}^2$ .
- 8: **return**  $(\mathbf{v}, \{K_i^0\}_{i \in [1, n-1]}, (K_n^0, K_n^1))$

**pcGGM.PuncFullEval**( $\alpha, \{K_i^{\bar{\alpha}_i}\}_{i \in [1, n]}, \gamma$ ): Given  $(\alpha, \{K_i^{\bar{\alpha}_i}\}_i, \gamma) \in \{0, 1\}^n \times \mathbb{K}^n \times \mathbb{K}$ ,

- 1:  $X_1^{\bar{\alpha}_1} := K_1^{\bar{\alpha}_1} \in \mathbb{F}_{2^\lambda}$ .
- 2: **for**  $i \in [2, n-1]$  **do**
- 3: **for**  $j \in [0, 2^{i-1})$ ,  $j \neq \alpha_1 \dots \alpha_{i-1}$  **do**
- 4:  $X_i^{2j} := H_S(X_{i-1}^j) \in \mathbb{F}_{2^\lambda}, X_i^{2j+1} := X_{i-1}^j \oplus X_i^{2j} \in \mathbb{F}_{2^\lambda}$ .
- 5:  $X_i^{\alpha_1 \dots \alpha_{i-1} \bar{\alpha}_i} := K_i^{\bar{\alpha}_i} \oplus (\bigoplus_{j \in [0, 2^{i-1}), j \neq \alpha_1 \dots \alpha_{i-1}} X_i^{2j+\bar{\alpha}_i}) \in \mathbb{F}_{2^\lambda}$ .
- 6: **for**  $j \in [0, 2^{n-1})$ ,  $j \neq \alpha_1 \dots \alpha_{n-1}$ ,  $b \in \{0, 1\}$  **do**
- 7:  $X_n^{2j+b} := \text{Convert}_{\mathbb{K}}(H_S(X_{n-1}^j \oplus b)) \in \mathbb{K}$ .
- 8:  $X_n^{\alpha_1 \dots \alpha_{n-1} \bar{\alpha}_n} := K_n^{\bar{\alpha}_n} - \sum_{j \in [0, 2^{n-1}), j \neq \alpha_1 \dots \alpha_{n-1}} X_n^{2j+\bar{\alpha}_n} \in \mathbb{K}$ .
- 9:  $X_n^\alpha := \gamma - \sum_{j \in [0, 2^n), j \neq \alpha} X_n^j \in \mathbb{K}, \mathbf{w} := (X_n^0, \dots, X_n^{2^n-1}) \in \mathbb{K}^{2^n}$ .
- 10: **return**  $\mathbf{w}$

Fig. 5: Two full-evaluation algorithms for pseudorandom correlated GGM tree.

**Theorem 2.** Given CCR function  $H : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ , function  $\text{Convert}_{\mathbb{K}} : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{K}$ , and keyed hash function  $H_S(x) := H(S \oplus x)$  with some key  $S \leftarrow \mathbb{F}_{2^\lambda}$ , protocol  $\Pi_{\text{spsVOLE-pcGGM}}$  (Figure 6) UC-realizes functionality  $\mathcal{F}_{\text{spsVOLE}}$  (Figure 2) without global-key queries against any semi-honest adversary in the  $(\mathcal{F}_{\text{COT}}, \mathcal{F}_{\text{sVOLE}})$ -hybrid model.

**Communication optimization.**  $\Pi_{\text{spsVOLE-pcGGM}}$  can be optimized as follows:

- The two random  $(c_1, \mu)$  to be sent by the sender in  $\Pi_{\text{spsVOLE-pcGGM}}$  can be compressed via the PRF technique for  $\Pi_{\text{spCOT}}$ . In  $t$  concurrent **Extend** executions, all such random messages can also be compressed in batch.
- The optimization for a large field  $\mathbb{F}$  in  $\Pi_{\text{spsVOLE-cGGM}}$  also applies.
- If  $\mathbb{F} = \mathbb{F}_2$ ,  $\Pi_{\text{spsVOLE-pcGGM}}$  degenerates to single-point COT and can do away with  $\mathcal{F}_{\text{sVOLE}}$  so that the receiver need not send a difference  $d \in \mathbb{F}$ . Instead, the sender locally samples  $\Gamma \in \mathbb{K}$  and masks this value with the sum of all last-level nodes in a pcGGM tree. This optimization has been used in [10].

**Complexity analysis.** Consider the complexity per execution when the PRF-based optimization is used in  $t$  concurrent **Extend** executions.  $\Pi_{\text{spsVOLE-pcGGM}}$  uses  $n$  precomputed COT tuples and one precomputed sVOLE tuple.  $P_0$  sends  $(n-2) \cdot \lambda + 3 \cdot \log |\mathbb{K}| + \frac{\lambda}{t}$  bits, and  $P_1$  sends  $\log |\mathbb{F}|$  bits. The computation is

**Protocol  $\Pi_{\text{spsVOLE-pcGGM}}$**

**Parameters:** Field  $\mathbb{F}$  and its extension field  $\mathbb{K}$ .

**Initialize:** This procedure is executed only once.

1.  $P_0$  and  $P_1$  send (init) to  $\mathcal{F}_{\text{COT}}$ , which returns  $\Delta \in \mathbb{F}_{2^\lambda}$  to  $P_0$ .
2.  $P_0$  and  $P_1$  send (init) to  $\mathcal{F}_{\text{sVOLE}}$ , which returns  $\Gamma \in \mathbb{K}$  to  $P_0$ .  $P_0$  outputs  $\Gamma$ .

**Extend:** This procedure can be executed many times.  $P_0$  and  $P_1$  input  $N = 2^n$  and use pcGGM (c.f. Figure 5) for  $n$ ,  $\mathbb{K}$ , keyed hash function  $H_S : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ , and function  $\text{Convert}_{\mathbb{K}} : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{K}$ .

3.  $P_0$  and  $P_1$  send (extend,  $n$ ) to  $\mathcal{F}_{\text{COT}}$ , which returns  $(K[r_1], \dots, K[r_n]) \in \mathbb{F}_{2^\lambda}^n$  to  $P_0$  and  $((r_1, \dots, r_n), (M[r_1], \dots, M[r_n])) \in \mathbb{F}_2^n \times \mathbb{F}_{2^\lambda}^n$  to  $P_1$  such that  $M[r_i] = K[r_i] \oplus r_i \cdot \Delta$  for  $i \in [1, n]$ .
4.  $P_0$  and  $P_1$  send (extend, 1) to  $\mathcal{F}_{\text{sVOLE}}$ , which returns  $K[s] \in \mathbb{K}$  to  $P_0$  and  $(s, M[s]) \in \mathbb{F} \times \mathbb{K}$  to  $P_1$  such that  $M[s] = K[s] + s \cdot \Gamma$ .
5.  $P_1$  samples  $\beta \leftarrow \mathbb{F}^*$ , sets  $M[\beta] := M[s]$ , and sends  $d := s - \beta \in \mathbb{F}$  to  $P_0$ .  $P_0$  sets  $K[\beta] := K[s] + d \cdot \Gamma$  such that  $M[\beta] = K[\beta] + \beta \cdot \Gamma$ .
6.  $P_0$  samples  $(c_1, \mu) \leftarrow \mathbb{F}_{2^\lambda}^2$  and sets  $k := K[r_1] \oplus c_1$ ,

$$(\mathbf{v}, \{K_i^0\}_{i \in [1, n-1]}, (K_n^0, K_n^1)) := \text{pcGGM.FullEval}(\Delta, k),$$

$$c_i := K[r_i] \oplus K_i^0 \text{ for } i \in [2, n-1], c_n^b := \text{Convert}_{\mathbb{K}}(H_S(\mu \oplus K[r_n] \oplus b \cdot \Delta)) + K_n^b$$

$$\text{for } b \in \{0, 1\}, \text{ and } \psi := K_n^0 + K_n^1 - K[\beta].$$

$$P_0 \text{ sends } (c_1, \dots, c_{n-1}, \mu, c_n^0, c_n^1, \psi) \text{ to } P_1.$$

7.  $P_1$  sets  $\alpha = \alpha_1 \dots \alpha_n := \bar{r}_1 \dots \bar{r}_n \in [0, N)$ ,  $K_i^{\bar{\alpha}_i} := M[r_i] \oplus c_i$  for  $i \in [1, n-1]$ ,  $K_n^{\bar{\alpha}_n} := c_n^{\bar{\alpha}_n} - \text{Convert}_{\mathbb{K}}(H_S(\mu \oplus M[r_n]))$ , and

$$\mathbf{u} := \text{unit}_{\mathbb{F}}(N, \alpha, \beta), \quad \mathbf{w} := \text{pcGGM.PuncFullEval}(\alpha, \{K_i^{\bar{\alpha}_i}\}_{i \in [1, n]}, \psi + M[\beta]).$$

8.  $P_0$  outputs  $\mathbf{v}$  and  $P_1$  outputs  $(\mathbf{u}, \mathbf{w})$ .

Fig. 6: pcGGM-based single-point sVOLE protocol in the  $(\mathcal{F}_{\text{COT}}, \mathcal{F}_{\text{sVOLE}})$ -hybrid model.

dominated by the tree expansion with  $1.5N$  RP calls for each party. Compared with the prior works [10, 43], our protocol roughly halve the communication, and the reduction in computation is 25%. This computation cost includes no PRG call in  $\text{Convert}_{\mathbb{K}}$ , which can be implemented from cheap modulo operations for the field size  $|\mathbb{K}|$  considered in many sVOLE applications, e.g., [43, 45, 40, 44].

## 5 DPF and DCF Correlation Generation

We model DPF/DCF correlation generation in functionality  $\mathcal{F}_{\text{FSS}}$  (Figure 7), which includes distributed key generation and local full-domain evaluation. By putting both procedures in the same functionality, we are able to model FSS as an ideal functionality and avoid caveats in the proof.  $\mathcal{F}_{\text{FSS}}$  focuses on  $N = 2^n$  for  $n \in \mathbb{N}$ , and we can define a similar functionality for a general  $N \in \mathbb{N}$ . Using padding, our protocols for  $\mathcal{F}_{\text{FSS}}$  also works in this general case.

**Functionality  $\mathcal{F}_{\text{FSS}}$**

**Parameters:** Ring  $\mathcal{R}$ .  $\text{FSS} \in \{\text{DPF}, \text{DCF}\}$  with domain  $[0, N)$ , where domain size  $N = 2^n$  for  $n \in \mathbb{N}$ , and range  $\mathcal{R}$ .

**Gen:** This functionality allows polynomially many (**gen**) commands. Upon receiving (**gen**,  $\langle \alpha \rangle_b, \langle \beta \rangle_b^A$ ) from  $P_b$  for each  $b \in \{0, 1\}$ , where  $(\langle \alpha \rangle_b, \langle \beta \rangle_b^A) \in [0, N) \times \mathcal{R}$ :

1. Set  $\alpha := \langle \alpha \rangle_0 \oplus \langle \alpha \rangle_1 \in [0, N)$ ,  $\beta := \langle \beta \rangle_0^A + \langle \beta \rangle_1^A \in \mathcal{R}$ , and  $\mathbf{r} \in \mathcal{R}^N$  such that
  - If  $\text{FSS} = \text{DPF}$ ,  $\mathbf{r}^{(j)} = 0$  for  $j \in [0, N)$ ,  $j \neq \alpha$ , and  $\mathbf{r}^{(\alpha)} = \beta$ .
  - If  $\text{FSS} = \text{DCF}$ ,  $\mathbf{r}^{(j)} = 0$  for  $j \in [0, N)$ ,  $j \geq \alpha$ , and  $\mathbf{r}^{(j)} = \beta$  otherwise.
2. If both parties are honest, sample  $\langle \mathbf{r} \rangle_0^A, \langle \mathbf{r} \rangle_1^A \leftarrow \mathcal{R}^N$  such that  $\langle \mathbf{r} \rangle_0^A + \langle \mathbf{r} \rangle_1^A = \mathbf{r}$ ; otherwise (i.e.,  $P_b$  is corrupted), receive  $\langle \mathbf{r} \rangle_b^A \in \mathcal{R}^N$  from the adversary and recompute  $\langle \mathbf{r} \rangle_{1-b}^A := \mathbf{r} - \langle \mathbf{r} \rangle_b^A \in \mathcal{R}^N$ .
3. Send  $\langle \mathbf{r} \rangle_0^A$  to  $P_0$  and  $\langle \mathbf{r} \rangle_1^A$  to  $P_1$ .

Fig. 7: Functionality for DPF/DCF correlation generation.

One can view  $\mathcal{F}_{\text{FSS}}$  as an alternative to the FSS key generation functionality that outputs each FSS key in the key pair to the designated party, who locally uses its key to evaluate its shares of the evaluation results at several points. We note that the full-domain evaluation included in  $\mathcal{F}_{\text{FSS}}$  does not complicate its implementation in contrast to the known protocols [22, 7] of the FSS key generation functionality. The reason is that, using the black-box evaluation technique [22], these protocols also perform full-domain evaluation. If FSS correlations are generated for immediate use without long-term storage (e.g., [22]),  $\mathcal{F}_{\text{FSS}}$  can be a drop-in replacement of the FSS key generation functionality. However, we also show in Appendix E.2 of the full version [29] that our protocols for  $\mathcal{F}_{\text{FSS}}$  can be adapted to realize this key generation functionality.

### 5.1 DPF and DCF Schemes

Note that DPF/DCF scheme may be used in not only distributed settings (e.g., [22]) but also the scenarios where a trusted dealer is available (e.g., two-server PIR [25, 13]). It would be better for us to present the two schemes alone.

We present in Figure 8 (resp., Figure 9) our DPF (resp., DCF) scheme, which is implicitly constructed from a *shared* pseudorandom correlated GGM tree. For simplicity of exposition, we slightly abuse the function  $\text{Convert}_{\mathbb{G}} : \{0, 1\}^* \rightarrow \mathbb{G}$  so that it can map random strings of either  $\lambda$  or  $\lambda - 1$  bits to pseudorandom group elements in  $\mathbb{G}$ . Our DCF scheme makes *non-black-box* use of our DPF scheme.

Note that our DPF and DCF schemes use a keyed hash function  $H_S$ . When there is a trusted dealer, the key  $S$  can be uniformly sampled by the dealer. In our DPF and DCF protocols in the upcoming sections, it can be jointly sampled by two parties using one-time public coin-tossing. This hash key can be reused across polynomially many FSS key pairs.

**Complexity analysis.** Consider the group  $\mathbb{G}$  (e.g., in [25, 13, 22, 14, 7]) with the PRG-free implementation of  $\text{Convert}_{\mathbb{G}}$  (c.f. Appendix F.1 of the full version [29]).

Our DPF scheme has a full-domain evaluation that takes  $1.5N$  RP calls, in contrast to the  $2N$  RP calls in the state-of-the-art construction of [13]. Its

**Parameters:** Domain size  $N = 2^n$  for  $n \in \mathbb{N}$ . Group  $\mathbb{G}$ . Keyed hash function  $H_S : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ . Function  $\text{Convert}_{\mathbb{G}} : \{0, 1\}^* \rightarrow \mathbb{G}$ .

**DPF.Gen**( $1^\lambda, (\alpha, \beta, n, \mathbb{G})$ ):

- 1: Parse  $\alpha = \alpha_1 \dots \alpha_n \in \{0, 1\}^n$  and  $\beta \in \mathbb{G}$ .
- 2: Sample  $\Delta \leftarrow \{0, 1\}^\lambda$  such that  $\text{lsb}(\Delta) = 1$ .
- 3: Sample  $\langle s_0 \parallel t_0 \rangle_0, \langle s_0 \parallel t_0 \rangle_1 \leftarrow \{0, 1\}^\lambda$  such that  $\langle s_0 \parallel t_0 \rangle_0 \oplus \langle s_0 \parallel t_0 \rangle_1 = \Delta$ .
- 4: **for**  $i \in [1, n-1]$  **do**
- 5:      $\text{CW}_i := H_S(\langle s_{i-1} \parallel t_{i-1} \rangle_0) \oplus H_S(\langle s_{i-1} \parallel t_{i-1} \rangle_1) \oplus \bar{\alpha}_i \cdot \Delta$
- 6:      $\langle s_i \parallel t_i \rangle_0 := H_S(\langle s_{i-1} \parallel t_{i-1} \rangle_0) \oplus \alpha_i \cdot \langle s_{i-1} \parallel t_{i-1} \rangle_0 \oplus \langle t_{i-1} \rangle_0 \cdot \text{CW}_i$
- 7:      $\langle s_i \parallel t_i \rangle_1 := H_S(\langle s_{i-1} \parallel t_{i-1} \rangle_1) \oplus \alpha_i \cdot \langle s_{i-1} \parallel t_{i-1} \rangle_1 \oplus \langle t_{i-1} \rangle_1 \cdot \text{CW}_i$
- 8:  $\langle \text{high}^\sigma \parallel \text{low}^\sigma \rangle_0 := H_S(\langle s_{n-1} \parallel t_{n-1} \rangle_0 \oplus \sigma)$  for  $\sigma \in \{0, 1\}$
- 9:  $\langle \text{high}^\sigma \parallel \text{low}^\sigma \rangle_1 := H_S(\langle s_{n-1} \parallel t_{n-1} \rangle_1 \oplus \sigma)$  for  $\sigma \in \{0, 1\}$
- 10:  $\text{HCW} := \langle \text{high}^{\alpha_n} \rangle_0 \oplus \langle \text{high}^{\alpha_n} \rangle_1$
- 11:  $\text{LCW}^0 := \langle \text{low}^0 \rangle_0 \oplus \langle \text{low}^0 \rangle_1 \oplus \bar{\alpha}_n$ ,      $\text{LCW}^1 := \langle \text{low}^1 \rangle_0 \oplus \langle \text{low}^1 \rangle_1 \oplus \alpha_n$
- 12:  $\text{CW}_n := (\text{HCW} \parallel \text{LCW}^0 \parallel \text{LCW}^1)$
- 13:  $\langle s_n \parallel t_n \rangle_0 := \langle \text{high}^{\alpha_n} \parallel \text{low}^{\alpha_n} \rangle_0 \oplus \langle t_{n-1} \rangle_0 \cdot (\text{HCW} \parallel \text{LCW}^{\alpha_n})$
- 14:  $\langle s_n \parallel t_n \rangle_1 := \langle \text{high}^{\alpha_n} \parallel \text{low}^{\alpha_n} \rangle_1 \oplus \langle t_{n-1} \rangle_1 \cdot (\text{HCW} \parallel \text{LCW}^{\alpha_n})$
- 15:  $\text{CW}_{n+1} := (\langle t_n \rangle_0 - \langle t_n \rangle_1) \cdot (\text{Convert}_{\mathbb{G}}(\langle s_n \rangle_1) - \text{Convert}_{\mathbb{G}}(\langle s_n \rangle_0) + \beta)$
- 16:  $k_b := (\langle s_0 \parallel t_0 \rangle_b, \{\text{CW}_i\}_{i \in [1, n+1]})$  for  $b \in \{0, 1\}$
- 17: **return**  $(k_0, k_1)$

**DPF.Eval**( $b, k_b, x$ ):

- 1: Parse  $k_b = (\langle s_0 \parallel t_0 \rangle_b, \{\text{CW}_i\}_{i \in [1, n+1]})$ ,  $\text{CW}_n = (\text{HCW} \parallel \text{LCW}^0 \parallel \text{LCW}^1)$ , and  $x = x_1 \dots x_n \in \{0, 1\}^n$ .
- 2: **for**  $i \in [1, n-1]$  **do**
- 3:      $\langle s_i^{x_1 \dots x_i} \parallel t_i^{x_1 \dots x_i} \rangle_b := H_S(\langle s_{i-1}^{x_1 \dots x_{i-1}} \parallel t_{i-1}^{x_1 \dots x_{i-1}} \rangle_b \oplus x_i \cdot \langle s_{i-1}^{x_1 \dots x_{i-1}} \parallel t_{i-1}^{x_1 \dots x_{i-1}} \rangle_b \oplus \langle t_{i-1}^{x_1 \dots x_{i-1}} \rangle_b \cdot \text{CW}_i)$
- 4:  $\langle \text{high} \parallel \text{low} \rangle_b := H_S(\langle s_{n-1}^{x_1 \dots x_{n-1}} \parallel t_{n-1}^{x_1 \dots x_{n-1}} \rangle_b \oplus x_n)$
- 5:  $\langle s_n^x \parallel t_n^x \rangle_b := \langle \text{high} \parallel \text{low} \rangle_b \oplus \langle t_{n-1}^{x_1 \dots x_{n-1}} \rangle_b \cdot (\text{HCW} \parallel \text{LCW}^{x_n})$
- 6: **return**  $y_b := (-1)^b \cdot (\text{Convert}_{\mathbb{G}}(\langle s_n^x \rangle_b) + \langle t_n^x \rangle_b \cdot \text{CW}_{n+1})$

Fig. 8: Our DPF scheme with domain  $[0, N)$  and range  $\mathbb{G}$ .

key generation algorithm uses about  $2n + 2$  RP calls while this figure is about  $4n$  in the prior work. In our scheme, the key size is  $n \cdot \lambda + (\lambda + 1) + \log |\mathbb{G}|$  bits, and the evaluation algorithm takes about  $n$  RP calls, both remaining the same complexity as those in the prior work. In our DCF scheme, the full-domain evaluation requires  $2.5N$  RP calls, in contrast to  $4N$  RP calls in the state-of-the-art construction [7]. Its key generation needs about  $4n + 2$  RP calls, in contrast to  $8n$  RP calls in the prior work. The key size is  $n \cdot \lambda + (\lambda + 1) + (n + 1) \cdot \log |\mathbb{G}|$  bits, and the evaluation requires about  $2n$  RP calls, without any improvement.

**Security.** We prove the following theorems in Appendix D.2 and Appendix D.3 of the full version [29]. These theorems turn to the intuition that  $\text{CW}_1, \dots, \text{CW}_n$  are masked by pseudorandom CCR outputs (as the root and the first  $n - 1$  on-path shared nodes are  $\Delta$ ), and  $\text{CW}_{n+1}, \text{VCW}_1, \dots, \text{VCW}_n$  are masked by some pseudorandom  $\text{Convert}_{\mathbb{G}}$  terms taking (pseudo)random CCR outputs as input.

**Parameters:** Domain size  $N = 2^n$  for  $n \in \mathbb{N}$ . Group  $\mathbb{G}$ . Keyed hash function  $H_S : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ . Function  $\text{Convert}_{\mathbb{G}} : \{0, 1\}^* \rightarrow \mathbb{G}$ .

**DCF.Gen**( $1^\lambda, (\alpha, \beta, n, \mathbb{G})$ ):

- 1: Parse  $\alpha = \alpha_1 \dots \alpha_n \in \{0, 1\}^n$  and  $\beta \in \mathbb{G}$ . Let  $\alpha_0 := 0$ .
- 2: Run  $(k'_0, k'_1) \leftarrow \text{DPF.Gen}(1^\lambda, (\alpha, -\alpha_n \cdot \beta, n, \mathbb{G}))$  and store its internal variables.
- 3: **for**  $i \in [1, n]$  **do**
- 4:    $\langle v_i \rangle_0 := H_S(\langle s_{i-1} \parallel t_{i-1} \rangle_0 \oplus 2)$
- 5:    $\langle v_i \rangle_1 := H_S(\langle s_{i-1} \parallel t_{i-1} \rangle_1 \oplus 2)$
- 6:    $\text{VCW}_i := (\langle t_{i-1} \rangle_0 - \langle t_{i-1} \rangle_1)$   
            $\cdot (\text{Convert}_{\mathbb{G}}(\langle v_i \rangle_1) - \text{Convert}_{\mathbb{G}}(\langle v_i \rangle_0) + (\alpha_i - \alpha_{i-1}) \cdot \beta)$
- 7:  $k_b := (k'_b, \{\text{VCW}_i\}_{i \in [1, n]})$  for  $b \in \{0, 1\}$
- 8: **return**  $(k_0, k_1)$

**DCF.Eval**( $b, k_b, x$ ):

- 1: Parse  $k_b = (k'_b, \{\text{VCW}_i\}_{i \in [1, n]})$ . Let  $V_b^0 := 0 \in \mathbb{G}$ .
- 2: Run  $y'_b := \text{DPF.Eval}(b, k'_b, x)$  and store its internal variables.
- 3: **for**  $i \in [1, n]$  **do**
- 4:    $\langle v_i^{x_1 \dots x_{i-1}} \rangle_b := H_S(\langle s_{i-1}^{x_1 \dots x_{i-1}} \parallel t_{i-1}^{x_1 \dots x_{i-1}} \rangle_b \oplus 2)$
- 5:    $V_b^i := V_b^{i-1} + (-1)^b \cdot (\text{Convert}_{\mathbb{G}}(\langle v_i^{x_1 \dots x_{i-1}} \rangle_b) + \langle t_{i-1}^{x_1 \dots x_{i-1}} \rangle_b \cdot \text{VCW}_i)$
- 6: **return**  $y_b := y'_b + V_b^n$

Fig. 9: Our DCF scheme with domain  $[0, N)$  and range  $\mathbb{G}$ .

**Theorem 3.** Given CCR function  $H : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ , function  $\text{Convert}_{\mathbb{G}} : \mathbb{F}_{2^{\lambda-1}} \rightarrow \mathbb{G}$ , and keyed hash function  $H_S(x) := H(S \oplus x)$  with some key  $S \leftarrow \mathbb{F}_{2^\lambda}$ , Figure 8 gives a DPF scheme with domain  $[0, N)$  and range  $\mathbb{G}$ .

**Theorem 4.** Given CCR function  $H : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ , function  $\text{Convert}_{\mathbb{G}} : \mathbb{F}_{2^\ell} \rightarrow \mathbb{G}$  with  $\ell \in \{\lambda - 1, \lambda\}$ , and keyed hash function  $H_S(x) := H(S \oplus x)$  with some key  $S \leftarrow \mathbb{F}_{2^\lambda}$ , Figure 9 gives a DCF scheme with domain  $[0, N)$  and range  $\mathbb{G}$ .

## 5.2 DPF Correlation Generation

We define a leveled evaluation algorithm  $\text{DPF.NextLevel}$  such that, on input a level index  $i \in [1, n]$ , all nodes on the  $(i - 1)$ -th level of the share of a shared pseudorandom correlated GGM tree, and the public correction word  $\text{CW}_i$  for the  $i$ -th level, outputs all nodes on the  $i$ -th level.

In Figure 10, we present our DPF correlation generation protocol  $\Pi_{\text{DPF}}$ . This protocol operates in the  $(\mathcal{F}_{\text{COT}}, \mathcal{F}_{\text{Rand}}, \mathcal{F}_{\text{OLE}})$ -hybrid model.  $\mathcal{F}_{\text{Rand}}$  is the standard coin-tossing functionality that outputs a uniform string to both parties.  $\mathcal{F}_{\text{OLE}}$  is the functionality for oblivious linear evaluation (OLE) on ring  $\mathcal{R}$ , where  $P_0$  (resp.,  $P_1$ ) is given *random*  $(\mathbf{x}_0, \mathbf{z}_0) \in \mathcal{R}^N \times \mathcal{R}^N$  (resp.,  $(\mathbf{x}_1, \mathbf{z}_1) \in \mathcal{R}^N \times \mathcal{R}^N$ ) such that  $\mathbf{z}_0 + \mathbf{z}_1$  equals the component-wise multiplication  $\mathbf{x}_0 \odot \mathbf{x}_1$ . We refer readers to Appendix F.2 and Appendix F.3 of the full version [29] for the definitions and instantiations of  $\mathcal{F}_{\text{Rand}}$  and  $\mathcal{F}_{\text{OLE}}$ . If  $\beta$  is a bit-string,  $\Pi_{\text{DPF}}$  never uses  $\mathcal{F}_{\text{OLE}}$ .

$\Pi_{\text{DPF}}$  requires  $\mathcal{F}_{\text{Rand}}$  for the following reason. Note that  $\Pi_{\text{DPF}}$  uses the same global offset  $\Delta$  as the roots of polynomially many shared trees, each of which defines a fresh DPF correlation. So, the two shares of this identical root should



**Protocol  $\Pi_{\text{DPF}}$**

**Parameters:** Domain size  $N = 2^n$  for  $n \in \mathbb{N}$ . Ring  $\mathcal{R}$ . Keyed hash function  $H_S : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ . Function  $\text{Convert}_{\mathcal{R}} : \{0, 1\}^* \rightarrow \mathcal{R}$ . Let  $H' := \text{hb} \circ H_S$ .

**DPF Gen:** This procedure can be executed many times. For each  $b \in \{0, 1\}$ ,  $P_b$  inputs  $(\langle \alpha \rangle_b, \langle \beta \rangle_b^A) \in [0, N) \times \mathcal{R}$  and proceeds as follows:

1. The two parties run sub-protocol  $\Pi_{\text{PREP}}$  (Figure 11), which, for each  $b \in \{0, 1\}$ , returns  $\langle \Delta \rangle_b$  and  $\{(\text{K}_b[\langle \alpha_i \rangle_{1-b}], \text{M}_b[\langle \alpha_i \rangle_b])\}_{i \in [1, n]}$  to  $P_b$  such that  $\text{lsb}(\langle \Delta \rangle_0 \oplus \langle \Delta \rangle_1) = 1$ , and  $\text{M}_b[\langle \alpha_i \rangle_b] = \text{K}_{1-b}[\langle \alpha_i \rangle_b] \oplus \langle \alpha_i \rangle_b \cdot \langle \Delta \rangle_{1-b}$  for  $i \in [1, n]$ .
2. The two parties send **(sample,  $\lambda$ )** to  $\mathcal{F}_{\text{Rand}}$ , which returns  $W \in \{0, 1\}^\lambda$  to them.
3.  $P_b$  computes  $\langle s_0^0 \parallel t_0^0 \rangle_b := \langle \Delta \rangle_b \oplus W$ . For  $i \in [1, n-1]$ ,  $P_b$  sends to  $P_{1-b}$

$$\begin{aligned} \langle \text{CW}_i \rangle_b &:= (\oplus_{j \in [0, 2^{i-1}]} H_S(\langle s_{i-1}^j \parallel t_{i-1}^j \rangle_b)) \\ &\quad \oplus \overline{\langle \alpha_i \rangle_b} \cdot \langle \Delta \rangle_b \oplus \text{K}_b[\langle \alpha_i \rangle_{1-b}] \oplus \text{M}_b[\langle \alpha_i \rangle_b], \end{aligned}$$

receives  $\langle \text{CW}_i \rangle_{1-b}$  from  $P_{1-b}$ , and computes  $\text{CW}_i := \langle \text{CW}_i \rangle_b \oplus \langle \text{CW}_i \rangle_{1-b}$  and

$$\{\langle s_i^j \parallel t_i^j \rangle_b\}_{j \in [0, 2^i)} := \text{DPF.NextLevel}(i, \{\langle s_{i-1}^j \parallel t_{i-1}^j \rangle_b\}_{j \in [0, 2^{i-1})}, \text{CW}_i).$$

4.  $P_b$  samples  $\mu_b \leftarrow \{0, 1\}^\lambda$ , computes

$$\langle \text{Xhigh}^\sigma \parallel \text{Xlow}^\sigma \rangle_b := \oplus_{j \in [0, 2^{n-1}]} H_S(\langle s_{n-1}^j \parallel t_{n-1}^j \rangle_b \oplus \sigma) \quad \text{for } \sigma \in \{0, 1\},$$

$d_b := H'(\mu_b \oplus \text{K}_b[\langle \alpha_n \rangle_{1-b}]) \oplus H'(\mu_b \oplus \text{K}_b[\langle \alpha_n \rangle_{1-b}] \oplus \langle \Delta \rangle_b) \oplus \langle \text{Xhigh}^0 \oplus \text{Xhigh}^1 \rangle_b$ , sends  $(\mu_b, d_b)$  to  $P_{1-b}$ , and receives  $(\mu_{1-b}, d_{1-b})$  from  $P_{1-b}$ . Then,  $P_b$  computes

$$\begin{aligned} \langle \text{HCW} \rangle_b &:= \langle \text{Xhigh}^{\overline{\langle \alpha_n \rangle_b}} \rangle_b \oplus H'(\mu_b \oplus \text{K}_b[\langle \alpha_n \rangle_{1-b}]) \\ &\quad \oplus H'(\mu_{1-b} \oplus \text{M}_b[\langle \alpha_n \rangle_b]) \oplus \langle \alpha_n \rangle_b \cdot d_{1-b}, \end{aligned}$$

$$\langle \text{LCW}^0 \rangle_b := \langle \text{Xlow}^0 \rangle_b \oplus \langle \alpha_n \rangle_b \oplus b, \quad \langle \text{LCW}^1 \rangle_b := \langle \text{Xlow}^1 \rangle_b \oplus \langle \alpha_n \rangle_b,$$

sends  $\langle \text{CW}_n \rangle_b := (\langle \text{HCW} \rangle_b \parallel \langle \text{LCW}^0 \rangle_b \parallel \langle \text{LCW}^1 \rangle_b)$  to  $P_{1-b}$ , receives  $\langle \text{CW}_n \rangle_{1-b}$  from  $P_{1-b}$ , and computes  $\text{CW}_n := \langle \text{CW}_n \rangle_b \oplus \langle \text{CW}_n \rangle_{1-b}$  and

$$\{\langle s_n^j \parallel t_n^j \rangle_b\}_{j \in [0, N)} := \text{DPF.NextLevel}(n, \{\langle s_{n-1}^j \parallel t_{n-1}^j \rangle_b\}_{j \in [0, 2^{n-1})}, \text{CW}_n).$$

5. **(Binary field  $\mathcal{R} = \mathbb{F}_{2^\ell}$ , without  $\mathcal{F}_{\text{OLE}}$ )**

$P_b$  computes  $\langle \text{CW}_{n+1} \rangle_b^A := (\sum_{j \in [0, N)} \text{Convert}_{\mathcal{R}}(\langle s_n^j \rangle_b)) + \langle \beta \rangle_b^A$ .

**(General ring  $\mathcal{R}$ , using  $\mathcal{F}_{\text{OLE}}$ )**

The two parties run sub-protocol  $\Pi_{\text{MULT}}$  (Figure 12), which, for each  $b \in \{0, 1\}$ , takes as input

$$\begin{aligned} \langle A \rangle_b^A &:= (-1)^b \cdot \sum_{j \in [0, N)} \langle t_n^j \rangle_b \in \mathcal{R}, \\ \langle B \rangle_b^A &:= (-1)^{1-b} \cdot \sum_{j \in [0, N)} \text{Convert}_{\mathcal{R}}(\langle s_n^j \rangle_b) + \langle \beta \rangle_b^A \in \mathcal{R}, \end{aligned}$$

and returns  $\langle \text{CW}_{n+1} \rangle_b^A$  to  $P_b$ .

In either case,  $P_b$  sends  $\langle \text{CW}_{n+1} \rangle_b^A$  to  $P_{1-b}$ , receives  $\langle \text{CW}_{n+1} \rangle_{1-b}^A$  from  $P_{1-b}$ , and computes  $\text{CW}_{n+1} := \langle \text{CW}_{n+1} \rangle_b^A + \langle \text{CW}_{n+1} \rangle_{1-b}^A$ .

6.  $P_b$  computes  $k_b := (\langle \Delta \rangle_b \oplus W, \{\text{CW}_i\}_{i \in [1, n+1]})$  and  $\langle \mathbf{r}^{(j)} \rangle_b^A := \text{DPF.Eval}(b, k_b, j)$  for  $j \in [0, N)$ , and outputs  $\langle \mathbf{r} \rangle_b^A \in \mathcal{R}^N$ .

Fig. 10: DPF correlation generation in the  $(\mathcal{F}_{\text{COT}}, \mathcal{F}_{\text{Rand}}, \mathcal{F}_{\text{OLE}})$ -hybrid model.

**Protocol  $\Pi_{\text{PREP}}$**

**Initialize:** This procedure is executed only once for each  $b \in \{0, 1\}$ . The two parties send (init) to  $\mathcal{F}_{\text{COT}}^b$  with identifier  $b$ , which returns  $\Delta'_b \in \{0, 1\}^\lambda$  to  $P_b$ .  $P_b$  sends  $\text{lsb}(\Delta'_b)$  to  $P_{1-b}$ , receives  $\text{lsb}(\Delta'_{1-b})$  from  $P_{1-b}$ , and sets  $\langle \Delta \rangle_b := \Delta'_b \oplus (0^{\lambda-1} \parallel (\text{lsb}(\Delta'_{1-b}) \oplus b))$  such that  $\text{lsb}(\langle \Delta \rangle_0 \oplus \langle \Delta \rangle_1) = 1$ .

For each  $b \in \{0, 1\}$ :  $P_b$  inputs  $\langle \alpha \rangle_b \in \{0, 1\}^n$  and proceeds as follows.

- 1-1. The two parties send (extend,  $n$ ) to  $\mathcal{F}_{\text{COT}}^b$  with identifier  $b$ , which returns  $\mathbf{k}_b \in \mathbb{F}_{2^\lambda}^n$  to  $P_b$  and  $(\mathbf{r}_{1-b}, \mathbf{m}_{1-b}) \in \mathbb{F}_2^2 \times \mathbb{F}_{2^\lambda}^n$  to  $P_{1-b}$  such that  $\mathbf{m}_{1-b} = \mathbf{k}_b \oplus \mathbf{r}_{1-b} \cdot \Delta'_b$ .
- 1-2.  $P_b$  sets  $\mathbf{g}_b := \langle \alpha \rangle_b \oplus \mathbf{r}_b$ , sends  $\mathbf{g}_b$  to  $P_{1-b}$ , and receives  $\mathbf{g}_{1-b}$  from  $P_{1-b}$ . For  $i \in [1, n]$ ,  $P_b$  sets

$$\mathbf{K}_b[\langle \alpha_i \rangle_{1-b}] := \mathbf{k}_b^{(i)} \oplus \mathbf{g}_{1-b}^{(i)} \cdot \langle \Delta \rangle_b,$$

$$\mathbf{M}_b[\langle \alpha_i \rangle_b] := \mathbf{m}_b^{(i)} \oplus \mathbf{r}_b^{(i)} \cdot (0^{\lambda-1} \parallel (\text{lsb}(\Delta'_b) \oplus (1-b))).$$

- 1-3.  $P_b$  outputs  $\langle \Delta \rangle_b$  and  $\{(\mathbf{K}_b[\langle \alpha_i \rangle_{1-b}], \mathbf{M}_b[\langle \alpha_i \rangle_b])\}_{i \in [1, n]}$ .

Fig. 11: Preprocessing sub-protocol for DPF/DCF correlation generation.

**Protocol  $\Pi_{\text{MULT}}$**

For each  $b \in \{0, 1\}$ :  $P_b$  inputs  $(\langle A \rangle_b^A, \langle B \rangle_b^A) \in \mathcal{R}^2$  and proceeds as follows.

1. The two parties send (extend, 2) to  $\mathcal{F}_{\text{OLE}}$ , which, for each  $b \in \{0, 1\}$ , returns  $(\mathbf{x}_b, \mathbf{z}_b) \in \mathcal{R}^2 \times \mathcal{R}^2$  to  $P_b$  such that  $\mathbf{z}_0 + \mathbf{z}_1 = \mathbf{x}_0 \cdot \mathbf{x}_1$ .
2.  $P_b$  computes  $(\gamma_b, \zeta_b) := (\langle A \rangle_b^A, \langle B \rangle_b^A) + (\mathbf{x}_b^{(b)}, \mathbf{x}_b^{(1-b)})$ , sends  $(\gamma_b, \zeta_b)$  to  $P_{1-b}$ , and receives  $(\gamma_{1-b}, \zeta_{1-b})$  from  $P_{1-b}$ .
3.  $P_b$  outputs  $\langle A \cdot B \rangle_b^A := \langle A \rangle_b^A \cdot \langle B \rangle_b^A + \langle A \rangle_b^A \cdot \zeta_{1-b} - \mathbf{x}_b^{(1-b)} \cdot \gamma_{1-b} + \mathbf{z}_b^{(0)} + \mathbf{z}_b^{(1)}$ .

Fig. 12: OLE-based multiplication sub-protocol.

be “re-randomized” to avoid the identical per-party shares of the defined correlations. The two parties do this re-randomization by calling  $\mathcal{F}_{\text{Rand}}$  for a public randomness  $W$  and XORing this value to their shares of  $\Delta$ , respectively.

In  $\Pi_{\text{DPF}}$ , the key  $S$  of the keyed hash function  $\mathbf{H}_S$  can be produced by one  $\mathcal{F}_{\text{Rand}}$  invocation before protocol execution, and we omit this setup for simplicity.

**Security.** We prove Theorem 5 in Appendix D.4 of the full version [29]. This proof will consider polynomially many concurrent **Gen** executions that uses the one-time initialized  $\Delta$ . Intuitively, the security primarily follows from the COT-based secure computation of correction words, where the COT tuples are related to the global offset  $\Delta$  so that the transcripts are masked by CCR responses. In particular, the intermediate transcript  $d_b$  is masked by a CCR response coming from a legal CCR query with overwhelming probability due to the uniform  $\mu_b$ .

**Theorem 5.** *Given CCR function  $\mathbf{H} : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ , function  $\text{Convert}_{\mathcal{R}} : \mathbb{F}_{2^{\lambda-1}} \rightarrow \mathcal{R}$ , and keyed hash function  $\mathbf{H}_S(x) := \mathbf{H}(S \oplus x)$  with some key  $S \leftarrow \mathbb{F}_{2^\lambda}$ , protocol  $\Pi_{\text{DPF}}$  (Figure 10) UC-realizes functionality  $\mathcal{F}_{\text{DPF}}$  (Figure 7) against any semi-*

		$n = 20$	$n = 22$	$n = 24$	$n = 26$	$n = 28$
$\mathcal{R} = \mathbb{F}_{2^{127}}$	LAN	50	120	397	1501	5920
	WAN	2752	3020	3492	4786	9355
$\mathcal{R} = \mathbb{F}_2$	LAN	29	30	34	52	120
	WAN	2930	3132	3337	3554	3823

Table 3: **The efficiency of distributed correlation generation for our DPF scheme.** All numbers are in milliseconds (*ms*).

*honest adversary in the  $(\mathcal{F}_{\text{COT}}, \mathcal{F}_{\text{Rand}}, \mathcal{F}_{\text{OLE}})$ -hybrid model. If  $\mathcal{R} = \mathbb{F}_{2^\ell}$  for  $\ell \in \mathbb{N}$ , protocol  $\Pi_{\text{DPF}}$  never invokes  $\mathcal{F}_{\text{OLE}}$ .*

**Communication optimization.**  $\Pi_{\text{DPF}}$  has the following two optimizations:

- For  $t$  concurrent **Gen** executions (e.g., in its applications to RAM-based computation [22], FSS-based MPC [7], and OLE extension [12], etc), each  $P_b$  can compress all  $\mu_b$ ’s in these executions via a PRF  $F : \mathbb{F}_{2^\lambda} \times \{0, 1\}^* \rightarrow \mathbb{F}_{2^\lambda}$  with a fresh key  $k_{\text{prf},b} \leftarrow \mathbb{F}_{2^\lambda}$  sampled *after receiving its COT outputs (from both  $\mathcal{F}_{\text{COT}}^b$  and  $\mathcal{F}_{\text{COT}}^{1-b}$ ) in all executions*. For each execution with sub-session ID  $\text{ssid}$ , the two parties define  $\mu_b := F(k_{\text{prf},b}, \text{ssid})$ .
- All invocations of  $\mathcal{F}_{\text{Rand}}$  can be compressed via another independent PRF key sampled *after the one-time initialization of  $\mathcal{F}_{\text{COT}}^b$  and  $\mathcal{F}_{\text{COT}}^{1-b}$*  so that the root of each  $P_b$ ’s tree is (pseudo)random.
- Another method to save the communication for random  $\mu_b$ ’s is to replace  $H_S$  by a hash function that meets “CCR for naturally derived keys” [47,28], which can also be implemented in one RP call. Note that  $\mu_b$  is introduced to prevent the replay attack, which results from the manipulation of COT outputs, against the hashing mask in  $d_b$ . The alternative hash function addresses this attack by adding non-repeating tweaks.

**Complexity analysis (binary field).** Consider the complexity per execution when the first PRF-based optimization is used in  $t$  concurrent **Gen** executions. The cost is symmetric.  $\Pi_{\text{DPF}}$  uses  $n$  COT tuples per party and one  $\mathcal{F}_{\text{Rand}}$  call. Each party sends  $(n+1) + (n+1) \cdot \lambda + \frac{\lambda}{t} + \log |\mathcal{R}|$  bits. The computation per party is dominated by the tree expansion in  $n$  DPF.NextLevel calls, or  $1.5N$  RP calls.  $\Pi_{\text{DPF}}$  runs in  $n+3$  rounds (without counting the one-time setup).

In contrast, the binary-field protocol [22] can be implemented from GMW-style 2PC and  $n$  string OTs each with  $(\lambda-1)$ -bit payloads. One can cast these string OTs into  $n$  precomputed COT tuples according to [34,3]. Using these tuples, each party sends  $n + n \cdot (3\lambda-1) + \log |\mathcal{R}|$  bits, and the computation per party is dominated by the  $2N$  RP calls in GGM tree expansion. This protocol can proceed in  $2n+2$  rounds: one for sending  $n$  masked choice bits, two for sharing and revealing each of the first  $n$  correction words, and one for revealing the  $(n+1)$ -th correction word. Our savings in computation, communication, and round complexity are about 25%, 66.6%, and 50%, respectively.

We implement  $\Pi_{\text{PREP}}$  and  $\Pi_{\text{DPF}}$  in C++, and perform benchmarks on a pair of Amazon EC2 R5.xlarge instances. We take binary fields  $\mathcal{R} = \mathbb{F}_{2^{127}}$  and  $\mathcal{R} = \mathbb{F}_2$  under computational security parameter  $\lambda \approx 128$ . The reported time include

both distributed key generation and full-domain evaluation. We set 1Gbps bandwidth with no latency as our LAN setting, and 20Mbps bandwidth with 100ms latency as our WAN setting. The results are shown in Table 3. We can see that our protocol is practically efficient, especially for two-server PIR. Although all numbers are reported based on one thread, performing one correlation generation for  $2^{28}$  127-bit values takes about 6 seconds, which is about 30% to 40% faster than the performance from a prior implementation in the same threads [22].

**Complexity analysis (general ring).** The two parties additionally need two precomputed OLE tuples for the secure multiplication. Overall, each party sends  $(n+1) + (n+1) \cdot \lambda + \frac{\lambda}{t} + 3 \cdot \log |\mathcal{R}|$  bits, and the protocol runs in  $n+4$  rounds.

In contrast, the binary-field protocol [22] can be adapted for the general-ring  $\text{CW}_{n+1}$  in the DPF scheme [13]. Securely computing this  $\text{CW}_{n+1}$  consumes two OLE tuples and needs the level-by-level 2PC, which leads to two additional bits in each OT payload per level, to share the last-level control bit  $\langle t_n \rangle_1$ . Each party sends at most  $n + n \cdot (3\lambda + 3) + 3 \cdot \log |\mathcal{R}|$  bits, and the protocol runs in  $2n+3$  rounds. The improvement is the same as the binary-field case.

### 5.3 DCF Correlation Generation

Our DCF protocol  $\Pi_{\text{DCF}}$  in Figure 13 extends  $\Pi_{\text{DPF}}$  by also computing  $n$  value correction words and defining the evaluation result as per our DCF scheme. If  $\beta$  is a bit-string, the two parties can compute  $n$  value correction words without using precomputed OLE tuples. Otherwise, for a general ring element  $\beta$ , these correction words are obtained from OLE-based secure multiplication.

**Security.** We prove Theorem 6 in Appendix D.5 of the full version [29], where polynomially many concurrent **Gen** executions are considered. The security is also based on the COT- and OLE-based secure computation of the  $n$  additional correction words of our DCF scheme. Note that the intermediate  $y_b^i$ 's are pseudorandom due the masking CCR responses, which are for the legal CCR queries with overwhelming probability in the presence of uniform  $x_b^i$ 's.

**Theorem 6.** *Given CCR function  $H : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ , function  $\text{Convert}_{\mathcal{R}} : \mathbb{F}_{2^\ell} \rightarrow \mathcal{R}$  for  $\ell \in \{\lambda-1, \lambda\}$ , and keyed hash function  $H_S(x) := H(S \oplus x)$  with some key  $S \leftarrow \mathbb{F}_{2^\lambda}$ , protocol  $\Pi_{\text{DCF}}$  (Figure 13) UC-realizes functionality  $\mathcal{F}_{\text{DCF}}$  (Figure 7) against any semi-honest adversary in the  $(\mathcal{F}_{\text{COT}}, \mathcal{F}_{\text{Rand}}, \mathcal{F}_{\text{OLE}})$ -hybrid model. If  $\mathcal{R} = \mathbb{F}_{2^\ell}$  for  $\ell \in \mathbb{N}$ , protocol  $\Pi_{\text{DCF}}$  never invokes  $\mathcal{F}_{\text{OLE}}$ .*

**Communication optimization.** The optimizations in Section 5.2 also applies to the DCF protocol  $\Pi_{\text{DCF}}$ . Moreover, the random elements  $\{x_b^i\}_{i \in [1, n]}$  in  $\Pi_{\text{DCF}}$  can also be compressed using the same technique for the random  $\mu_b$ 's.

**Complexity analysis (binary field).** Consider the complexity per execution when the first PRF-based optimization is used in  $t$  concurrent **Gen** executions. The cost is symmetric.  $\Pi_{\text{DCF}}$  consumes  $n$  COT tuples per party and one  $\mathcal{F}_{\text{Rand}}$  call. Each party sends  $(n+1) + (n+1) \cdot \lambda + \frac{\lambda}{t} + (2n+1) \cdot \log |\mathcal{R}|$  bits, and the computation per party comes from the  $2.5N$  RP calls in the tree expansion.  $\Pi_{\text{DCF}}$  has round complexity  $n+3$ , the same as  $\Pi_{\text{DPF}}$  in the binary-field case.

In contrast, the state-of-the-art protocol of [7] requires  $n$  string OTs to run GMW-style 2PC. The string OTs consume  $n$  precomputed COT tuples and

**Protocol  $\Pi_{\text{DCF}}$**

**Parameters:** Domain size  $N = 2^n$  for  $n \in \mathbb{N}$ . Ring  $\mathcal{R}$ . Keyed hash function  $H_S : \mathbb{F}_{2^\lambda} \rightarrow \mathbb{F}_{2^\lambda}$ . Function  $\text{Convert}_{\mathcal{R}} : \{0, 1\}^* \rightarrow \mathcal{R}$ . Let  $H^* := \text{Convert}_{\mathcal{R}} \circ H_S$ .

**DCF Gen:** This procedure can be executed many times. For each  $b \in \{0, 1\}$ ,  $P_b$  inputs  $(\langle \alpha \rangle_b, \langle \beta \rangle_b^A) \in [0, N) \times \mathcal{R}$  and proceeds as in  $\Pi_{\text{DPF}}$  (Figure 8), with the same Step 1, 2 and the following modifications to the subsequent steps:

3. Along with  $\langle CW_i \rangle_b$  for  $i \in [1, n-1]$ ,  $P_b$  samples  $x_b^i \leftarrow \{0, 1\}^\lambda$ , computes  $y_b^i := H^*(x_b^i \oplus K_b[\langle \alpha_i \rangle_{1-b}]) - H^*(x_b^i \oplus K_b[\langle \alpha_i \rangle_{1-b}] \oplus \langle \Delta \rangle_b) + \langle \beta \rangle_b^A - 2 \cdot \langle \alpha_i \rangle_b \cdot \langle \beta \rangle_b^A$ , sends  $(x_b^i, y_b^i)$  to  $P_{1-b}$ , receive  $(x_{1-b}^i, y_{1-b}^i)$  from  $P_{1-b}$ , and computes  $\langle \alpha_i \cdot \beta \rangle_b^A := \langle \alpha_i \rangle_b \cdot \langle \beta \rangle_b^A - H^*(x_b^i \oplus K_b[\langle \alpha_i \rangle_{1-b}]) + H^*(x_{1-b}^i \oplus M_b[\langle \alpha_i \rangle_b]) + \langle \alpha_i \rangle_b \cdot y_{1-b}^i$ .
4. Along with  $\langle CW_n \rangle_b$ ,  $P_b$  repeats Step 3 for  $i = n$  and computes  $\langle \alpha_n \cdot \beta \rangle_b^A$ .
5. For  $i \in [1, n]$  and  $j \in [0, 2^{i-1})$ ,  $P_b$  computes  $\langle v_i^j \rangle_b := H_S(\langle s_{i-1}^j \parallel t_{i-1}^j \rangle_b \oplus 2)$  and  $\langle \alpha_0 \cdot \beta \rangle_b^A := 0$ .  $P_b$  computes  $\langle CW_{n+1} \rangle_b^A$  by using  $\langle \alpha_n \cdot \beta \rangle_b^A$  instead of  $\langle \beta \rangle_b^A$ , and:
 

**(Binary field  $\mathcal{R} = \mathbb{F}_{2^\ell}$ , without  $\mathcal{F}_{\text{OLE}}$ )** For  $i \in [1, n]$  in parallel:  
 $P_b$  computes  $\langle VCW_i \rangle_b^A := (\sum_{j \in [0, 2^{i-1})} \text{Convert}_{\mathcal{R}}(\langle v_i^j \rangle_b)) + \langle \alpha_i \cdot \beta \rangle_b^A - \langle \alpha_{i-1} \cdot \beta \rangle_b^A$ .

**(General ring  $\mathcal{R}$ , using  $\mathcal{F}_{\text{OLE}}$ )** For  $i \in [1, n]$  in parallel:  
 The two parties run sub-protocol  $\Pi_{\text{MULT}}$  (Figure 12), which, for each  $b \in \{0, 1\}$ , takes as input

$$\langle A_i \rangle_b^A := (-1)^b \cdot \sum_{j \in [0, 2^{i-1})} \langle t_{i-1}^j \rangle_b \in \mathcal{R},$$

$$\langle B_i \rangle_b^A := (-1)^{1-b} \cdot \sum_{j \in [0, 2^{i-1})} \text{Convert}_{\mathcal{R}}(\langle v_i^j \rangle_b) + \langle \alpha_i \cdot \beta \rangle_b^A - \langle \alpha_{i-1} \cdot \beta \rangle_b^A \in \mathcal{R},$$
 and returns  $\langle VCW_i \rangle_b^A$  to  $P_b$ .

In either case, along with  $\langle CW_{n+1} \rangle_b^A$ ,  $P_b$  sends  $\langle VCW_i \rangle_b^A$  to  $P_{1-b}$ , receives  $\langle VCW_i \rangle_{1-b}^A$  from  $P_{1-b}$ , and computes  $\text{VCW}_i := \langle VCW_i \rangle_b^A + \langle VCW_i \rangle_{1-b}^A$ .
6.  $P_b$  computes  $k_b := (\langle \Delta \rangle_b \oplus W, \{CW_i\}_{i \in [1, n+1]}, \{VCW_i\}_{i \in [1, n]})$  and  $\langle \mathbf{r}^{(j)} \rangle_b^A := \text{DCF.Eval}(b, k_b, j)$  for  $j \in [0, N)$ , and outputs  $\langle \mathbf{r} \rangle_b^A \in \mathcal{R}^N$ .

Fig. 13: DCF correlation generation in the  $(\mathcal{F}_{\text{COT}}, \mathcal{F}_{\text{Rand}}, \mathcal{F}_{\text{OLE}})$ -hybrid model.

have payloads of  $(\lambda - 1) + 2 \cdot \log |\mathcal{R}|$  bits. Using  $n$  COT tuples, each party sends  $n + n \cdot (3\lambda - 1 + 5 \cdot \log |\mathcal{R}|) + \log |\mathcal{R}|$  bits, and the computation per party is dominated by the  $4N$  RP calls in GGM tree expansion in  $2n + 2$  rounds. Our savings in computation and round complexity are 37.5% and 50%, respectively. For a typical ring  $\mathcal{R}$  with size  $|\mathcal{R}| \approx 2^\lambda$ , the communication reduction is about 62.5%. When  $\mathcal{R}$  is sufficiently small, this reduction can be 66.6%.

**Complexity analysis (general ring).**  $\Pi_{\text{DCF}}$  also works for general  $\mathcal{R}$  at the cost of additionally using  $2n + 2$  precomputed OLE tuples. This general-ring version proceeds in  $n + 4$  rounds, and the overall outgoing communication per party is  $(n + 1) + (n + 1) \cdot \lambda + \frac{\lambda}{t} + (4n + 3) \cdot \log |\mathcal{R}|$  bits.

In contrast, the OT-based protocol [7] can run in  $2n + 3$  rounds. Each party sends at most  $n + n \cdot (3\lambda + 3 + 4 \cdot \log |\mathcal{R}|) + (3n + 3) \cdot \log |\mathcal{R}|$  bits and uses  $2n + 2$  OLE tuples. Our savings in communication and round complexity are about 50%  $\sim$  66.6% and 50%, respectively, for typical ring size  $|\mathcal{R}| \leq 2^\lambda$ .

## Acknowledgements

Work of Kang Yang is supported by the National Key Research and Development Program of China (Grant No. 2022YFB2702000), and by the National Natural Science Foundation of China (Grant Nos. 62102037, 61932019). Work of Xiao Wang is supported by DARPA under Contract No. HR001120C0087, NSF award #2016240, and research awards from Meta and Google. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. Work of Jiang Zhang is supported by the National Key Research and Development Program of China (Grant No. 2022YFB2702000), and by the National Natural Science Foundation of China (Grant Nos. 62022018, 61932019). Work of Zheli Liu is supported by the National Natural Science Foundation of China (Grant No. 62032012).

## References

1. Asharov, G., Lindell, Y., Schneider, T., Zohner, M.: More efficient oblivious transfer and extensions for faster secure computation. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013. pp. 535–548. ACM Press (Nov 2013). <https://doi.org/10.1145/2508859.2516738>
2. Baum, C., Malozemoff, A.J., Rosen, M.B., Scholl, P.: Mac’n’Cheese: Zero-Knowledge Proofs for Boolean and Arithmetic Circuits with Nested Disjunctions. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part IV. LNCS, vol. 12828, pp. 92–122. Springer, Heidelberg, Virtual Event (Aug 2021). [https://doi.org/10.1007/978-3-030-84259-8\\_4](https://doi.org/10.1007/978-3-030-84259-8_4)
3. Beaver, D.: Precomputing Oblivious Transfer. In: Coppersmith, D. (ed.) CRYPTO’95. LNCS, vol. 963, pp. 97–109. Springer, Heidelberg (Aug 1995). [https://doi.org/10.1007/3-540-44750-4\\_8](https://doi.org/10.1007/3-540-44750-4_8)
4. Bellare, M., Hoang, V.T., Keelveedhi, S., Rogaway, P.: Efficient Garbling from a Fixed-Key Blockcipher. In: 2013 IEEE Symposium on Security and Privacy. pp. 478–492. IEEE Computer Society Press (May 2013). <https://doi.org/10.1109/SP.2013.39>
5. Bhattacharya, S., Nandi, M.: Full Indifferentiable Security of the Xor of Two or More Random Permutations Using the  $\chi^2$  Method. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part I. LNCS, vol. 10820, pp. 387–412. Springer, Heidelberg (Apr / May 2018). [https://doi.org/10.1007/978-3-319-78381-9\\_15](https://doi.org/10.1007/978-3-319-78381-9_15)
6. Boneh, D., Boyle, E., Corrigan-Gibbs, H., Gilboa, N., Ishai, Y.: Lightweight Techniques for Private Heavy Hitters. In: 2021 IEEE Symposium on Security and Privacy. pp. 762–776. IEEE Computer Society Press (May 2021). <https://doi.org/10.1109/SP40001.2021.00048>
7. Boyle, E., Chandran, N., Gilboa, N., Gupta, D., Ishai, Y., Kumar, N., Rathee, M.: Function Secret Sharing for Mixed-Mode and Fixed-Point Secure Computation. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part II. LNCS, vol. 12697, pp. 871–900. Springer, Heidelberg (Oct 2021). [https://doi.org/10.1007/978-3-030-77886-6\\_30](https://doi.org/10.1007/978-3-030-77886-6_30)
8. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y.: Compressing Vector OLE. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018. pp. 896–912. ACM Press (Oct 2018). <https://doi.org/10.1145/3243734.3243868>
9. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Resch, N., Scholl, P.: Correlated Pseudorandomness from Expand-Accumulate Codes. In: Dodis, Y., Shrimp-

- ton, T. (eds.) CRYPTO 2022, Part II. LNCS, vol. 13508, pp. 603–633. Springer, Heidelberg (Aug 2022). [https://doi.org/10.1007/978-3-031-15979-4\\_21](https://doi.org/10.1007/978-3-031-15979-4_21)
10. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Rindal, P., Scholl, P.: Efficient Two-Round OT Extension and Silent Non-Interactive Secure Computation. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) ACM CCS 2019. pp. 291–308. ACM Press (Nov 2019). <https://doi.org/10.1145/3319535.3354255>
11. 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, Part III. LNCS, vol. 11694, pp. 489–518. Springer, Heidelberg (Aug 2019). [https://doi.org/10.1007/978-3-030-26954-8\\_16](https://doi.org/10.1007/978-3-030-26954-8_16)
12. 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, Part II. LNCS, vol. 12171, pp. 387–416. Springer, Heidelberg (Aug 2020). [https://doi.org/10.1007/978-3-030-56880-1\\_14](https://doi.org/10.1007/978-3-030-56880-1_14)
13. Boyle, E., Gilboa, N., Ishai, Y.: Function Secret Sharing: Improvements and Extensions. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016. pp. 1292–1303. ACM Press (Oct 2016). <https://doi.org/10.1145/2976749.2978429>
14. Boyle, E., Gilboa, N., Ishai, Y.: Secure Computation with Preprocessing via Function Secret Sharing. In: Hofheinz, D., Rosen, A. (eds.) TCC 2019, Part I. LNCS, vol. 11891, pp. 341–371. Springer, Heidelberg (Dec 2019). [https://doi.org/10.1007/978-3-030-36030-6\\_14](https://doi.org/10.1007/978-3-030-36030-6_14)
15. Canetti, R.: Universally Composable Security: A New Paradigm for Cryptographic Protocols. In: 42nd FOCS. pp. 136–145. IEEE Computer Society Press (Oct 2001). <https://doi.org/10.1109/SFCS.2001.959888>
16. Chen, S., Steinberger, J.P.: Tight Security Bounds for Key-Alternating Ciphers. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 327–350. Springer, Heidelberg (May 2014). [https://doi.org/10.1007/978-3-642-55220-5\\_19](https://doi.org/10.1007/978-3-642-55220-5_19)
17. Choi, S.G., Katz, J., Kumaresan, R., Zhou, H.S.: On the Security of the “Free-XOR” Technique. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 39–53. Springer, Heidelberg (Mar 2012). [https://doi.org/10.1007/978-3-642-28914-9\\_3](https://doi.org/10.1007/978-3-642-28914-9_3)
18. Couteau, G., Rindal, P., Raghuraman, S.: Silver: Silent VOLE and Oblivious Transfer from Hardness of Decoding Structured LDPC Codes. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part III. LNCS, vol. 12827, pp. 502–534. Springer, Heidelberg, Virtual Event (Aug 2021). [https://doi.org/10.1007/978-3-030-84252-9\\_17](https://doi.org/10.1007/978-3-030-84252-9_17)
19. Damgård, I., Nielsen, J.B., Nielsen, M., Ranellucci, S.: The TinyTable Protocol for 2-Party Secure Computation, or: Gate-Scrambling Revisited. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 167–187. Springer, Heidelberg (Aug 2017). [https://doi.org/10.1007/978-3-319-63688-7\\_6](https://doi.org/10.1007/978-3-319-63688-7_6)
20. Damgård, I., Pastro, V., Smart, N.P., Zakarias, S.: Multiparty Computation from Somewhat Homomorphic Encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 643–662. Springer, Heidelberg (Aug 2012). [https://doi.org/10.1007/978-3-642-32009-5\\_38](https://doi.org/10.1007/978-3-642-32009-5_38)
21. Dittmer, S., Ishai, Y., Ostrovsky, R.: Line-Point Zero Knowledge and Its Applications. In: 2nd Conference on Information-Theoretic Cryptography (2021). <https://doi.org/10.4230/LIPIcs.ITC.2021.5>



22. Doerner, J., shelat, a.: Scaling ORAM for Secure Computation. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017. pp. 523–535. ACM Press (Oct / Nov 2017). <https://doi.org/10.1145/3133956.3133967>
23. Garimella, G., Pinkas, B., Rosulek, M., Trieu, N., Yanai, A.: Oblivious Key-Value Stores and Amplification for Private Set Intersection. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part II. LNCS, vol. 12826, pp. 395–425. Springer, Heidelberg, Virtual Event (Aug 2021). [https://doi.org/10.1007/978-3-030-84245-1\\_14](https://doi.org/10.1007/978-3-030-84245-1_14)
24. Ghosh, S., Nielsen, J.B., Nilges, T.: Maliciously Secure Oblivious Linear Function Evaluation with Constant Overhead. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part I. LNCS, vol. 10624, pp. 629–659. Springer, Heidelberg (Dec 2017). [https://doi.org/10.1007/978-3-319-70694-8\\_22](https://doi.org/10.1007/978-3-319-70694-8_22)
25. 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 (May 2014). [https://doi.org/10.1007/978-3-642-55220-5\\_35](https://doi.org/10.1007/978-3-642-55220-5_35)
26. Goldreich, O., Goldwasser, S., Micali, S.: How to Construct Random Functions (Extended Abstract). In: 25th FOCS. pp. 464–479. IEEE Computer Society Press (Oct 1984). <https://doi.org/10.1109/SFCS.1984.715949>
27. Goldreich, O., Micali, S., Wigderson, A.: How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In: Aho, A. (ed.) 19th ACM STOC. pp. 218–229. ACM Press (May 1987). <https://doi.org/10.1145/28395.28420>
28. Guo, C., Katz, J., Wang, X., Yu, Y.: Efficient and Secure Multiparty Computation from Fixed-Key Block Ciphers. In: 2020 IEEE Symposium on Security and Privacy. pp. 825–841. IEEE Computer Society Press (May 2020). <https://doi.org/10.1109/SP40000.2020.00016>
29. Guo, X., Yang, K., Wang, X., Zhang, W., Xie, X., Zhang, J., Liu, Z.: Half-Tree: Halving the Cost of Tree Expansion in COT and DPF. Cryptology ePrint Archive, Report 2022/1431 (2022), <https://eprint.iacr.org/2022/1431>
30. Gupta, K., Kumaraswamy, D., Chandran, N., Gupta, D.: LLAMA: A Low Latency Math Library for Secure Inference. Privacy Enhancing Technologies Symposium (PETS 2022) (2022). <https://doi.org/10.56553/popets-2022-0109>
31. Hazay, C., Lindell, Y.: Efficient Secure Two-Party Protocols - Techniques and Constructions. ISC, Springer, Heidelberg (2010). <https://doi.org/10.1007/978-3-642-14303-8>
32. Hazay, C., Scholl, P., Soria-Vazquez, E.: Low Cost Constant Round MPC Combining BMR and Oblivious Transfer. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part I. LNCS, vol. 10624, pp. 598–628. Springer, Heidelberg (Dec 2017). [https://doi.org/10.1007/978-3-319-70694-8\\_21](https://doi.org/10.1007/978-3-319-70694-8_21)
33. Heath, D., Kolesnikov, V.: One Hot Garbling. In: Vigna, G., Shi, E. (eds.) ACM CCS 2021. pp. 574–593. ACM Press (Nov 2021). <https://doi.org/10.1145/3460120.3484764>
34. Ishai, Y., Kilian, J., Nissim, K., Petrank, E.: Extending Oblivious Transfers Efficiently. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 145–161. Springer, Heidelberg (Aug 2003). [https://doi.org/10.1007/978-3-540-45146-4\\_9](https://doi.org/10.1007/978-3-540-45146-4_9)
35. Keller, M., Orsini, E., Scholl, P.: MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016. pp. 830–842. ACM Press (Oct 2016). <https://doi.org/10.1145/2976749.2978357>

36. Keller, M., Pastro, V., Rotaru, D.: Overdrive: Making SPDZ Great Again. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 158–189. Springer, Heidelberg (Apr / May 2018). [https://doi.org/10.1007/978-3-319-78372-7\\_6](https://doi.org/10.1007/978-3-319-78372-7_6)
37. Kolesnikov, V., Schneider, T.: Improved Garbled Circuit: Free XOR Gates and Applications. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 486–498. Springer, Heidelberg (Jul 2008). [https://doi.org/10.1007/978-3-540-70583-3\\_40](https://doi.org/10.1007/978-3-540-70583-3_40)
38. Nielsen, J.B., Nordholt, P.S., Orlandi, C., Burra, S.S.: A New Approach to Practical Active-Secure Two-Party Computation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 681–700. Springer, Heidelberg (Aug 2012). [https://doi.org/10.1007/978-3-642-32009-5\\_40](https://doi.org/10.1007/978-3-642-32009-5_40)
39. Patarin, J.: The “Coefficients H” Technique (Invited Talk). In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. LNCS, vol. 5381, pp. 328–345. Springer, Heidelberg (Aug 2009). [https://doi.org/10.1007/978-3-642-04159-4\\_21](https://doi.org/10.1007/978-3-642-04159-4_21)
40. Rindal, P., Schoppmann, P.: VOLE-PSI: Fast OPRF and Circuit-PSI from Vector-OLE. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part II. LNCS, vol. 12697, pp. 901–930. Springer, Heidelberg (Oct 2021). [https://doi.org/10.1007/978-3-030-77886-6\\_31](https://doi.org/10.1007/978-3-030-77886-6_31)
41. Rogaway, P., Steinberger, J.P.: Constructing Cryptographic Hash Functions from Fixed-Key Blockciphers. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 433–450. Springer, Heidelberg (Aug 2008). [https://doi.org/10.1007/978-3-540-85174-5\\_24](https://doi.org/10.1007/978-3-540-85174-5_24)
42. Schoppmann, P., Gascón, A., Reichert, L., Raykova, M.: Distributed Vector-OLE: Improved Constructions and Implementation. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) ACM CCS 2019. pp. 1055–1072. ACM Press (Nov 2019). <https://doi.org/10.1145/3319535.3363228>
43. Weng, C., Yang, K., Katz, J., Wang, X.: Wolverine: Fast, Scalable, and Communication-Efficient Zero-Knowledge Proofs for Boolean and Arithmetic Circuits. In: 2021 IEEE Symposium on Security and Privacy. pp. 1074–1091. IEEE Computer Society Press (May 2021). <https://doi.org/10.1109/SP40001.2021.00056>
44. Weng, C., Yang, K., Xie, X., Katz, J., Wang, X.: Mystique: Efficient Conversions for Zero-Knowledge Proofs with Applications to Machine Learning. In: Bailey, M., Greenstadt, R. (eds.) USENIX Security 2021. pp. 501–518. USENIX Association (Aug 2021)
45. Yang, K., Sarkar, P., Weng, C., Wang, X.: QuickSilver: Efficient and Affordable Zero-Knowledge Proofs for Circuits and Polynomials over Any Field. In: Vigna, G., Shi, E. (eds.) ACM CCS 2021. pp. 2986–3001. ACM Press (Nov 2021). <https://doi.org/10.1145/3460120.3484556>
46. Yang, K., Weng, C., Lan, X., Zhang, J., Wang, X.: Ferret: Fast Extension for Correlated OT with Small Communication. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020. pp. 1607–1626. ACM Press (Nov 2020). <https://doi.org/10.1145/3372297.3417276>
47. Zahur, S., Rosulek, M., Evans, D.: Two Halves Make a Whole - Reducing Data Transfer in Garbled Circuits Using Half Gates. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 220–250. Springer, Heidelberg (Apr 2015). [https://doi.org/10.1007/978-3-662-46803-6\\_8](https://doi.org/10.1007/978-3-662-46803-6_8)