# Gold OPRF: Post-Quantum Oblivious Power-Residue PRF

Yibin Yang*, Fabrice Benhamouda†, Shai Halevi†, Hugo Krawczyk† and Tal Rabin†
*Georgia Institute of Technology, USA
Emails: yyang811@gatech.edu
†Amazon Web Services, USA
Emails: {fabrice.benhamouda, shai.halevi, hugokraw}@gmail.com, talr@seas.upenn.edu

***Abstract*—We propose plausible post-quantum (PQ) oblivious pseudorandom functions (OPRFs) based on the Power-Residue PRF (Damgård CRYPTO'88), a generalization of the Legendre PRF. For security parameter $\lambda$, we consider the PRF $\mathsf{Gold}_k(x)$ that maps an integer $x$ modulo a public prime $p = 2^\lambda \cdot g + 1$ to the element $(k+x)^g \bmod p$, where $g$ is public and $\log g \approx 2\lambda$.**

**At the core of our constructions are efficient novel methods for evaluating Gold within two-party computation (2PC-Gold), achieving different security requirements. Here, the server $\mathcal{P}_s$ holds the PRF key $k$ whereas the client $\mathcal{P}_c$ holds the PRF input $x$, and they jointly evaluate Gold in 2PC. 2PC-Gold uses standard Vector Oblivious Linear Evaluation (VOLE) correlations and is *information-theoretic* and *constant-round* in the (V)OLE-hybrid model. We show:**

- **For a semi-honest $\mathcal{P}_s$ and a malicious $\mathcal{P}_c$: a 2PC-Gold that just uses a single (V)OLE correlation, and has a communication complexity of 3 field elements (2 field elements if we only require a uniformly sampled key) and a computational complexity of $\mathcal{O}(\lambda)$ field operations. We refer to this as *half-malicious* security.**
- **For malicious $\mathcal{P}_s$ and $\mathcal{P}_c$: a 2PC-Gold that just uses $\frac{\lambda}{4} + \mathcal{O}(1)$ VOLE correlations, and has a communication complexity of $\frac{\lambda}{4} + \mathcal{O}(1)$ field elements and a computational complexity of $\mathcal{O}(\lambda)$ field operations.**

**These constructions support additional features and extensions, e.g., batched evaluations with better amortized costs where $\mathcal{P}_c$ repeatedly evaluates the PRF under the same key.**

**Furthermore, we extend 2PC-Gold to Verifiable OPRFs and use the methodology from Beullens et al. (Eurocrypt'25) to get strong OPRF security in the universally composable setting.**

**All the protocols are efficient in practice. We implemented 2PC-Gold—with (PQ) VOLEs—and benchmarked them. For example, our half-malicious (resp. malicious) $n$-batched PQ OPRFs incur about 100B (resp. 1.9KB) of amortized communication for $\lambda = 128$.**

## 1. Introduction

**OPRF.** *Pseudorandom functions* [39] (PRFs) are essential tools in cryptography. Modern applications often require the evaluation of PRFs in a distributed and privacy-preserving manner. This leads to the notion of *oblivious pseudorandom functions* (OPRFs) [37], [65].

Concretely, consider a PRF $F$, a server $\mathcal{P}_s$ that holds a PRF key $k$, and a client $\mathcal{P}_c$ that holds a PRF input $x$. Through the OPRF protocol, $\mathcal{P}_c$ learns the output $F_k(x)$ without obtaining any additional information about $k$ while $\mathcal{P}_s$ learns nothing about $x$ (including nothing about the output $F_k(x)$). Alternatively, an OPRF can be viewed as a specialized instance of *secure two-party computation* (2PC) [84] tailored for evaluating a PRF.

It is important to note that OPRFs are sometimes defined with more stringent requirements than a standard 2PC over a PRF, i.e., they demand additional properties. However, the core component—2PC over PRFs—underlies also the stronger notions and we focus on the design of such a component. We then show how to extend the design to obtain stronger forms of OPRFs.

**Post-Quantum OPRF.** Currently, the most widely deployed OPRF protocols rely on Diffie-Hellman-type assumptions, with the 2HashDH OPRF [50] being a notable example. While these protocols are lightweight and highly efficient, they are *insecure* against quantum adversaries [75], [76]. In light of the potential threats posed by quantum computing, it is imperative to develop OPRF protocols that remain secure in the quantum era.

Over the past five years, numerous proposals for *post-quantum* (PQ) OPRFs have emerged, e.g., [1], [2], [3], [4], [7], [13], [17], [29], [33], [34], [48], [49], [74]. However, a significant efficiency gap remains: these PQ OPRFs are primarily of theoretical interest and are not yet practical for widespread deployment. Addressing this efficiency gap is crucial for practical adoption of PQ OPRFs, and our work focuses on constructing novel PQ OPRFs to bridge this gap.

**Challenges.** Constructing an efficient PQ OPRF presents two primary and intertwined challenges. First, it requires identifying or designing a suitable PQ PRF. Second, it involves developing methods to efficiently evaluate the chosen PQ PRF within 2PC based on PQ assumptions.

### 1.1. Our Results

In this work, we aim to bridge the efficiency gap by leveraging the *Power-Residue PRF* [24], a generalization

*The full version [83] is available online: https://eprint.iacr.org/2024/1955.*

of the Legendre PRF. Let $\lambda$ denote the security parameter (e.g., $\lambda = 128$, which produces 128-bit outputs, targeting the NIST Security Strength Category 1 for PQ cryptography). To get an $\mathcal{O}(\lambda)$-bit output, we consider the following PRF:

$$F_k(x) := (k + x)^g \bmod p$$

where $p = 2^\lambda \cdot g + 1$ is a *non-secret* prime, $k \in \mathbb{F}_p$ is the key, and $x \in \mathbb{F}_p$ is the input to the PRF.

Crucially, when $p$ is sufficiently large ($\log p \approx 3\lambda$), there are currently *no* known quantum attacks against this PRF.[1]

**Gold.** We name this PRF Gold since it has a private base and a public exponent, whereas the well-known discrete logarithm has a public base and a private exponent—reversing the order of characters in dlog gives Gold.

**2PC-Gold.** Our **first major contribution** is developing a family of efficient 2PC protocols, called 2PC-Gold, for evaluating the Gold function across various settings between a server $\mathcal{P}_s$ that inputs a key $k$ and a client $\mathcal{P}_c$ with input $x$. 2PC-Gold protocols leverage standard *Vector Oblivious Linear Evaluation* (VOLE) correlations (recalled in Section 2.4) in a *black-box* manner and are *information-theoretically* secure and *constant-round* in the VOLE-hybrid model.[2] Our 2PC-Gold protocols include:

- A protocol secure against an *unbounded semi-honest* $\mathcal{P}_s$ and an *unbounded malicious* $\mathcal{P}_c$ that uses a single (V)OLE correlation, and has a communication complexity of 3 field elements (2 if only a uniformly sampled key is required) and a computational complexity of $\mathcal{O}(\lambda)$ field operations (essentially, the cost of a single exponentiation). This protocol requires 3 rounds (2 if only a uniformly sampled key is required). We refer to this as our *half-malicious* protocol. It is useful in scenarios where the server is trusted not to depart from its intended behavior.
- An enhanced protocol secure against *unbounded malicious* $\mathcal{P}_s$ and $\mathcal{P}_c$ that uses $\lambda + 8$ VOLE correlations, and has a communication complexity of $\lambda + 15$ field elements and a computational complexity of $\mathcal{O}(\lambda)$ field operations. We further show that for any small constant $\phi$ that, w.l.o.g., divides $\lambda$, the communication complexity can be improved to $\frac{\lambda}{\phi} + 2^\phi + 13$ field elements, and the number of VOLE correlations required is reduced to $\frac{\lambda}{\phi} + \mathcal{O}(1)$. This protocol requires 5 rounds (3 with the Fiat-Shamir transformation [36]). We refer to this as our *malicious* protocol.

The above protocols have additional features and extensions useful in different applications, including:

- **Offline-Online Mode:** Our half-malicious and malicious 2PC-Gold support an offline-online mode. The generation of VOLE correlations can be performed during an *input/key-independent* offline phase. Moreover, in our malicious 2PC-Gold, the majority of the work can be further shifted to the offline phase, resulting in an *online communication complexity of only* 6 $\mathbb{F}_p$ *elements.*
- **Batching:** Our protocols support batched evaluations for better amortized costs, allowing $\mathcal{P}_c$ to repeatedly evaluate the PRF under the *same* key with arbitrary inputs. For example, the online communication of our malicious 2PC-Gold can be reduced to amortized $2 + \frac{4}{n}$ $\mathbb{F}_p$ elements for $n$-batched evaluations.
- **Classical and PQ Instantiations:** Our half-malicious and malicious 2PC-Gold rely solely on standard VOLE correlations and are *information-theoretically secure* in the VOLE-hybrid model.[3] By employing appropriate methods to generate these VOLE correlations, our unmodified protocols can be instantiated to achieve either classical or post-quantum security in the plain model. We implement both options and report performance in Section 7. This performance will further improve in the future with ongoing optimizations in VOLE generation.
- **Key Verification:** Malicious 2PC-Gold supports *zero-knowledge proofs* [40] of *any* NP relation over the key. Notably, we build efficient verifiable OPRFs based on this property.

See Section 3 for a concise technical overview of 2PC-Gold.

**O-Gold and UC-Gold.** OPRFs have been defined in multiple ways in the literature. In its basic form, just the 2PC over a PRF between server and client (i.e., 2PC-Gold) suffices for the OPRF definition. This, however, is not sufficient for some applications and stronger OPRF notions have been formulated. Our **second major contribution** is augmenting 2PC-Gold to achieve these notions.

Let $\mathsf{H}_1$ be a hash function mapping arbitrary strings to $\mathbb{F}_p$ elements and $\mathsf{H}_2$ be another hash function producing $\{0,1\}^{2\lambda}$ elements. We define the function $\mathsf{O\text{-}Gold}_k(x)$ as $\mathsf{H}_2(x, \mathsf{Gold}_k(\mathsf{H}_1(x)))$. A simple 2PC over O-Gold (leaking[4] $\mathsf{Gold}_k(\mathsf{H}_1(x))$ to $\mathcal{P}_c$) can be implemented by $\mathcal{P}_c$ inputting $\mathsf{H}_1(x)$ to 2PC-Gold, obtaining $y := \mathsf{Gold}_k(\mathsf{H}_1(x))$ and locally computing $\mathsf{H}_2(y)$. O-Gold function inherits the 2PC security of Gold and, in addition, supports arbitrary inputs, avoids collisions under the same $k$, and allows modeling its outputs as a random oracle. It also serves as a basis for a *verifiable* OPRF and to achieve strong *universal-composable* security.

We show how different forms of *verifiability* can be efficiently implemented for our OPRFs, including the approach to achieve standard verifiable OPRFs [50]; see Section 6.1. Finally, we show how to extend O-Gold to achieve strong UC security by following the methodology from [13] in Section 6.1; we call the resultant OPRF UC-Gold.[5]

We refer to the three constructions—2PC-Gold, O-Gold, and UC-Gold—as OPRFs, each distinguished by its features.

1. There exists an efficient quantum distinguisher if it is allowed to make quantum queries [72], [79]. In this work, we restrict ourselves to (quantum) attacks with classical queries and justify this choice in Section 2.3.

2. To clarify, our 2PC protocols for evaluating Gold is *information-theoretically* secure in the VOLE-hybrid model. However, when we employ this protocol as an OPRF, we rely on computational assumptions to guarantee Gold being a secure PRF and to generate VOLE correlations.

3. We note that looking ahead, our simulation does not need *rewinding* in the VOLE-hybrid model.

4. This leakage to $\mathcal{P}_c$ is usually harmless in OPRF applications.

5. We note that [13]'s proof considers the classical random oracle model; extension to the Quantum ROM [15] is an interesting topic for future work.

**Implementation.** All our protocols are efficient in practice. *We implemented half-malicious and malicious* 2PC-Gold*, with classical and PQ VOLEs.*[6] We report the performance in Section 7. The cost of O-Gold is similar to 2PC-Gold, and we estimate concrete overheads for UC-Gold.

Consider $\lambda = 128$. Our half-malicious (resp. malicious) non-batched PQ 2PC-Gold needs 774KB (resp. 970KB) of communication, 568ms (resp. 1.1s) of wall-clock time in a WAN-like network, and 163ms (resp. 510ms) of wall-clock time in a LAN-like network. Most of the overhead comes from generating standard PQ VOLE correlations: it only needs 96B (resp. 2.7KB) in the VOLE-hybrid world. Indeed, by deploying the sublinear generation of large enough PQ VOLE correlations ($n \approx 10^7$ for half-malicious and $n \approx 3 \times 10^5$ for malicious in our experiments, but $n$ can be much smaller; see Section 7.2) from [18], our half-malicious (resp. malicious) batched PQ 2PC-Gold only needs 100B (resp. 1.9KB) of amortized communication, 87µs (resp. 1.6ms) of amortized wall-clock time in a WAN-like network, and 57µs (resp. 1ms) of amortized wall-clock time in a LAN-like network. Finally, achieving UC-Gold requires an additional (amortizable) 899KB of communication.

**Full Version.** The full version of this paper, which includes all proofs, is available at [83].

## 1.2. Related Work

**OPRF and Its Applications.** The concept of obliviously evaluating PRFs dates back to [65] and was later refined and formalized as Oblivious PRF (OPRF) by [37].

Informally, OPRFs increase the entropy of the client's inputs, making them essential components in many real-world privacy-preserving applications. Notable examples include private set intersection (e.g., [21], [58]), key management (e.g., [52]), anonymous tokens (e.g., [27], [59]), and password-based key-exchange (e.g., [53]). Many OPRF applications (e.g., PSI and key management) can benefit from the batched OPRF.

One of the most successful OPRF constructions is the 2HashDH OPRF [50]. This construction relies on Diffie-Hallman-type assumptions and is extremely efficient: the basic version requires only 2 group elements of communication per evaluation. For more details, we refer the reader to a *Systematization of Knowledge* work on OPRFs [20].

**PQ OPRF Constructions.** Research on PQ OPRF schemes has emerged in the past five years.

One approach is to build PQ OPRFs based on *isogenies*. This was initiated by [17], relying on the SIDH and CSIDH assumptions. Subsequent work [7], [48] optimized these constructions and addressed vulnerabilities exposed by SIDH attacks [8], [61] over [17]'s constructions. However, isogeny-based OPRFs remain computationally intensive.

Another approach is to build PQ OPRFs based on *lattices*, initiated by the work [3]. This is merely a feasibility result: [3]'s constructions require more than 2MB (resp. 128GB) of communication per evaluation for semi-honest (resp. malicious) security. Very recent work [4], [33] optimized the communication cost to amortized under 200KB for malicious security over batched evaluations. To our best knowledge, the only available implementation of lattice-based OPRF protocols is the one in [3] for semi-honest security.

An alternative, straightforward approach to constructing PQ OPRFs is to apply Yao's *Garbled Circuits* [84] to, e.g., AES with the help of post-quantum *oblivious transfers* (e.g., [32], [62]). This idea was explored in [34]. However, the size of the garbled circuit for AES is relatively large, making further optimization inherently hard.

This approach, however, opens up the possibility of building PQ OPRFs based on PQ MPC-friendly PRFs. The literature has explored this method using two such PRFs:

- **"Crypto Dark Matter" PRFs [16]:** These PRFs are specifically designed for efficient use within MPC. Several proposals have explored the oblivious evaluations of these PRFs. In particular, [29] studied preprocessing 2PC over them, and [2] studied TFHE over them. Most recently, [1] has aggressively improved the constructions of these PRFs and showed that the amortized communication cost per evaluation over batched evaluations can only be ≈1Kbit for semi-honest security. However, these PRF constructions rely on new assumptions that need further study. Moreover, achieving malicious security may require substantial effort, e.g., [2] relies on a heuristic argument to achieve it. Furthermore, these PRFs are only *weak* PRFs (i.e., they are only secure when the input is chosen uniformly at random). In the semi-honest case, this is not an issue as the input can be hashed; however, in the malicious case, additional mechanisms are required.
- **Legendre PRFs [24]:** Another type of MPC-friendly PRF is the *Legendre PRF*, originally proposed by Damgård in 1988. This PRF was first identified as MPC-friendly by [42] and has since been employed to construct PQ OPRFs in works [13], [74]. The state-of-the-art construction presented in the recent work of [13] incurs approximately 900KB of (non-amortizable) communication per evaluation for malicious security.

Our work constructs PQ OPRFs using the generalized Legendre PRF, i.e., the Power-Residue PRF also introduced in [24]. See Section 2.3 for more detailed discussions of the Legendre and Power-Residue PRFs.

Recent work [54] considers building PQ distributed OPRF from the Legendre PRF. More specifically, [54] considers distributed OPRF with multiple ($\geq$3) non-colluding servers (i.e., the threshold setting), while we consider the classical OPRF setting with a single server. Meanwhile, for $n$ servers, the cost in [54] is exponential in $n$, making it unsuitable for many servers. Extending our protocols to the distributed OPRF setting is an interesting future direction.

We note that our work may appear to extend the recent study by [13], which also utilizes VOLE correlations; however, this is not the case. In particular, our protocols use VOLE correlations in the opposite direction to build a novel customized 2PC for PRF evaluation, enabling efficient

6. Our implementation: https://github.com/gconeice/PR-OPRF.

extension to batched (amortized) evaluations and malicious security. Moreover, our protocols rely solely on black-box usage of the *standard* VOLE correlations, whereas [13] requires non-black-box modifications of the VOLE functionality. Finally, our work provides a complete end-to-end implementation along with a comprehensive benchmark. At the same time, our work borrows from [13] the elegant methodology for building the strong UC variant of Gold.

**VOLE and VOLE-Based ZK.** Our protocols leverage the standard VOLE correlations; see Section 2.4. In particular, our malicious OPRFs utilize efficient zero-knowledge proofs based on VOLE, commonly referred to as VOLE-based ZK; see Section 2.5.

**Concrete Performance Comparison.** Section 7.3 includes a concrete performance comparison between our protocols and prior work.

## 2. Preliminaries

### 2.1. Notation

- $\lambda$ is the security parameter (e.g., 128).
- The server is $\mathcal{P}_s$. We refer to $\mathcal{P}_s$ by she, her, hers...
- The client is $\mathcal{P}_c$. We refer to $\mathcal{P}_c$ by he, his, him...
- $x := y$ denotes that $y$ is *assigned* to $x$.
- We denote sets by upper-case letters. We denote that $x$ is uniformly drawn from a set $S$ by $x \xleftarrow{\$} S$.
- $[m] := \{1, \ldots, m\}$ and $[a, b) := \{a, a+1, \ldots, b-1\}$.
- We use $\mathbb{F}$ to denote a finite field. For a prime $p$, we use $\mathbb{F}_p$ to denote the standard modular-$p$ arithmetic field.
- We use $\mathbb{Z}$ to denote the integer ring.
- We use $\mathbb{G}$ to denote a cyclic group.
- We denote row vectors by bold lower-case letters (e.g., $\boldsymbol{a}$), where $a_i$ (or $a[i]$) denotes the $i$-th component of $\boldsymbol{a}$ (starting from 1). $|\boldsymbol{a}|$ denotes the length of $\boldsymbol{a}$.
- For vectors $\boldsymbol{a}, \boldsymbol{b}$ where $|\boldsymbol{a}| = |\boldsymbol{b}|$. $\langle \boldsymbol{a}, \boldsymbol{b} \rangle$ denotes the inner product; $\boldsymbol{a} \odot \boldsymbol{b}$ denotes the element-wise product.
- $x\boldsymbol{a} := (xa_1, \ldots xa_m)$ for some $\boldsymbol{a} = (a_1, \ldots, a_m)$.
- We use $n$ to denote the number of OPRF evaluations.

### 2.2. Security Model

We formalize our protocols in the universally composable (UC) framework [19]. We consider static corruptions. In particular, our basic OPRF protocol is secure against a semi-honest $\mathcal{P}_s$ and a malicious $\mathcal{P}_c$. We refer to this protocol as the *half-malicious* OPRF. On the other hand, our malicious protocol is secure against malicious $\mathcal{P}_s$ and $\mathcal{P}_c$. For simplicity, we omit standard UC (sub-)session IDs. We also omit that each time a party sends a message to the functionality, it provides a receipt to the simulator, and that whenever the functionality needs to deliver an output to a party, it waits for the simulator's instruction to do so. Note, this includes the standard *security with abort*.

### 2.3. Legendre and Power-Residue PRFs

Pseudorandom Functions (PRFs) [39] are deterministic keyed functions that are indistinguishable from random functions, when the key is chosen uniformly at random.

Our protocols essentially obliviously evaluate the Power-Residue PRF, which is a generalization of the Legendre PRF. Both PRFs can be traced back to the work of Damgård [24] in 1988. This section briefly reviews these two PRFs and existing (classical and quantum) attacks.

**Legendre PRF.** Let $p$ be a prime. For $a \in \mathbb{F}_p$, the Legendre symbol $\left(\frac{a}{p}\right)$ is defined as: $\left(\frac{a}{p}\right) := a^{\frac{p-1}{2}} \pmod p$.

It is conjectured that for a sufficiently large *public* prime $p$, the function $F_k(x) := \left(\frac{k+x}{p}\right)$ is a PRF defined over $\mathbb{F}_p \mapsto \{1, -1\}$, with key $k \in \mathbb{F}_p$. The output of $F_k$ on $-k$ is zero, which is neither 1 nor $-1$. However, since $k$ is supposed to be chosen uniformly, an adversary cannot find the single input $x = -k$ that results in an output of zero.[7]

**Power-Residue PRF.** A glaring downside of the Legendre PRF is that it only produces a single-bit output. In many scenarios, (oblivious) PRF is only useful when it can produce $\mathcal{O}(\lambda)$ bits. A direct way to achieve this is by repeating the Legendre PRF $\mathcal{O}(\lambda)$ times (e.g., on $\mathcal{O}(\lambda)$ different keys; see [13]). This is clearly undesirable. Instead, we exploit the higher order residues, a natural generalization of the Legendre PRF also proposed by [24].

In more detail, consider a sufficiently large prime $p = eg + 1$ where $p, e, g$ are public. For a random key $k \xleftarrow{\$} \mathbb{F}_p$, the PRF $F$ on input $x \in \mathbb{F}_p$ is defined as:

$$F_k(x) := (k+x)^{\frac{p-1}{e}} = (k+x)^g \pmod p$$

This PRF can produce $e$ different non-zero outputs (and a unique input results in a zero output). This is because, as indicated by its name, each non-zero output is a $g$-th residue (aka an $e$-th root of unity). This is crucial for efficiency, as now one evaluation can produce a $\lfloor \log e \rfloor$-bit output. Note, for $e = 2$, the Power-Residue PRF is the Legendre PRF.

In this work, we set $e$ to be $2^\lambda$ and name this PRF Gold. One evaluation of Gold can produce $\lambda$-bit entropy. See Section 4.1 for how to transfer the output to $\{0,1\}^\lambda$, or $\mathcal{O}(\lambda)$ bits in general, and for associated formal assumptions.

**Statistical Properties.** The study of the distribution of Legendre symbols dates to Davenport, 1931 [25], followed by extensive work (e.g., [26], [28], [45], [63], [67], [78]). It points to statistical properties of Legendre symbols that resemble the behavior of fair coin tosses. Some of these properties can be extended to Power-Residue symbols [26].

**Attacks and Quantum Resistance.** Several works attempt to break the Legendre and Power-Residue PRFs. This trend has become popular recently, mainly because of Legendre's candidate usage in the Ethereum 2.0 blockchain.

7. We note that in (UC) OPRF constructions, the UC environment may choose the input $x = -k$. This is also true for OPRFs based on the Power-Residue PRF, and we handle this case carefully.

**Functionality $\mathcal{F}_{\mathsf{VOLE}}^{a,b}$**

$\mathcal{F}_{\mathsf{VOLE}}^{a,b}$, parameterized by a field $\mathbb{F}$, running with a sender $\mathcal{P}_a$, a receiver $\mathcal{P}_b$ and an adversary $\mathcal{S}$, proceeds as follows:

**Initialize.** Upon receiving (`init`) from $\mathcal{P}_a$ and $\mathcal{P}_b$, if $\mathcal{P}_b$ is honest, set $\Delta^{(b)} \xleftarrow{\$} \mathbb{F}$, else receive $\Delta^{(b)}$ from $\mathcal{S}$. Store $\Delta^{(b)}$. Send $\Delta^{(b)}$ to $\mathcal{P}_b$. Ignore subsequent (`init`).

**Extend.** Upon receiving (`extend`, $m$) from $\mathcal{P}_a$ and $\mathcal{P}_b$:

- If $\mathcal{P}_b$ is honest, set $\boldsymbol{v} \xleftarrow{\$} \mathbb{F}^m$, else receive $\boldsymbol{v} \in \mathbb{F}^m$ from $\mathcal{S}$.
- If $\mathcal{P}_a$ is honest, set $\boldsymbol{u} \xleftarrow{\$} \mathbb{F}^m$ and compute $\boldsymbol{w} := \boldsymbol{v} - \boldsymbol{u}\Delta^{(b)} \in \mathbb{F}^m$, else receive $\boldsymbol{u} \in \mathbb{F}^m$ and $\boldsymbol{w} \in \mathbb{F}^m$ from $\mathcal{S}$ and compute $\boldsymbol{v} := \boldsymbol{w} + \boldsymbol{u}\Delta^{(b)} \in \mathbb{F}^m$.
- Send $(\boldsymbol{u}, \boldsymbol{w})$ to $\mathcal{P}_a$ and $\boldsymbol{v}$ to $\mathcal{P}_b$.

Figure 1: The VOLE correlation functionality [81].

Both Legendre and Power-Residue PRF can be broken in quantum polynomial time with a single *quantum superposition query access* to the PRF [72], [79]. However, currently, cryptographic protocols and algorithms are only implemented on classical computers, including OPRFs. As long as all systems holding the OPRF key are classical, an adversary cannot make a quantum superposition of queries. Thus, we restrict ourselves to attacks with classical queries.

For the Legendre PRF, classical and quantum attacks include [12], [38], [55], [56], [57], [64], [74]. In particular, the best-known (quantum) attack by [38] has time complexity $2^{\mathcal{O}(\log p)} \cdot p^{1/3}$. For the Power-Residue PRF, the additional (quantum) attack by [12] has time complexity $\mathcal{O}(\frac{p \log^2 p}{Le \log^2 e})$ where $L \leq p^{1/4}$ denotes the number of queries.

Finally, a recent elegant work [23] considers an extension of the Legendre PRF to any modulo $N$ instead of a prime modulo $p$. The authors break the pseudorandomness of such a function when $N$ is composite and can be factored, which implies PQ insecurity [75]. However, this factoring-based attack (and subsequent work [22]) does *not* apply to the Legendre and Power-Residue PRFs modulo a prime $p$.

**Legendre vs. Power-Residue with $e = 2^\lambda$.** When $e = 2^\lambda$, breaking the Power-Residue PRF is not harder than breaking the Legendre PRF. Indeed, one can trivially compute $(k+x)^{\frac{p-1}{2}}$ from $(k+x)^{\frac{p-1}{e}}$. Conversely, we do not know of any reduction from breaking the Legendre PRF to the Power-Residue PRF. We leave it as an interesting open problem.

Nevertheless, considering the previously cited works, we believe that our Power-Residue PRF Gold is a natural, reasonable assumption (see the formal definition in Section 4.1) and encourage further study on its pseudorandomness.

## 2.4. Vector OLE

Our OPRFs use Vector Oblivious Linear Evaluation (VOLE) correlations: they are in the VOLE-hybrid model. VOLE correlations of size $m$ (or $m$ VOLE correlations) over a field $\mathbb{F}$ are correlated random vectors held by a sender $\mathcal{P}_a$ and a receiver $\mathcal{P}_b$. The sender $\mathcal{P}_b$ obtains a uniformly sampled scalar $\Delta^{(b)} \xleftarrow{\$} \mathbb{F}$ and a vector $\boldsymbol{v} \xleftarrow{\$} \mathbb{F}^m$. The receiver $\mathcal{P}_a$ obtains a uniformly sampled vector $\boldsymbol{u} \xleftarrow{\$} \mathbb{F}^m$ and a correlated vector $\boldsymbol{w} := \boldsymbol{v} - \boldsymbol{u}\Delta^{(b)}$.

Our protocols rely solely on a weaker form of VOLE correlations known as *endemic* [62] VOLE correlations, where the adversary is allowed to choose its own correlated shares. See Figure 1 for the formal definition.

Looking ahead, our half-malicious OPRFs use $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$: $\mathcal{P}_s$ plays the role of the VOLE receiver, whereas $\mathcal{P}_c$ plays the role of the VOLE sender. Our malicious OPRFs also use $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$ (with $\mathcal{P}_s$'s and $\mathcal{P}_c$'s roles are switched in VOLE).

Multiple ways exist to realize $\mathcal{F}_{\mathsf{VOLE}}$ with malicious security. In this work, we adopt the approaches in [13] and [81]. Due to space constraints, we defer the review of these techniques to the full version [83]; see Section 7.1 for some implementation details. Here, we only highlight that these approaches only rely on minicrypt-type cryptographic primitives (e.g., PRF, PRG) and/or the *Learning Parity with Noise* (LPN) assumption [14], given access to oblivious transfers (OTs) [69]. Thus, when using post-quantum OTs (e.g., [32], [62]), our instantiations of $\mathcal{F}_{\mathsf{VOLE}}$ are PQ secure:

**Lemma 1** (Informal)**.** *Assuming the existence of one-way functions or the hardness of LPN, there exists a protocol that UC-emulates $\mathcal{F}_{\mathsf{VOLE}}$ (Figure 1) in the OT-hybrid model.*

## 2.5. VOLE-Based ZK

Our malicious OPRFs also use VOLE correlations to build lightweight ZKPs to protect against a malicious $\mathcal{P}_s$, aligning with recent progress on VOLE-based ZK (e.g., [30], [47], [82]). Here, $\mathcal{P}_s$ is the prover, whereas $\mathcal{P}_c$ is the verifier. We review related techniques in this section.

**IT-MAC.** Consider a single VOLE correlation generated by $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$ defined over $\mathbb{F}_p$. That is, $\mathcal{P}_s$ holds $u$ and $w_u$ whereas $\mathcal{P}_c$ holds $\Delta^{(c)}$ and $v_u$, such that $v_u = w_u + u\Delta^{(c)}$. This correlation can be rather interpreted as the *information-theoretic message authentication code* (IT-MAC) [11], [66] commitment over the value $u$, from $\mathcal{P}_s$ to $\mathcal{P}_c$. We denote this correlation as $\langle (u, w_u), v_u \rangle_{\Delta^{(c)}}$ or $[u]_{\Delta^{(c)}}$ in short. IT-MAC commitments have the following properties:

- **Perfect Hiding:** $v_u$ and $\Delta^{(c)}$, held by $\mathcal{P}_c$, are independent of the committed value $u$.
- **Statistical Binding:** $\mathcal{P}_s$ can open $[u]_{\Delta^{(c)}}$ by sending $u$ and $w_u$, where $\mathcal{P}_c$ would check if $v_u \stackrel{?}{=} w_u + u\Delta^{(c)}$. For a malicious $\mathcal{P}_s$ to open it to a different value $u' \neq u$, she has to guess $\Delta^{(c)}$—succeed with probability $\frac{1}{p}$.
- **Linear Homomorphism:** Suppose $\mathcal{P}_s$ and $\mathcal{P}_c$ hold $\langle (x, w_x), v_x \rangle_{\Delta^{(c)}}$ and $\langle (y, w_y), v_y \rangle_{\Delta^{(c)}}$. For any public $\alpha, \beta, \gamma \in \mathbb{F}_p$, parties can locally generate $[\alpha x + \beta y + \gamma]_{\Delta^{(c)}}$ as follows: $\mathcal{P}_s$ computes $\alpha w_x + \beta w_y$, whereas $\mathcal{P}_c$ computes $\alpha v_x + \beta v_y + \gamma\Delta^{(c)}$.

Note, the linear homomorphism property implies that, for a random IT-MAC $[u]_{\Delta^{(c)}}$ (from VOLE), $\mathcal{P}_s$ can send $z - u$ to commit to $z$, where $u$ acts as a one-time pad [10].

**Line-Point Zero-Knowledge.** Our malicious OPRFs need $\mathcal{P}_s$ to prove in ZK that three IT-MAC commitments form

a multiplication triple. Namely, $\mathcal{P}_s$ and $\mathcal{P}_c$ hold $[x]_{\Delta^{(c)}} = \langle (x, w_x), v_x \rangle_{\Delta^{(c)}}$, $[y]_{\Delta^{(c)}} = \langle (y, w_y), v_y \rangle_{\Delta^{(c)}}$, $[z]_{\Delta^{(c)}} = \langle (z, w_z), v_z \rangle_{\Delta^{(c)}}$, where $\mathcal{P}_s$ wants to prove in ZK to $\mathcal{P}_c$ that $z = xy$. This can be done using the *line-point zero-knowledge* (LPZK) technique [30], [82]. LPZK relies on the following equality (let $\Delta = \Delta^{(c)}$ for simplicity):

$$\begin{aligned}
&\overbrace{v_x v_y - v_z \Delta}^{\text{known by } \mathcal{P}_c} \\
=&(x\Delta + w_x)(y\Delta + w_y) - (z\Delta + w_z)\Delta \\
=&\underbrace{(xy - z)}_{0}\Delta^2 + \underbrace{(xw_y + yw_x - w_z)}_{\text{known by } \mathcal{P}_s}\Delta + \underbrace{w_x w_y}_{\text{known by } \mathcal{P}_s}
\end{aligned}$$

Hence, if ZK is not required, $\mathcal{P}_s$ can send two coefficients $C_1 = xw_y + yw_x - w_z$ and $C_0 = w_x w_y$, and $\mathcal{P}_c$ will check if $v_x v_y - v_z \Delta \stackrel{?}{=} C_1 \Delta + C_0$. This is sound because the equality holds only when $\Delta$ happens to be the root of a quadratic equation if $xy \neq z$. If $\Delta$ has full entropy to $\mathcal{P}_s$, this will only happen with probability $\frac{2}{p}$ since this is a degree-2 polynomial [73], [86]. Of course, this is not ZK as $C_1, C_0$ are correlated with $x, y, z$. ZK can be recovered by consuming one random IT-MAC $[r]_{\Delta^{(c)}}$ (from VOLE). That is, $\mathcal{P}_s$ sends $C_1 = xw_y + yw_x - w_z + r$ and $C_0 = w_x w_y + w_r$ and $\mathcal{P}_c$ checks if $v_x v_y - v_z \Delta + v_r \stackrel{?}{=} C_1 \Delta + C_0$.

LPZK technique can be optimized in the batched settings. Namely, parties hold $[\boldsymbol{x}]_{\Delta^{(c)}}, [\boldsymbol{y}]_{\Delta^{(c)}}, [\boldsymbol{z}]_{\Delta^{(c)}}$ each of length $m$, and $\mathcal{P}_s$ wants to prove in ZK that $\boldsymbol{z} = \boldsymbol{x} \odot \boldsymbol{y}$. In this case, instead of sending $2m$ coefficients, the $m$ verifications can be batched as follows: $\mathcal{P}_c$ sends a random linear combination, $\mathcal{P}_s$ evaluates the linear combination on the $m$ coefficients $C_1$ and $m$ coefficients $C_0$, and then sends resulting 2 aggregated coefficients. One random IT-MAC $[r]_{\Delta^{(c)}}$ is needed as before for ZK. In this work, this random linear combination is generated as the powers of a random field element, achieving information-theoretical security.

**Generalized LPZK.** As observed by [82], the LPZK technique can be generalized to higher degree polynomials. In this work, we use this generalized technique to improve concrete efficiency in the following setting: $\mathcal{P}_s$ and $\mathcal{P}_c$ hold $[x]_{\Delta^{(c)}} = \langle (x, w_x), v_x \rangle_{\Delta^{(c)}}, [y]_{\Delta^{(c)}} = \langle (y, w_y), v_y \rangle_{\Delta^{(c)}}$, where $\mathcal{P}_s$ wants to prove in ZK to $\mathcal{P}_c$ that $y = x^e$ for some public positive integer $e$. The core of this technique is the following (generalized) equality:

$$\begin{aligned}
&\overbrace{v_x^e - v_y \Delta^{e-1}}^{\text{known by } \mathcal{P}_c} \\
=&(x\Delta + w_x)^e - (y\Delta + w_y)\Delta^{e-1} \\
=&(x^e - y)\Delta^e + (ex^{e-1}w_x - w_y)\Delta^{e-1} + \sum_{i=0}^{e-2} \binom{e}{i} x^i w_x^{e-i} \Delta^i
\end{aligned}$$

where $\binom{e}{i}$ denotes the binomial coefficient "$e$ chooses $i$". Crucially, all coefficients before each $\Delta$ term in the last row are known by $\mathcal{P}_s$. Hence, if we do not need ZK, $\mathcal{P}_s$ can send $e$ coefficients where $\mathcal{P}_c$ checks the equality. Similar to the LPZK, a malicious $\mathcal{P}_s$ can only create a wrong proof with probability $\frac{e}{p}$. ZK can be recovered using $e-1$ random IT-MACs (from VOLE). See [46], [82] for details. The batched optimization can also be applied: with $m$ $e$-th exponential proofs, the total communication cost is $e+1$ $\mathbb{F}_p$ elements.

The generalized LPZK can be used to improve the following task (used by our malicious OPRFs): parties hold $[x]_{\Delta^{(c)}}$ and want to obtain $[x^{2^{128}}]_{\Delta^{(c)}}$. The naïve approach would require $\mathcal{P}_s$ to commit to 128 intermediate results with a batched multiplication LPZK proof—the total communication is 131 $\mathbb{F}_p$ elements. By deploying the generalized LPZK, $\mathcal{P}_s$ can only commit to, e.g., $x^{2^4}, x^{2^8}, x^{2^{12}}, \cdots$, with a batched $2^4$-th exponential relation proof—the total communication decreases to $128/4 + 2^4 + 1 = 49$ $\mathbb{F}_p$ elements. Note that this optimization does not improve communication asymptotically. It also increases the computation concretely.

All the proof techniques shown in this section provide *information-theoretic* security in the VOLE-hybrid model.

## 3. Technical Overview

In this section, we present our half-malicious and malicious 2PC-Gold with sufficient detail to understand our first major contribution. Other Gold OPRFs, formalization, and detailed analysis are presented in subsequent sections.

We focus on efficiently and obliviously evaluating Gold in the VOLE-hybrid model (see Figure 1). Namely, for a public prime $p = 2^\lambda \cdot g + 1$ where $g$ is of $\mathcal{O}(\lambda)$ bits, $\mathcal{P}_s$ holds a PRF key $k \in \mathbb{F}_p$ whereas $\mathcal{P}_c$ holds a PRF input $x \in \mathbb{F}_p$, and the objective is to allow $\mathcal{P}_c$ to learn $(k+x)^g$ over $\mathbb{F}_p$ obliviously, with the help of VOLE correlations over $\mathbb{F}_p$. See Section 4 for the choice of $g$ and how to efficiently convert the output into a $\mathcal{O}(\lambda)$-bit element.

We first address scenarios where $\mathcal{P}_s$ uses a (fresh) uniformly sampled PRF key. Later, we show how to adapt the protocol to accommodate a server-specified key at the cost of $\mathcal{P}_s$ sending a single extra field element.

We note that Section 2.1 includes the used notation.

### 3.1. Overview of Half-Malicious 2PC-Gold

We give an overview of our half-malicious 2PC-Gold, namely, providing security against a malicious client and a semi-honest server. The protocol is depicted in Figure 2.

**Random Root Technique.** Our protocol relies heavily on a simple observation [35], [42]: To evaluate $(k+x)^g$ obliviously, it is sufficient to let $\mathcal{P}_c$ obliviously learn $\alpha^{2^\lambda} \cdot (k+x)$, where $\alpha \in \mathbb{F}_p$ is uniformly sampled and *unknown* to $\mathcal{P}_c$. Then, $\mathcal{P}_c$ can locally compute $\left(\alpha^{2^\lambda} \cdot (k+x)\right)^g$ over $\mathbb{F}_p$. This works because (1) for any $\alpha \in \mathbb{F}_p^*$, $\alpha^{2^\lambda \cdot g} = 1$ over $\mathbb{F}_p$ (which ensures correctness), and (2) $\alpha^{2^\lambda} \cdot (k+x)$ is distributed as a uniformly chosen root of the equation $X^g = (x+k)^g$ over $\mathbb{F}_p$, which can be simulated by choosing such a random root (see Section 4.2).

**Protocol.** Our half-malicious 2PC-Gold (see Figure 2) only requires two messages in the VOLE-hybrid model. Informally, the first message is sent by $\mathcal{P}_c$ and encrypts the

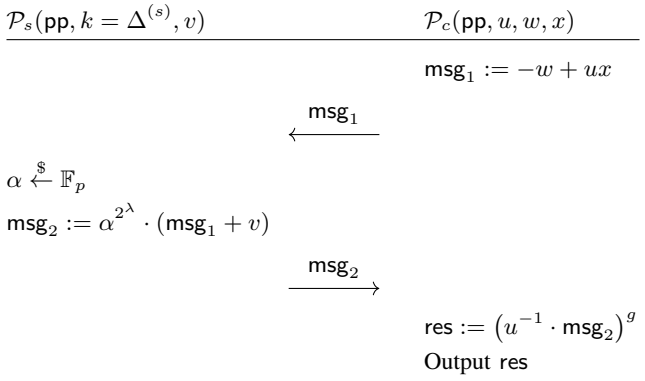


Figure 2: The diagram of our half-malicious 2PC-Gold between $\mathcal{P}_s$ and $\mathcal{P}_c$. $\mathsf{pp} := (1^\lambda, p, g)$ denotes the public parameters. $\mathcal{P}_s$ holds the PRF key $k$ whereas $\mathcal{P}_c$ holds the PRF input $x$. The objective is to allow $\mathcal{P}_c$ obliviously obtain $(k+x)^g$. Recall that $k, v, u, w$ form a single VOLE correlation, i.e., $v = w + uk$.

evaluated input $x$. Then, the second message is sent by $\mathcal{P}_s$ and allows $\mathcal{P}_c$ to recover $\alpha^{2^\lambda} \cdot (k+x)$ over $\mathbb{F}_p$ where $\alpha$ is sampled uniformly by $\mathcal{P}_s$.

In detail, $\mathcal{P}_s$ and $\mathcal{P}_c$ use $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$ (Figure 1) to obtain one VOLE correlation: $\mathcal{P}_s$ holds $\Delta^{(s)}, v$ whereas $\mathcal{P}_c$ holds $u, w$ such that $v = w + u\Delta^{(s)}$. Since $\Delta^{(s)}$ is uniformly sampled, it can be used as the PRF key $k$.[8] We make the following remarks w.r.t. correctness, security, and performance:

- **Correctness:** Recall that $v = w + u\Delta^{(s)} = w + uk$. The correctness relies on the following equality:

$$\begin{aligned} u^{-1} \cdot \mathsf{msg}_2 &= u^{-1} \cdot \alpha^{2^\lambda} \cdot (\mathsf{msg}_1 + v) \\ &= u^{-1} \cdot \alpha^{2^\lambda} \cdot (-w + ux + v) \\ &= u^{-1} \cdot \alpha^{2^\lambda} \cdot (uk + ux) \\ &= \alpha^{2^\lambda} \cdot (k + x) \end{aligned}$$

  Note that $u$ is zero with negligible probability $\frac{1}{p}$ since it is uniformly sampled by $\mathcal{F}_{\mathsf{VOLE}}$.
- **Semi-honest $\mathcal{P}_s$:** This protocol is secure against a semi-honest $\mathcal{P}_s$ in the VOLE-hybrid model. Indeed, consider $\mathsf{msg}_1 + v = u \cdot (k+x)$. If $k + x \neq 0$, since $u$ is uniformly random in the view of $\mathcal{P}_s$, so is $u \cdot (k+x)$. Thus, it can be simulated by a random value in $\mathbb{F}_p$. In case $k + x = 0$, $\mathcal{P}_s$ learns this fact, which we formalize as a 1-bit leakage in our UC treatment (Figure 4). See Section 3.4 for simple and efficient ways to eliminate this leakage.
- **Malicious $\mathcal{P}_c$:** This protocol is secure against a malicious $\mathcal{P}_c$ in the VOLE-hybrid model. Indeed, the only place a malicious $\mathcal{P}_c$ can cheat is by choosing an arbitrary $\widetilde{\mathsf{msg}_1}$. Let us first assume that $u \neq 0$. Then the simulator $\mathcal{S}$ can extract a corresponding input $\widetilde{x}$ from $\widetilde{\mathsf{msg}_1}$ as follows: since $\mathcal{S}$ knows $u$ and $w$ (by emulating $\mathcal{F}_{\mathsf{VOLE}}$), it can compute $\widetilde{x} := (\widetilde{\mathsf{msg}_1} + w) \cdot u^{-1}$. Then, $\mathcal{S}$ gets from the ideal functionality the PRF output $(k+\widetilde{x})^g$. Finally, $\mathcal{S}$ can use the random root technique described above to simulate $\mathsf{msg}_2$ as a random root $X$ of $X^g = (k+\widetilde{x})^g$.
  If $u = 0$, we remark that $v = w$. Thus, $\mathcal{S}$ knows $v$ and can trivially simulate $\mathsf{msg}_2 = \alpha^{2^\lambda} \cdot (\widetilde{\mathsf{msg}_1} + v)$.
- **Performance:** Our half-malicious 2PC-Gold is extremely efficient. It has near-optimal communication: 2 messages, each consisting of a single $\mathbb{F}_p$ element. It consumes a single VOLE correlation per OPRF invocation, and the time complexity of each party is dominated by the time to perform a single exponentiation modulo $p$ (with $\mathcal{O}(\lambda)$-bit exponents). In more detail, $\mathcal{P}_c$ and $\mathcal{P}_s$ also perform one multiplication modulo $p$, and $\mathcal{P}_c$ performs an inversion modulo $p$, which can be done using the standard extended GCD algorithm in time $\tilde{O}(\log^2 p)$.

Note that the security of our protocol in the VOLE-hybrid model is information-theoretic.

**Batched Evaluations.** Our protocol can support $n$-batched evaluations using $n$ VOLE correlations (which can be generated more efficiently as a batch). That is, we can execute the protocol in Figure 2 in $n$ parallel instances: for each $i \in [n]$, $\mathcal{P}_c$ sends $-w_i + u_i x_i$; then $\mathcal{P}_s$ sends $\alpha_i^{2^\lambda} \cdot u_i \cdot (k + x_i)$; $\boldsymbol{v}, \boldsymbol{w}, \boldsymbol{u} \in \mathbb{F}_p^n$ are output by $\mathcal{F}_{\mathsf{VOLE}}$ s.t. $\boldsymbol{v} = \boldsymbol{w} + \boldsymbol{u}k$. As noted in footnote 8, if $\mathcal{P}_s$ wants to adjust the key to a given value $k^*$ instead of using the $\mathcal{F}_{\mathsf{VOLE}}$'s $\Delta^{(s)}$, it suffices to send $k^* - \Delta^{(s)}$ only once. We remark that these $n$ inputs do not need to be different.

### 3.2. Overview of Malicious 2PC-Gold

We give an overview of our malicious 2PC-Gold, which is information-theoretically secure against a malicious $\mathcal{P}_s$ and a malicious $\mathcal{P}_c$ in the VOLE-hybrid model.

Here, we focus on 2PC-Gold with malicious security. This is different from the *public verifiability* property of an OPRF, which is covered in Section 6.1.

**Malicious $\mathcal{P}_s$ Attack Surface.** To see how our malicious 2PC-Gold works, it is instructive to analyze what a malicious $\mathcal{P}_s$ can do in our half-malicious 2PC-Gold. The only place the malicious $\mathcal{P}_s$ can inject an error is in the second message $\mathsf{msg}_2$ sent from $\mathcal{P}_s$ to $\mathcal{P}_c$. As a result, to protect against a malicious $\mathcal{P}_s$, it suffices to require $\mathcal{P}_s$ to prove in ZK that $\mathsf{msg}_2$ is generated correctly. That is,

1) $\mathcal{P}_s$ adds the value $v$, output by $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$, to $\mathsf{msg}_1$, yielding the intermediate result $\mathsf{int} = \mathsf{msg}_1 + v$.
2) $\mathcal{P}_s$ multiplies $\mathsf{int}$ with $\alpha^{2^\lambda}$ to generate $\mathsf{msg}_2 = \mathsf{int} \cdot \alpha^{2^\lambda}$ for some $\alpha \in \mathbb{F}_p$.

Even with such zero-knowledge proof, we note that a malicious $\mathcal{P}_s$ is still able to learn whether $k + x$ equals zero. This is the 1-bit leakage we mentioned in Section 3.1. See Section 3.4 for details on how to eliminate it efficiently.

**Deploying ZK.** Malicious 2PC-Gold exploits the VOLE-based ZK techniques (see Section 2.5) to ensure a well-formed $\mathsf{msg}_2$. Namely, $\mathcal{P}_s$ and $\mathcal{P}_c$ will use another VOLE correlation functionality, *with reversed roles* as $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$, in

8. In applications where $\mathcal{P}_s$ inputs an arbitrary key $k^*$ (e.g., a key committed earlier), $\mathcal{P}_s$ will send $k^* - \Delta^{(s)}$, allowing $\mathcal{P}_c$ to adjust the correlation as $v = w + uk^*$. That is, $\mathcal{P}_c$ sets $w := w - (k^* - \Delta^{(s)}) \cdot u$.

addition to the correlations required by the half-malicious 2PC-Gold from Figure 2. These new correlations can be viewed as IT-MAC commitments from $\mathcal{P}_s$ to $\mathcal{P}_c$ over a $\mathcal{P}_s$-known random element $r \in \mathbb{F}_p$, denoted as $[r]_{\Delta^{(c)}}$. Note, $\Delta^{(c)}$ is known by $\mathcal{P}_c$ and can be viewed as the verifier private coins in the ZK proof. (Instead, $\Delta^{(s)}$ in our half-malicious protocol is used as the OPRF key.) If parties can get $[v]_{\Delta^{(c)}}$ and $[\alpha]_{\Delta^{(c)}}$, since $\mathsf{msg}_1$ and $\mathsf{msg}_2$ are public, parties can directly use a VOLE-based ZK to prove that $\mathsf{msg}_2$ is equal to $\alpha^{2^\lambda} \cdot (\mathsf{msg}_1 + v)$. In other words, with $[v]_{\Delta^{(c)}}$, $[\alpha]_{\Delta^{(c)}}$ and $\mathsf{msg}_1$, parties can generate $[\mathsf{msg}_2 = \alpha^{2^\lambda} \cdot (\mathsf{msg}_1 + v)]_{\Delta^{(c)}}$; then $\mathcal{P}_s$ can bindingly open $\mathsf{msg}_2$ to $\mathcal{P}_c$. While this blueprint is relatively straightforward, there are several challenges to tackle:

1) $[v]_{\Delta^{(c)}}$ needs to be generated correctly, even with a malicious $\mathcal{P}_s$. Note, $v$ is not an arbitrary value but rather is part of the VOLE correlation from $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$, used by our half-malicious protocol. That is, we not only need $\mathcal{P}_s$ to commit to $v$ but to commit to a consistent $v$. More importantly, recall that $v = w + uk$ where $u$ will be used as a one-time pad during the OPRF evaluation; hence, we must generate $[v]_{\Delta^{(c)}}$ without revealing any information correlated to $u$.
2) $[\alpha]_{\Delta^{(c)}}$ needs to be generated. This is easy as $\alpha$ is only used to protect $\mathcal{P}_s$'s OPRF key. In fact, $[\alpha]_{\Delta^{(c)}}$ can be directly generated as one correlation from $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$ since the committed value $\alpha$ is pseudorandom.
3) Generating $[\alpha^{2^\lambda}]_{\Delta^{(c)}}$ from $[\alpha]_{\Delta^{(c)}}$ requires $\lambda + 3$ field elements of communication: $\lambda$ of them are used to commit to the intermediate results (i.e., $[\alpha^{2^1}]_{\Delta^{(c)}}$, $[\alpha^{2^2}]_{\Delta^{(c)}}$, ..., $[\alpha^{2^\lambda}]_{\Delta^{(c)}}$), and 3 of them are used to deploy the batched LPZK technique (see Section 2.5) to ensure each squaring is computed correctly.
   We could further reduce this cost via the generalized LPZK technique (see Section 2.5). Namely, w.l.o.g., for any constant $\phi$ that divides $\lambda$, $\mathcal{P}_s$ commits to $\frac{\lambda}{\phi}$ intermediate results as $[\alpha^{2^\phi}]_{\Delta^{(c)}}$, $[\alpha^{2^{2\phi}}]_{\Delta^{(c)}}$, ..., $[\alpha^{2^\lambda}]_{\Delta^{(c)}}$ and proves in ZK that each $2^\phi$-powering is computed correctly. Now, the communication cost to generate $[\alpha^{2^\lambda}]_{\Delta^{(c)}}$ is reduced to $\frac{\lambda}{\phi} + 2^\phi + 1$ field elements.

**Generating Consistent $[v]_{\Delta^{(c)}}$.** We now show how to tackle Challenge 1. Recall that $v = w + uk$ where $v, k$ are known by $\mathcal{P}_s$ and $u, w$ are known by $\mathcal{P}_c$. We aim to generate $[v]_{\Delta^{(c)}}$. We first present an ineffective yet insightful approach. $\mathcal{P}_s$ commits to $[v]_{\Delta^{(c)}}$ and $[k]_{\Delta^{(c)}}$; then $\mathcal{P}_c$ reveals $u$ and $w = v - uk$ and requires $\mathcal{P}_s$ to show that $[v]_{\Delta^{(c)}} - w - u[k]_{\Delta^{(c)}} = [v - w - uk]_{\Delta^{(c)}}$ commits to zero. We now argue why this check ensures a correct $[v]_{\Delta^{(c)}}$ and $[k]_{\Delta^{(c)}}$. Assume $\mathcal{P}_s$ committed to $[\widetilde{v} \neq v]_{\Delta^{(c)}}$ or $[\widetilde{k} \neq k]_{\Delta^{(c)}}$. Then, if the IT-MAC $[\widetilde{v}]_{\Delta^{(c)}} - w - u[\widetilde{k}]_{\Delta^{(c)}}$ commits to zero, based on the binding property of IT-MAC, we know w.h.p. $\widetilde{v} - w - u\widetilde{k} = 0$. This implies $\widetilde{v} - u\widetilde{k} = v - uk$. If $k = \widetilde{k}$, then $v = \widetilde{v}$. Otherwise, we have $u = (v - \widetilde{v})/(k - \widetilde{k})$. This means that $\mathcal{P}_s$ can successfully guess $u$ *before* $u$ is revealed. Note that since $u$ is uniformly sampled when $\mathcal{P}_c$ is honest by $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$ (see Figure 1), this can only happen with probability $\frac{1}{p}$. However, although this approach guarantees a well-formed $[v]_{\Delta^{(c)}}$, it is ineffective. This is because $u$ is revealed and, crucially, is later used as a one-time pad.

Therefore, we must ensure a well-formed $[v]_{\Delta^{(c)}}$ without revealing any information correlated to $u$ (and $w$). We notice that this can be done by exploiting a random linear combination to "sacrifice" another $u_{\mathsf{sa}}$, which will be discarded immediately after the check to ensure $u$ remains full entropy. In detail, let parties hold another VOLE correlation from $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$: $v_{\mathsf{sa}} = w_{\mathsf{sa}} + u_{\mathsf{sa}}k$ where $v_{\mathsf{sa}}, k$ are known by $\mathcal{P}_s$ and $u_{\mathsf{sa}}, w_{\mathsf{sa}}$ are known by $\mathcal{P}_c$. Now, let $\mathcal{P}_s$ commit to $[v_{\mathsf{sa}}]_{\Delta^c}$, $[v]_{\Delta^c}$ and $[k]_{\Delta^c}$. Then, $\mathcal{P}_c$ samples a uniform $\chi \xleftarrow{\$} \mathbb{F}_p$ and reveals $\chi, \chi u + u_{\mathsf{sa}}$ and $\chi w + w_{\mathsf{sa}}$. Finally, $\mathcal{P}_c$ requires $\mathcal{P}_s$ to show that the following IT-MAC

$$\begin{aligned}&\chi[v]_{\Delta^{(c)}} + [v_{\mathsf{sa}}]_{\Delta^{(c)}} - (\chi w + w_{\mathsf{sa}}) - (\chi u + u_{\mathsf{sa}})[k]_{\Delta^{(c)}}\\ =&[\chi(v - w - uk) + (v_{\mathsf{sa}} - w_{\mathsf{sa}} - u_{\mathsf{sa}}k)]_{\Delta^{(c)}}\end{aligned}$$

commits to zero. Note, this can be viewed as a consistency check over the value $\chi v + v_{\mathsf{sa}}$. That is, it ensures that $\chi\widetilde{v} + \widetilde{v_{\mathsf{sa}}} = \chi v + v_{\mathsf{sa}}$ and $\widetilde{k} = k$, where $\widetilde{v}, \widetilde{v_{\mathsf{sa}}}, \widetilde{k}$ are arbitrary (potentially inconsistent) values committed by $\mathcal{P}_s$. Moreover, since $\chi$ is sampled and revealed only after $\widetilde{v}, \widetilde{v_{\mathsf{sa}}}$ are chosen, $\chi\widetilde{v} + \widetilde{v_{\mathsf{sa}}} = \chi v + v_{\mathsf{sa}}$ implies $\widetilde{v} = v$ and $\widetilde{v_{\mathsf{sa}}} = v_{\mathsf{sa}}$ with overwhelming probability based on the well-known *Schwartz-Zippel* lemma [73], [86]. Crucially, the values revealed by $\mathcal{P}_c$ (i.e., $\chi, \chi u + u_{\mathsf{sa}}$ and $\chi w + w_{\mathsf{sa}}$) are *independent* of $u$, so $u$ remains full entropy.

One essential remark is that we also need to ensure a well-formed $\chi, \chi u + u_{\mathsf{sa}}$ and $\chi w + w_{\mathsf{sa}}$ for a malicious $\mathcal{P}_c$. These new messages also provide a new opportunity for a malicious $\mathcal{P}_c$ to exploit—indeed, a malicious $\mathcal{P}_c$ can learn the OPRF key with ill-formed messages. Interestingly, this can be resolved as follows: recall that $u$ and $u_{\mathsf{sa}}$ can also be viewed as committed values inside IT-MAC but rather from $\mathcal{P}_c$ to $\mathcal{P}_s$; thus, $\chi w + w_{\mathsf{sa}}$ can be used as a proof to force correct opening of $\chi u + u_{\mathsf{sa}}$. That is, an honest $\mathcal{P}_s$ can check if $\chi v + v_{\mathsf{sa}}$ is equal to $\chi w + w_{\mathsf{sa}} + (\chi u + u_{\mathsf{sa}})k$.

We emphasize that this check relies on $k$ with enough entropy *to the UC environment*. Hence, when $k$ is input from $\mathcal{P}_s$, this does not work. Instead, parties can perform the consistency check with the committed $\Delta^{(s)}$ rather than $k$ (see Footnote 8 and Section 5.3). However, note that $\Delta^{(s)}$ will also be revealed to the UC environment (1) as $\mathcal{P}_s$'s output when the OPRF uses a uniformly sampled key, or (2) from $k^* - \Delta^{(s)}$ when the OPRF uses a server-specified key. Therefore, the step of generating consistent $[v]_{\Delta^{(c)}}$ must take place before the above step (1) or (2) is performed. This can be easily enforced if parties only need to execute the OPRF once but it requires additional care in other cases, e.g., batched offline phases presented in Section 3.3.

Finally, we remark that $[k]_{\Delta^{(c)}}$ is an IT-MAC of the OPRF key and thus can be used to prove any NP relation over it using regular VOLE-based ZK. For example, it can be used to prove in ZK that the server reuses the same committed OPRF key across different invocations as how we do to achieve verifiable OPRFs in Section 6.1.

**Batched Evaluations.** Our malicious protocol can be extended to the batched variant naturally using more VOLE correlations in both directions (i.e., $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$ and $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$). For $n$-batched evaluations, the step to generate IT-MACs of $[v_1]_{\Delta^{(c)}}, \ldots, [v_n]_{\Delta^{(c)}}$ can be improved (i.e., addressing Challenge 1) by "sacrificing" a single correlation. Namely, after $\mathcal{P}_s$ commits to $\boldsymbol{v}$, $\mathcal{P}_c$ can reveal $\chi \xleftarrow{\$} \mathbb{F}_p$, $\left(\sum_{i=1}^{n} \chi^i u_i\right) + u_{\mathsf{sa}}$ and $\left(\sum_{i=1}^{n} \chi^i w_i\right) + w_{\mathsf{sa}}$, reducing the communication of this step from $6n+1$ to $n+6$ field elements. Note that the $\Delta^{(s)}$ must remain full entropy to the UC environment at this point. Furthermore, the batched (generalized) LPZK can be applied across $n$ evaluations to generate $[\alpha^{2^\lambda}_{i\in[n]}]_{\Delta^{(c)}}$ and prove the correct $\mathsf{msg}_{2,i}$.

### 3.3. Offline-Online Mode

Our protocols support an *offline-online* mode where most portions of computation and communication can be pushed into an *input-independent* (and PRF-key-*independent*) offline (aka preprocessing) phase. Clearly, the generation of (V)OLE correlations is a hybrid input-independent functionality, so it can be pushed into the offline phase. This is true even when $\mathcal{P}_s$ wants to select its own OPRF key $k^*$, since this can be done via sending $k^* - \Delta^{(s)}$ in the online phase.

Moreover, for our malicious protocol, the parts generating $[v]_{\Delta^{(c)}}$ and $[\alpha^{2^\lambda}]_{\Delta^{(c)}}$ can also be pushed into the offline phase. Again, this is true even when the $\mathcal{P}_s$ wants to select its own OPRF key $k^*$. Crucially, this does *not* affect the ability to obtain the committed key. That is, the parties in the offline phase would obtain $[\Delta^{(s)}]_{\Delta^{(c)}}$, which can be locally adjusted to $[k^*]_{\Delta^{(c)}}$ via $k^* - \Delta^{(s)}$. The online phase cost of our malicious protocol is very lean compared to the online phase cost of our half-malicious protocol: it only needs 3 more field elements (i.e., a batched LPZK), regardless of $n$.

In all, our protocols are very efficient in the amortized offline-online setting and reasonably efficient even *with* offline cost included. See Section 7.

**Batched Offline Phases.** Our protocols support batched offline phases. That is, offline phases of $n_{\mathsf{batch}}$ evaluations can be prepared together for better amortized costs. Then, each $n$-sub-batched evaluation—where various values of $n$ add up to $n_{\mathsf{batch}}$—can be processed immediately without needing the entire $n_{\mathsf{batch}}$ batch, paying an $n$-amortized online phase. Note that once the online phase starts, the UC environment learns $\Delta^{(s)}$, making the consistency check in our malicious 2PC-Gold no longer UC-secure; see Section 3.2. Thus, in our malicious OPRFs, $n_{\mathsf{batch}}$ must be predetermined, and the batched offline phases must be completed in advance. We formalize batched offline and sub-batched online phases as our ideal functionality in Section 5.1.

In certain applications, fixing $n_{\mathsf{batch}}$ in advance is either impractical or impossible. In Section 6.1, we demonstrate how to remove this limitation easily via *private verifiability*.

### 3.4. Eliminating 1-bit Leakage in Our Protocols

The protocols presented in Sections 3.1 and 3.2 incur a 1-bit leakage per evaluation. In particular, for each evaluation, $\mathcal{P}_s$ learns $u \cdot (k+x)$, which allows her to determine whether $k+x$ equals zero.

Namely, our ideal functionality to capture these protocols has to leak this 1-bit information to the simulator if $\mathcal{P}_s$ is corrupted. Note, if we consider UC security, this 1-bit leakage is unavoidable even for a semi-honest $\mathcal{P}_s$ with a uniformly sampled key. This is because the environment can always set the honest $\mathcal{P}_c$'s input as $-k$.

While this leakage may not be a problem for many applications (in particular, this leakage also exists in Dodis-Yampolskiy PRF [31] and some corresponding OPRFs), we now demonstrate how to eliminate this leakage, depending on whether we need an OPRF with a uniformly sampled key or one specified by $\mathcal{P}_s$.

**Uniformly Sampled Key.** If only a uniformly sampled PRF key is needed, we first show how to enforce a malicious $\mathcal{P}_s$ to use a uniformly sampled key. Note, in $\mathcal{F}_{\mathsf{VOLE}}$ (see Figure 1), the malicious receiver (i.e., $\mathcal{P}_s$) is allowed to select an arbitrary scalar $\Delta^{(s)}$ (i.e., an arbitrary PRF key). Our key observation is that this scalar $\Delta^{(s)}$ can be randomized by $\mathcal{P}_c$. That is, $\mathcal{P}_c$ can send a uniformly sampled $\Delta' \xleftarrow{\$} \mathbb{F}_p$ to set the scalar as $\Delta^{(s)} + \Delta'$—which will be used as the PRF key that is guaranteed to be uniformly sampled—and adjust his VOLE shares locally.

By enforcing a uniformly sampled key $k$, this leakage can be naturally eliminated in the standalone setting. However, in the UC setting, the leakage persists because the environment can always set $x$ as $-k$ after learning $k$. To further eliminate it in the UC setting, we can let $\mathcal{P}_c$ hash the input before executing the OPRF protocol(s). Intuitively, this prevents the leakage as the environment cannot find an $x$ s.t. $\mathsf{H}(x) = -k$ when $\mathsf{H}(\cdot)$ is modeled as a random oracle.

**Server-Specified Key.** If we need an OPRF where the key is fully specified by $\mathcal{P}_s$, the previously mentioned fix does not apply. Instead, we can eliminate the leakage by restricting the PRF key to be chosen from elements in $\mathbb{F}_p$ that have two leading zero bits (i.e., the two most significant bits are 0s) and by limiting the PRF input to the subset of $\mathbb{F}_p$ consisting of all non-zero elements with two leading zero bits. In essence, this ensures that $k+x$ is never zero.

If $\mathcal{P}_s$ is malicious, we can further require $\mathcal{P}_s$ to provide a VOLE-based ZK proof over IT-MAC $[k^*]_{\Delta^{(c)}}$ to show it has two leading zero bits. This proof requires communicating $\mathcal{O}(\lambda)$ field elements and performing $\mathcal{O}(\lambda)$ field operations. It can be amortized in batched evaluations.

Note that this fix requires a minor modification of the underlying PRF assumption to sample the key accordingly; see Section 4.1.

## 4. Gold PRF Basics

### 4.1. Formal Hardness Assumptions

We define the underlying computation problem—the Decisional Shifted Power-Residue Symbol—as follows:

**Definition 1** (Decisional Shifted Power-Residue Symbol (DSPRS) Problem)**.** *Let $p = p(\lambda)$ be a family of prime*

*numbers, where each $p = e(\lambda) \cdot g(\lambda) + 1$. Let $k$ be uniformly sampled from $\mathbb{F}_p$. Let $\mathcal{D} := \{a^g \bmod p \mid a \in \mathbb{F}_p^*\}$. Let $O_{\mathsf{PR}}$ be an oracle that on input $x \in \mathbb{F}_p$ outputs $(x+k)^g \bmod p$, and $O_{\mathsf{R}}$ be a random oracle that maps elements from $\mathbb{F}_p$ to $\mathcal{D}$. The DSPRS problem asks to distinguish between $O_{\mathsf{PR}}$ and $O_{\mathsf{R}}$ given $1^\lambda$ and $p, g$, with classical queries.*

The post-quantum security of our Gold PRF relies on the following hardness assumption:

**Assumption 1.** *For any $p = p(\lambda) = e(\lambda) \cdot g(\lambda) + 1$, where $e(\lambda)$ is smooth and $g(\lambda) = \Omega(\lambda)$, there is no probabilistic polynomial time (quantum) algorithm for the DSPRS problem (Definition 1) with non-negligible advantage in $\lambda$.*

In this work, we set $e = 2^\lambda$ and $\log g = 2\lambda + \mathcal{O}(1)$. Gold is essentially $O_{\mathsf{PR}}$ in Definition 1. Existing attacks presented in Section 2.3 guide this choice. These parameters are very conservative—the best-known quantum attack for Gold by [12] has time complexity $\mathcal{O}(\frac{g}{L})$ using $L \leq p^{1/4}$ queries; the best-known quantum attack for the Legendre PRF over $p$ (which also applies for Gold) by [38] has time complexity $2^{\mathcal{O}(\log p)} \cdot p^{1/3}$ using $p^{1/3}$ queries.

Finally, we note that a minor modification of Definition 1 is needed to eliminate the one-bit leakage when the OPRF key is chosen by $\mathcal{P}_s$ (see Section 3.4). Namely, we only need to sample the PRF key from the subspace of $\mathbb{F}_p$ consisting of all elements with two leading zero bits.

### 4.2. Power-Residue Subgroup

Let $p = 2^\lambda \cdot g + 1$ be a prime of $\mathcal{O}(\lambda)$ bits. We review several useful properties of the $g$-th power-residue subgroup $\mathbb{G}$ of the multiplicative group $\mathbb{F}_p^*$ of $\mathbb{F}_p$. That is,

$$\mathbb{G} := \{x^g \bmod p \mid x \in \mathbb{F}_p^*\}$$

where the group operation is multiplication modulo $p$.

**Fact 1.** *$\mathbb{G}$ is a finite cyclic group of order $2^\lambda$. Let $h$ be a generator of $\mathbb{G}$. Then, for any element $x \in \mathbb{G}$, the Pohlig-Hellman algorithm [68] can solve the discrete logarithm (DLOG) of $x$ in base $h$ in $O(\lambda^2)$ group operations.*

**Fact 2.** *There exist $2^{\lambda-1}$ generators in $\mathbb{G}$. In particular, there exists an efficient way to find a generator $h$: repeatedly sample $r \xleftarrow{\$} \mathbb{F}_p^*$, and set $h := r^g$, until $h^{2^{\lambda-1}} \neq 1$.*

**Converting Output.** Fact 1 implies a straightforward deterministic way to transfer the output of 2PC-Gold from $\mathbb{G}$ to $\{0,1\}^\lambda$. I.e., $\mathcal{P}_c$ can locally solve the DLOG of 2PC-Gold output to a publicly agreed base $h$. Note that when $k+x = 0$, $\mathcal{P}_c$ will abort, so there is no need to solve the DLOG. Of course, there exist other standard ways to post-process output, such as applying a hash function over it, in particular, if more than $\lambda$ bits of output are required (e.g., O-Gold).

**Solving Equation $X^g \equiv a \pmod p$.** Facts 1 and 2 can be exploited to design an efficient algorithm (defined in Figure 3) to find a solution $X$ of the equation $X^g \equiv a \pmod p$ for some $a \in \mathbb{G}$. This algorithm is essential for arguing the security of our protocols: the simulator uses it.

SolveEq$(1^\lambda, p, g, a)$

1 : Apply Fact 2 to get a generator $h = r^g$ of $\mathbb{G}$
2 : Apply Fact 1 to get $z$ such that $h^z = a$
3 : **return** $r^z$

Figure 3: An efficient algorithm to find a solution $X$ of the equation $X^g \equiv a \pmod p$ where $a \in \mathbb{G}$, $\mathbb{G} := \{x^g \bmod p \mid x \in \mathbb{F}_p^*\}$, and $p = 2^\lambda \cdot g + 1$ is a prime of length $\mathcal{O}(\lambda)$. Correctness comes from: $h^z = (r^g)^z = (r^z)^g = a$.

**Functionality $\mathcal{F}_{\mathsf{2PC\text{-}Gold}}$**

$\mathcal{F}_{\mathsf{2PC\text{-}Gold}}$, parameterized by a field $\mathbb{F}_p$ where (1) $p = p(\lambda) = 2^\lambda \cdot g + 1$ is a prime and (2) $g = g(\lambda)$ is an integer of $2\lambda + O(1)$ bits, running with a server $\mathcal{P}_s$, a client $\mathcal{P}_c$ and an adversary $\mathcal{S}$, proceeds as follows:

**Initialize.** Upon receiving $(\texttt{init}, n_{\mathsf{batch}})$ from $\mathcal{P}_s$ and $\mathcal{P}_c$ where $n_{\mathsf{batch}} \in \mathbb{Z}^+$, if $\mathcal{P}_s$ is honest (or semi-honest), sample $k \xleftarrow{\$} \mathbb{F}_p$, else receive $k$ from $\mathcal{S}$. Store $n_{\mathsf{batch}}$ and $k$. Send $k$ to $\mathcal{P}_s$. Set $n_{\mathsf{eval}} := 0$. Ignore subsequent $(\texttt{init}, \cdot)$.

**Evaluate.** Upon receiving $(\texttt{eval}, n)$ from $\mathcal{P}_s$ and $(\texttt{eval}, n, x_1, x_2, \ldots, x_n)$ from $\mathcal{P}_c$, where $n \in \mathbb{Z}^+$ and each $x_{i\in[n]} \in \mathbb{F}_p$. If $n_{\mathsf{eval}} + n > n_{\mathsf{batch}}$, ignore the instruction; otherwise, set $n_{\mathsf{eval}} := n_{\mathsf{eval}} + n$ and proceed as follows:

1) For each $i \in [n]$, compute $y_i := (k + x_i)^g$ over $\mathbb{F}_p$.
2) (**One-bit leakage**) If $\mathcal{P}_s$ is corrupted, for each $i \in [n]$, compute $\ell_i := (k + x_i) \stackrel{?}{=} 0$, send $(\texttt{eval}, \boldsymbol{\ell})$ to $\mathcal{S}$. If $\exists i \in [n]$ such that $\ell_i$ is 1, send $(\texttt{abort})$ to $\mathcal{P}_c$ and halt.
3) Send $(\texttt{eval}, \boldsymbol{y})$ to $\mathcal{P}_c$.

Figure 4: The 2PC-Gold functionality.

## 5. Formalization and Analysis of 2PC-Gold

We formalize our 2PC-Gold protocols in the UC framework. Our other OPRFs and their corresponding ideal functionalities, theorems, and proofs can be derived through straightforward modifications or, in the case of UC-Gold in section 6.1, via the elegant transformation from [13].

### 5.1. Ideal Functionality for 2PC-Gold

The ideal functionality $\mathcal{F}_{\mathsf{2PC\text{-}Gold}}$ achieved by our 2PC-Gold is defined in Figure 4. For simplicity, we assume the PRF key is uniformly sampled when $\mathcal{P}_s$ is honest. We remark that our protocols can support a server-specified key almost for free (see Section 5.3). We directly define $\mathcal{F}_{\mathsf{2PC\text{-}Gold}}$ (and realizations) with batched offline phases (i.e., `init` instruction) and sub-batched online phases (i.e., `eval` instruction). The non-batched evaluation setting is the particular case where $n_{\mathsf{batch}} = n = 1$.

Note that `init` can be called only once per session with a predetermined $n_{\mathsf{batch}}$, capturing the maximum number of evaluations (if more are needed `init` needs to be called again). See Section 3.3 for the reason behind this restriction and Section 6.1 for how to remove it easily. Our

half-malicious 2PC-Gold in fact does not even have this restriction. Specifically, Step 3 in our protocol in Figure 5 can be executed repeatedly, even after the online phase has been performed. However, for simplicity, since our plain malicious protocol has this restriction on `init` (specifically Step 5 cannot be executed again after `init` ends), we adopt a unified $\mathcal{F}_{\text{2PC-Gold}}$ that restricts `init` in all settings.

**One-Bit Leakage.** Recall that protocols presented in Section 3 incur a one-bit leakage per evaluation. We show how to remove this leakage efficiently in Section 3.4. Since this functionality (with the one-bit leakage in Step 2) is already useful for many applications (justified below and in Section 3.4), we stick to this definition. Essentially, the changes to upgrade the functionality/protocols/proofs to the version without the one-bit leakage are inexpensive.

**One-Bit Leakage OPRF Applications.** For some applications, the one-bit leakage is harmless, for example:

1) **PSI, where $\mathcal{P}_s$ is semi-honest:** In the case of OPRF-based PSI where the server is semi-honest, the server computes an OPRF with the client over the client's set elements. Since the server uses a uniformly random key, it is unlikely to gain additional information via collision. Indeed, a recent PSI work [77] bases its protocols on an OPRF (not Gold) that features the same type of leakage, indicating that our OPRF with one-bit leakage can directly instantiate their PSI protocols.
2) **Password-based protocols, where $\mathcal{P}_s$ is malicious:** In the case of password-based protocols from OPRF (e.g., OPAQUE [53]), the client computes OPRF over its password and then stores certain public functions of this value on the server. Since $\mathcal{P}_s$ can simply test the client's password locally, the one-bit leakage is harmless.

### 5.2. Formal 2PC-Gold Protocols and Theorems

We defer the reader to Section 3 for concise overviews of 2PC-Gold, including intuitive arguments regarding security. In this section, we formalize our protocols as $\Pi_{\text{2PC-Gold}}$ in Figure 5 with corresponding theorems.

**Theorem 1** (Half-Malicious)**.** *Protocol $\Pi_{\text{2PC-Gold}}$ (Figure 5 without gray boxes) information-theoretically UC-realizes $\mathcal{F}_{\text{2PC-Gold}}$ (Figure 4) in the $\mathcal{F}_{\text{VOLE}}^{c,s}$-hybrid model (Figure 1), in the presence of static corruptions, where a corrupt $\mathcal{P}_s$ can be semi-honest and a corrupt $\mathcal{P}_c$ can be malicious.*

**Theorem 2** (Malicious)**.** *Protocol $\Pi_{\text{2PC-Gold}}$ (Figure 5 with gray boxes) information-theoretically UC-realizes $\mathcal{F}_{\text{2PC-Gold}}$ (Figure 4) in the $(\mathcal{F}_{\text{VOLE}}^{c,s}, \mathcal{F}_{\text{VOLE}}^{s,c})$-hybrid model (Figure 1), in the presence of static corruptions, where a corrupt $\mathcal{P}_s$ or $\mathcal{P}_c$ can be malicious.*

Due to space constraints, we defer proofs to our full version [83]. We note that Section 3 includes intuitive arguments regarding security, i.e., how to construct simulators.

### 5.3. Some Details

**Costs.** We list costs of $\Pi_{\text{2PC-Gold}}$ in the VOLE-hybrid model when $n_{\text{batch}} = n$ (details deferred to the full version [83]):

- **Half-malicious** $\Pi_{\text{2PC-Gold}}$ uses $n$ VOLE correlations with $\mathcal{O}(n\lambda)$ (resp. $\mathcal{O}(n)$) field operations in offline (resp. online) phase for $\mathcal{P}_s$, $\mathcal{O}(n\lambda)$ field operations in online phase for $\mathcal{P}_c$, $2n$ field elements (in online phase) and 2 rounds for communication.
- **Malicious** $\Pi_{\text{2PC-Gold}}$ uses $n\lambda+3n+5$ VOLE correlations with $\mathcal{O}(n\lambda)$ (resp. $\mathcal{O}(n)$) field operations in offline (resp. online) phase for $\mathcal{P}_s$, $\mathcal{O}(n\lambda)$ field operations in both offline and online phases for $\mathcal{P}_c$, $n\lambda+n+9$ (resp. $2n+3$) field elements in offline (resp. online) phase and 5 rounds for communication.

**Optimization via Generalized LPZK.** The communication bottleneck of our malicious protocol lies in Sub-step 4c where $\mathcal{P}_s$ needs to commit to $n_{\text{batch}} \cdot \lambda$ field elements representing each intermediate result of computing $\alpha_{i\in[n_{\text{batch}}]}^{2^\lambda}$. We can exploit the generalized LPZK technique [82] (see Section 2.5) to optimize this step. In particular, w.l.o.g., let $\phi$ be some constant that divides $\lambda$. We can raise each element to the power of $2^\phi$ rather than 2. This will reduce the communication cost of this step to $\frac{n_{\text{batch}}\cdot\lambda}{\phi}$. Note that this will also increase the coefficients $\mathcal{P}_s$ needs to send in Sub-step 4d to $2^\phi$. Hence, $\phi$ can only be a small constant. In conclusion, for any constant $\phi$ divides $\lambda$, we can improve the offline communication cost to $\frac{n_{\text{batch}}\cdot\lambda}{\phi} + n_{\text{batch}} + 7 + 2^\phi$ field elements, with reduced VOLE correlations required.

Using the generalized LPZK will also increase the computation *concretely*—the hidden constant will increase by a factor of $\approx 2^\phi$ in asymptotic. Moreover, while it saves the required VOLE correlations for committing to the intermediate results by $\phi\times$, it requires $2^\phi - 2$ more VOLE correlations (see [46]) to randomize the coefficients sent in ZKP. Finally, the statistical advantage needs to be adjusted.

**Key Selection.** The protocols in Figure 5 assume a uniformly sampled PRF key $k := \Delta^{(s)}$. In some applications, $\mathcal{P}_s$ may instead want to specify this key $k^*$. This can be done by $\mathcal{P}_s$ sending $k^* - \Delta^{(s)}$ at the very beginning of the online phase, followed by $\mathcal{P}_c$ adjusting his $\boldsymbol{w}$. Crucially, $k^* - \Delta^{(s)}$ is one-time padded by uniformly sampled $\Delta^{(s)}$ *and* the simulator can trivially extract $k^*$.

**ZKP over the Committed Key.** In our malicious protocol, Sub-step 5a generates $[\Delta^{(s)}]_{\Delta^{(c)}}$ and Step 5 ensures the unforgability of $\Delta^{(s)}$. This is an IT-MAC over the PRF key. (Note, this is true even in the case where $\mathcal{P}_s$ selects her own PRF key since $[k^*]_{\Delta^{(c)}} = [\Delta^{(s)}]_{\Delta^{(c)}} + (k^* - \Delta^{(s)})$.) Thus, $\mathcal{P}_s$ can execute the standard VOLE-based ZKP using this IT-MAC to demonstrate any properties over the used PRF key. For example, $\mathcal{P}_s$ can prove (1) the leading two bits of the key are zero to get rid of the one-bit leakage in $\mathcal{F}_{\text{2PC-Gold}}$ (Figure 4) as discussed in Section 3.4; and (2) the key binds to a public verification key as discussed in Section 6.1.

## 6. Verifiability and Strong UC Security

### 6.1. Verifiability: Ensuring a Consistent Key

In this section, we discuss how to add different forms of verifiability to our protocols, which ensures the same key is

**Protocol $\Pi_{\mathsf{2PC\text{-}Gold}}$**

$\Pi_{\mathsf{2PC\text{-}Gold}}$, parameterized by a field $\mathbb{F}_p$ where (1) $p = p(\lambda) = 2^\lambda \cdot g + 1$ is a prime and (2) $g = g(\lambda)$ is an integer of $2\lambda + O(1)$ bits, running with a server $\mathcal{P}_s$, a client $\mathcal{P}_c$, with hybrid access to $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$ and $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$ over $\mathbb{F}_p$, proceeds as follows:

**Offline Phase**

**Initialize.** $\mathcal{P}_s$ and $\mathcal{P}_c$, each on receiving $(\texttt{init}, n_{\mathsf{batch}})$ where $n_{\mathsf{batch}} \in \mathbb{Z}^+$, proceed as follows:

1) **Sample OPRF key:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{init})$ to $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$, where $\mathcal{P}_s$ receives the uniformly sampled OPRF key $k := \Delta^{(s)} \in \mathbb{F}_p$.
2) **Sample ZKP private coins:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{init})$ to $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$, where $\mathcal{P}_c$ receives the uniformly sampled $\Delta^{(c)} \in \mathbb{F}_p$.
3) **Sample VOLE correlations used for OPRF evaluations:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{extend}, n_{\mathsf{batch}})$ to $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$, and then $\mathcal{P}_s$ receives $\boldsymbol{v}$ whereas $\mathcal{P}_c$ receives $\boldsymbol{u}, \boldsymbol{w}$ such that $\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{w} \in \mathbb{F}_p^{n_{\mathsf{batch}}}, \boldsymbol{v} = \boldsymbol{w} + \boldsymbol{u}\Delta^{(s)}$.
4) **Prepare the committed $g$-th residues:**
   a) **Sample the hiding pads:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{extend}, n_{\mathsf{batch}})$ to $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$, and then parties receive length-$n_{\mathsf{batch}}$ random IT-MACs as $[\boldsymbol{\alpha}]_{\Delta^{(c)}}$. (Note: when $\mathcal{P}_s$ is honest, $\boldsymbol{\alpha}$ looks uniform to $\mathcal{P}_c$.)
   b) **Sample the VOLE correlations used for raising powers:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{extend}, n_{\mathsf{batch}} \cdot \lambda)$ to $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$, and then parties receive length-$(n_{\mathsf{batch}} \cdot \lambda)$ random IT-MACs as $[\delta_{i,j}]_{\Delta^{(c)}}$ for each $i \in [n_{\mathsf{batch}}], j \in [\lambda]$.
   c) **Raising powers:** for each $i \in [n_{\mathsf{batch}}], j \in [\lambda]$, $\mathcal{P}_s$ sends $\alpha_i^{2^j} - \delta_{i,j}$, then parties compute $[\alpha_i^{2^j}]_{\Delta^{(c)}} := [\delta_{i,j}]_{\Delta^{(c)}} + (\alpha_i^{2^j} - \delta_{i,j})$.
   d) **ZKP check on raising powers:** $\mathcal{P}_s$ and $\mathcal{P}_c$ deploy the batched LPZK check (see Section 2.5) to ensure for each $i \in [n_{\mathsf{batch}}], j \in [\lambda]$, $[\alpha_i^{2^{j-1}}]_{\Delta^{(c)}}$ and $[\alpha_i^{2^j}]_{\Delta^{(c)}}$ satisfy the relation: $\left(\alpha_i^{2^{j-1}}\right)^2 = \alpha_i^{2^j}$. Note, this needs a random VOLE correlation from $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$. If the check fails, $\mathcal{P}_c$ aborts the protocol.
5) **Prepare the committed $v$:**
   a) **Commit the OPRF key $k = \Delta^{(s)}$:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{extend}, 1)$ to $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$, and then parties receive length-1 random IT-MACs as $[\beta]_{\Delta^{(c)}}$. Next, $\mathcal{P}_s$ sends $\Delta^{(s)} - \beta$ to $\mathcal{P}_c$, then parties compute $[k = \Delta^{(s)}]_{\Delta^{(c)}} := [\beta]_{\Delta^{(c)}} + (\Delta^{(s)} - \beta)$.
   b) **Commit $\boldsymbol{v}$:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{extend}, n_{\mathsf{batch}})$ to $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$, and then parties receive length-$n_{\mathsf{batch}}$ random IT-MACs as $[\boldsymbol{\zeta}]_{\Delta^{(c)}}$. Next, for each $i \in [n_{\mathsf{batch}}]$, $\mathcal{P}_s$ sends $v_i - \zeta_i$ to $\mathcal{P}_c$, then parties compute $[v_i]_{\Delta^{(c)}} := [\zeta_i]_{\Delta^{(c)}} + (v_i - \zeta_i)$.
   c) **Sample the sacrificed VOLE correlation:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{extend}, 1)$ to $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$, and then $\mathcal{P}_s$ receives $v_{\mathsf{sa}}$ whereas $\mathcal{P}_c$ receives $u_{\mathsf{sa}}, w_{\mathsf{sa}}$ such that $v_{\mathsf{sa}}, u_{\mathsf{sa}}, w_{\mathsf{sa}} \in \mathbb{F}_p, v_{\mathsf{sa}} = w_{\mathsf{sa}} + u_{\mathsf{sa}}\Delta^{(s)}$.
   d) **Commit $v_{\mathsf{sa}}$:** $\mathcal{P}_s$ and $\mathcal{P}_c$ each sends $(\texttt{extend}, 1)$ to $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$, and then parties receive length-1 random IT-MACs as $[\gamma]_{\Delta^{(c)}}$. Next, $\mathcal{P}_s$ sends $v_{\mathsf{sa}} - \gamma$ to $\mathcal{P}_c$, then parties compute $[v_{\mathsf{sa}}]_{\Delta^{(c)}} := [\gamma]_{\Delta^{(c)}} + (v_{\mathsf{sa}} - \gamma)$.
   e) **Perform the consistency check:** $\mathcal{P}_c$ samples $\chi \in \mathbb{F}_p$, then computes $u_{\mathsf{poly}} := u_{\mathsf{sa}} + \sum_{i=1}^{n_{\mathsf{batch}}} u_i \cdot \chi^i$ and $w_{\mathsf{poly}} := w_{\mathsf{sa}} + \sum_{i=1}^{n_{\mathsf{batch}}} w_i \cdot \chi^i$. Next, $\mathcal{P}_c$ sends $\chi, u_{\mathsf{poly}}, w_{\mathsf{poly}}$ to $\mathcal{P}_s$, and then $\mathcal{P}_s$ computes $v_{\mathsf{poly}} := v_{\mathsf{sa}} + \sum_{i=1}^{n_{\mathsf{batch}}} v_i \cdot \chi^i$; $\mathcal{P}_s$ checks if $v_{\mathsf{poly}} \stackrel{?}{=} w_{\mathsf{poly}} + u_{\mathsf{poly}}\Delta^{(s)}$, if not, $\mathcal{P}_s$ aborts the protocol. Finally, parties *locally* compute the IT-MAC

$$\begin{aligned}[v_{\mathsf{poly}} - w_{\mathsf{poly}} - u_{\mathsf{poly}} \cdot \Delta^{(s)}]_{\Delta^{(c)}} &= [v_{\mathsf{poly}}]_{\Delta^{(c)}} - w_{\mathsf{poly}} - u_{\mathsf{poly}} \cdot [\Delta^{(s)}]_{\Delta^{(c)}} \\ &= [v_{\mathsf{sa}}]_{\Delta^{(c)}} + \sum_{i=1}^{n} \chi^i \cdot [v_i]_{\Delta^{(c)}} - w_{\mathsf{poly}} - u_{\mathsf{poly}} \cdot [\Delta^{(s)}]_{\Delta^{(c)}}\end{aligned}$$

   and $\mathcal{P}_s$ proves to $\mathcal{P}_c$ that it commits to a zero via opening the IT-MAC. If the check fails, $\mathcal{P}_c$ aborts the protocol.

Finally, $\mathcal{P}_s$ and $\mathcal{P}_c$ each sets $n_{\mathsf{eval}} := 0$, ignoring the subsequent $(\texttt{init}, \cdot)$. In the half-malicious protocol, $\mathcal{P}_s$ samples $\boldsymbol{\alpha} \stackrel{\$}{\leftarrow} \mathbb{F}_p^{n_{\mathsf{batch}}}$ and computes $\alpha_i^{2^\lambda}$ for each $i \in [n_{\mathsf{batch}}]$ (in the malicious case, $\boldsymbol{\alpha}$ is sampled and $\boldsymbol{\alpha}^{2^\lambda}$ is computed in Sub-step 4a). Then, $\mathcal{P}_s$ outputs $k$.

**Online Phase**

**Evaluate.** $\mathcal{P}_s$ on receiving $(\texttt{eval}, n)$ and $\mathcal{P}_c$ on receiving $(\texttt{eval}, n, \boldsymbol{x})$, where $n \in \mathbb{Z}^+$ and $\boldsymbol{x} \in \mathbb{F}_p^n$. If $n_{\mathsf{eval}} + n > n_{\mathsf{batch}}$, ignore the instruction; otherwise, proceed as follows:

6) $\mathcal{P}_c$ **sends the first message:** For each $i \in [n]$, $\mathcal{P}_c$ sends $\mathsf{msg}_{1,i} := -w_{n_{\mathsf{eval}}+i} + u_{n_{\mathsf{eval}}+i} \cdot x_i$ to $\mathcal{P}_s$.
7) $\mathcal{P}_s$ **sends the second message:** For each $i \in [n]$, $\mathcal{P}_s$ sends $\mathsf{msg}_{2,i} := \alpha_{n_{\mathsf{eval}}+i}^{2^\lambda} \cdot (\mathsf{msg}_{1,i} + v_{n_{\mathsf{eval}}+i})$ to $\mathcal{P}_c$. If $\exists i \in [n], \mathsf{msg}_{2,i} = 0$, $\mathcal{P}_c$ aborts the protocol. (Note, the computing of $\boldsymbol{\alpha}^{2^\lambda}$ has been performed in the offline phase.)
8) **ZK proofs on a well-formed second message:** $\mathcal{P}_s$ and $\mathcal{P}_c$ deploy the batched LPZK check (see Section 2.5) to ensure for each $i \in [n]$, $[v_{n_{\mathsf{eval}}+i}]_{\Delta^{(c)}}$ (from Sub-step 5b), $[\alpha_{n_{\mathsf{eval}}+i}^{2^\lambda}]$ (from Sub-step 4c), $\mathsf{msg}_{1,i}$ (from Step 6), and $\mathsf{msg}_{2,i}$ (from Step 7) satisfy the relation: $\alpha_{n_{\mathsf{eval}}+i}^{2^\lambda} \cdot (\mathsf{msg}_{1,i} + v_{n_{\mathsf{eval}}+i}) = \mathsf{msg}_{2,i}$. Note, this needs a random VOLE correlation from $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$, which can be generated in the offline phase. If the check fails, $\mathcal{P}_c$ aborts the protocol.
9) $\mathcal{P}_c$ **computes the final output:** for each $i \in [n]$, $\mathcal{P}_c$ computes $y_i := \left(u_{n_{\mathsf{eval}}+i}^{-1} \cdot \mathsf{msg}_{2,i}\right)^g \in \mathbb{F}_p$. $\mathcal{P}_c$ outputs $(\texttt{eval}, \boldsymbol{y})$.

Finally, $\mathcal{P}_s$ and $\mathcal{P}_c$ each sets $n_{\mathsf{eval}} := n_{\mathsf{eval}} + n$.

Figure 5: Our 2PC-Gold protocols in the $(\mathcal{F}_{\mathsf{VOLE}}^{c,s}, \mathcal{F}_{\mathsf{VOLE}}^{s,c})$-hybrid model. The descriptions include our half-malicious and malicious protocols. In particular, the malicious protocol includes the steps (and corresponding sub-steps) in gray boxes. Note that for the malicious `eval` with $n = 1$, it suffices to use (cheaper) non-batched LPZK technique in Step 8.

used across multiple sessions and/or clients.

Verifiability is important, sometimes essential, in many applications. Naturally, this requires the client to keep verification information across OPRF invocations.

**Private Verifiability.** We show how to achieve *private verifiability*, namely, the verification information is specific to a client. This client wants to participate in an OPRF protocol multiple times and ensure that the server, which can be malicious, only uses a single key.

Note, since the PRF key is embedded in the VOLE correlations generated by $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$, if $\mathcal{P}_c$ can save all correlations (i.e., malicious batched offline phases), it already achieves private verifiability—the ZKP with a wrong key fails w.h.p.

We highlight that $\mathcal{P}_c$ can also achieve this by saving only one *fresh* correlation related to $k$. Specifically, consider $\mathcal{P}_c$ saves $u$ and $w$, whereas $\mathcal{P}_s$ has $k$ and $v = uk + w$. Now, during the new invocation, we can start a new $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$ where $\mathcal{P}_s$ inputs her key $k'$ (via sending $k' - \Delta_{\mathsf{new}}^{(s)}$) and the client verifies that $k' = k$. For this, let $v' = u'k' + w'$ be a correlation generated by this new $\mathcal{F}_{\mathsf{VOLE}}^{c,s}$. Parties proceed:

1) $\mathcal{P}_c$ sends $\alpha := u - u'$.
2) $\mathcal{P}_s$ commits $\beta := v - v' - \alpha \cdot k$ via a commitment scheme.
3) $\mathcal{P}_c$ sends $\gamma := w - w'$. If $\gamma \neq \beta$, $\mathcal{P}_s$ aborts.
4) $\mathcal{P}_s$ opens $\beta$, If $\beta \neq \gamma$, $\mathcal{P}_c$ aborts.

Here, the commitment scheme prevents a malicious $\mathcal{P}_c$ to learn $k$. Additionally, since $u$ and $u'$ are uniformly sampled, $w - w' = v - v' - (uk - u'k')$ is uniformly random to $\mathcal{P}_s$ given $v, v', u - u'$ when $k \neq k'$. After verification, these two involved correlations are discarded. Note that the $\mathcal{F}_{\mathsf{VOLE}}^{s,c}$ to support ZK (and Steps 4 and 5 in Figure 5) must also be re-executed. These steps need to be executed *before* $k' - \Delta_{\mathsf{new}}^{(s)}$ is sent. This ensures that $\Delta_{\mathsf{new}}^{(s)}$ remains full entropy to the UC environment to protect against a malicious client.

We note that the UC environment can learn $k$ from the honest $\mathcal{P}_s$ output, so a malicious $\mathcal{P}_c$ can easily pass the equality check in Step 3 with maliciously-chosen $\alpha^* = u - u' + \delta$ and $\gamma^* = w - w' - \delta k$ for some $\delta \neq 0$. While this implies the extraction of $k$ by the UC simulator, we still need to simulate the honest $\mathcal{P}_s$'s abort. This can be done by adding one extra instruction in the ideal functionality, defined as follows:

> "**Global-key query.** If $\mathcal{P}_c$ is corrupted, receive $(\texttt{guess}, k')$ from $\mathcal{S}$: if $k = k'$, send `success` to $\mathcal{S}$ and ignore any subsequent query; otherwise, send `abort` to both parties."

This global-key query is harmless as the $\mathcal{P}_c$ with PRF evaluation can always verify his guess on the key locally.

Finally, we remark that the above technique to achieve private verifiability with a single correlation can be used to remove the predetermined $n_{\mathsf{batch}}$ limit: it allows the execution of a new offline phase with a new $n_{\mathsf{batch}}$ and ensures the same key is reused. Actually, we can directly adopt the commit-verify-open idea to the consistency check for $[v]_{\Delta^{(c)}}$ (particularly the proof of a zero IT-MAC; see Section 3.2) in a way that it can be executed even after the key $k$ (or $\Delta^{(s)}$) is revealed to the UC environment.

**Public Verifiability.** We show how to achieve *public verifiability*, namely, the verification information is public (denoted as $\mathsf{VK}$) and can be used by multiple clients. $\mathsf{VK}$ can be viewed as a public key of the server, hiding and binding to her PRF key $k$, and it needs to be obtained by the clients from a reliable source (e.g., authenticated by the server, via certificates, etc.). This corresponds to the notion of *Verifiable* OPRF (VOPRF) in the literature.

We achieve this inspired by [13, Lemma 7 and 8], which focuses on building (V)OPRFs from Legendre PRFs. We first port the core Lemma into our $\mathsf{Gold}$ setting as follows (proof is deferred to the full version [83]):

**Lemma 2.** *For a prime $p$ and any positive integer $g$ that divides $p - 1$, for $m$ uniformly sampled elements $\ell_1, \ldots, \ell_m \xleftarrow{\$} \mathbb{F}_p$, the following statement holds:*

$$\Pr[\exists k \neq k', \forall i, (k + \ell_i)^g = (k' + \ell_i)^g] \leq \frac{(p-1)(g-1)^m}{2p^{m-1}}$$

In our approach, let $\ell_1, \ldots, \ell_m \in \mathbb{F}_p$ be public parameters which are assumed to be uniform; $\mathcal{P}_s$ publishes

$$\mathsf{VK}_m := (k + \ell_1)^g, \ldots, (k + \ell_m)^g \tag{1}$$

as the public key. Based on Lemma 2, this achieves statistical binding for large enough $m$. Concretely, since $\log p \approx 3\lambda$ and $\log g \approx 2\lambda$, we can set $m = 7$. For a uniformly sampled $k$, this also achieves computational hiding based on the hardness of the DSPRS problem (see Assumption 1). We only consider the uniform key since this is typical of how the OPRF key is chosen in the first place.

To validate, we exploit $[k]_{\Delta^{(c)}}$, the IT-MAC of $k$, generated intermediately in our malicious 2PC-$\mathsf{Gold}$ (i.e. Substep 5a in Figure 5). With it, $\mathcal{P}_s$ can prove that for each $i \in [m]$, $(k + \ell_i)^g$ meets the $i$-th entry of $\mathsf{VK}_m$ where $\ell_i$ is public. Note, instead of raising to the power $g$ directly, the random root technique (see Section 3.1) can be applied for better efficiency: $\mathcal{P}_s$ can let $\mathcal{P}_c$ get $r_i^{2^\lambda} \cdot (k + \ell_i)$ for $r_i \xleftarrow{\$} \mathbb{F}_p$. In detail, parties generate $[r_i^{2^\lambda}]_{\Delta^{(c)}}$, $i \in [m]$, which can be merged into Step 4; then $\mathcal{P}_s$ sends $r_i^{2^\lambda} \cdot (k + \ell_i)$ and proves each multiplication is done correctly, which can be merged into Step 8. This results in $m\frac{\lambda}{\phi} + m$ additional $\mathbb{F}_p$ elements of communication (independently of $n$) and $m\frac{\lambda}{\phi} + m$ additional VOLE correlations, in the VOLE-hybrid model. See Appendix A for concrete cost, including VOLE.

Practically, $\boldsymbol{\ell}$ can be generated by applying a public hash function $\mathsf{H}$ over $[m]$. In this case, it is essential to hash the inputs to $\mathsf{Gold}$ with an independent hash function (this is $\mathsf{H}_1$ in the O-$\mathsf{Gold}$ function discussed below). Otherwise, $\mathsf{VK}_m$ would reveal $m$ $\mathsf{Gold}$ evaluations.

**Stateless Verifiability.** Either private or public verifiability discussed above requires the client to be stateful or be able to receive authenticated $\mathsf{VK}$. However, in some applications, the client is stateless and has no means to guarantee an authenticated $\mathsf{VK}$. Interestingly, in applications such as Password-Protected Secret Sharing (PPSS) [51], a weaker form of verification suffices: if a malicious $\mathcal{P}_s$ changes the key with $\mathcal{P}_c$ using the same input, the output is different.

2PC-Gold does not ensure this property, particularly due to the collision issues discussed in Section 6.1. The same is the case for the function O-Gold that we defined (see section 1.1) as the 2Hash mode of Gold, namely,

$$\mathsf{O\text{-}Gold}_k(x) := \mathsf{H}_2(x, \mathsf{Gold}_k(\mathsf{H}_1(x))).$$

However, the above property can be provided by including the verification value $\mathsf{VK}_6$ in the computation of O-Gold as

$$\mathsf{H}_2(x, \mathsf{Gold}_k(\mathsf{H}_1(x)), \mathsf{Gold}_k(\ell_1), \ldots, \mathsf{Gold}_k(\ell_6)) \quad (2)$$

where $\ell_1, \ldots, \ell_6 \xleftarrow{\$} \mathbb{F}_p$ are public parameters. This value $\mathsf{VK}_6$ and its binding property (i.e., Lemma 2) ensures two keys produce two different public keys, ultimately resulting in two different OPRF outputs. Note that we only need $\mathsf{VK}_6$ rather than $\mathsf{VK}_7$ since $\mathsf{H}_1(x)$ acts as an extra point. Looking ahead, this modified O-Gold implements a strong UC OPRF [13], which also explains why it supports PPSS [51].

### 6.2. Avoiding Collisions in Gold

In this section, we discuss collisions in the Gold function and their effect in defining a Gold-based UC OPRF in the strong sense (e.g., [13], [50], [51]). By the end of this section, we show how to build such an OPRF (UC-Gold).

Collisions in Gold—pairs $(k, x) \neq (k', x')$ s.t. $\mathsf{Gold}_k(x) = \mathsf{Gold}_{k'}(x')$—are trivial to find. However, when considering the function $\mathsf{O\text{-}Gold}_k(x) := \mathsf{H}_2(x, \mathsf{Gold}_k(\mathsf{H}_1(x)))$ where $\mathsf{H}_1$ and $\mathsf{H}_2$ are modeled as random oracles, the latter with $2\lambda$ bits of output, finding collisions with $x \neq x'$ is infeasible. Yet, O-Gold inherits one form of Gold collisions, namely, colliding pairs $(k, x), (k', x)$ for $k \neq k'$. To see that such pairs exist (and can be computed), consider a procedure similar to the algorithm in Figure 3 that on input $k, x$ finds $X$ such that $X^g = (k + x)^g$ and then outputs $k' := X - x$.

While collisions of the form $(k, x), (k', x)$ do not violate the standard security of PRFs, OPRF applications often require stronger properties. In particular, strong UC formulations of OPRF model these functions as random oracles with outputs independent for any two pairs $(k, x) \neq (k', x')$. Therefore, collisions of the form $(k, x), (k', x)$ are not allowed (some OPRF applications, e.g., [51], are actually insecure in the presence of such collisions).

Hence, to achieve such strong UC OPRFs, we need to eliminate these collisions. Thanks to the fact (Lemma 2) that $\mathsf{Gold}_k(\ell_1), \ldots, \mathsf{Gold}_k(\ell_6)$ act as a commitment to a single $k$, the function defined in Equation (2) achieves this.

This function can be proven UC secure using the formalism and methodology from [13]. Informally, they show that the 2Hash mode applied to a keyed function $F$, namely, $\mathsf{H}_2(x, F_k(\mathsf{H}_1(x)))$, results in a strong UC-secure OPRF (in their corresponding OPRF functionality) provided the following properties hold for $F$: (1) $F$ has a secure UC two-party computation between a server with input a key $k$ and a client with input $x$ (similar to the leakage-free version of 2PC-Gold functionality from Figure 4); (2) $F$ is one-more unpredictable; and (3) collisions $F_k(x) = F_{k'}(x)$ are hard to find for any $x$ and any $k \neq k'$.

For our case, we consider the function $F_k(x)$ defined as the concatenation of the 7 values:

$$\mathsf{Gold}_k(\mathsf{H}_1(x)), \mathsf{Gold}_k(\ell_1), \ldots, \mathsf{Gold}_k(\ell_6).$$

This $F$ inherits from Gold its UC 2PC security (Section 5.1) and one-more unpredictability (implied by the hardness of the DSPRS problem; see Assumption 1). In addition, it is collision-resistant (in the sense of condition (3) above) based on Lemma 2. Thus, if we consider $F$ in 2Hash mode, we obtain exactly the function defined in Equation (2), and the results from [13, Theorem 1] imply this function is a UC OPRF in the formalization of [13]. For concreteness and future use, we define:

$$\mathsf{UC\text{-}Gold}_k(x) := \mathsf{H}_2\big(x, \mathsf{Gold}_k(\mathsf{H}_1(x)), \mathsf{Gold}_k(\mathsf{H}_0(1)), \ldots, \mathsf{Gold}_k(\mathsf{H}_0(6))\big)$$

where $\mathsf{H}_0, \mathsf{H}_1, \mathsf{H}_2$ are hash functions modeled as independent random oracles with ranges $\mathbb{F}_p$, $\mathbb{F}_p$ and $\{0,1\}^{2\lambda}$, resp.

## 7. Implementation and Benchmark

We implemented our half-malicious and malicious 2PC-Gold with both non-batched and batched variants (Figure 5) using C++. In particular, since the performance difference is nuanced, we consider batched variants with a single batch, i.e., $n_{\mathsf{batch}} = n$. In this section, we discuss this implementation and provide a comprehensive benchmark.

We only implemented 2PC-Gold since it is sufficient for many applications and reflects our performance. Note, 2PC-Gold and O-Gold have similar concrete performance, and we estimate additional costs needed for UC-Gold in Appendix A. One can easily extend our compact and modular implementation to O-Gold and UC-Gold.

We only consider the case where the key is uniformly sampled since extra costs to support a server-specified key are negligible, as discussed in Appendix A.

In this paper, we use the convention: 1KB$= 2^{10}$B.

**Open-Source Implementation.** Our implementation is public and available at https://github.com/gconeice/PR-OPRF.

### 7.1. Setup

**Security Level.** Our implementation considers $\lambda = 128$, aiming at the NIST Security Strength Category 1 for post-quantum cryptography. In particular, we use 128-bit OTs, and AES-128 to implement, e.g., the PRNG. We chose this based on baselines in the literature. We note that the famous Grover algorithm [44] can indeed provide a square-root quantum attack over AES-128, but this is not considered as a practical attack (see, e.g., [43], [85]). We can upgrade our implementation to be based on 256-bit OTs and AES-256 with some engineering efforts, and our estimated performance overhead is only $\leq 2\times$.

**Prime $p$.** Our implementation sets $p = 2^{128} \cdot g + 1$ as a 384-bit prime to produce 128-bit OPRF outputs. We choose $g = 2^{256} - 33375$, which is the largest 256-bit prime, ensuring $p$ is also a prime. Observe that for our selected $p$, hash

functions or PRNGs generating outputs in $\{0,1\}^{384}$ suffice to produce $\mathbb{F}_p$ elements.[9] We remark that $g$ does not need to be a prime, but we chose a prime $g$ as a conservative option. We use the `GMP` library [41] for $\mathbb{F}_p$ operations.

**Functionality $\mathcal{F}_{\mathsf{VOLE}}$.** Recall that our protocols are designed in the VOLE-hybrid (Figure 1) model. To generate these VOLE correlations, we deploy the following malicious-secure VOLE protocols:

- **Non-batched variant:** Our non-batched (i.e., single-input) variant, either half-malicious or malicious, only requires a small number of VOLE correlations. Hence, we implemented the VOLE protocol in [13], which relies on [9], [71]. In short, in this protocol, besides a one-time setup to generate random OTs and some GGM trees, each VOLE correlation requires $\approx \frac{\log p + 2\sigma}{t}$ elements of $\mathbb{F}_p$, where $\sigma = 64$ is the statistical security parameter and $t = 8$ is a communication-computation trade-off parameter—a per-correlation cost of $\approx 3$KB.
- **Batched variant:** Our $n$-batched (particularly for a sufficiently large $n$) variant, either half-malicious or malicious, requires a large number of VOLE correlations. Hence, we deployed the [81]'s VOLE protocol to extend a small amount of VOLE correlations into a large amount of VOLE correlations with sublinear communication (in length) based on OTs and the post-quantum LPN assumption. That is, amortized communication is almost free. [81]'s implementation is open-sourced in the `EMP-Toolkit` [80], and we adopted it with adjustments toward $\mathbb{F}_p$ and PQ OTs.

**OTs.** We utilize the `libOTe` library [70] to implement OTs. These (random) OTs are generated using the state-of-the-art malicious-secure post-quantum OT extension technique [71] relied on base (random) OTs, which we generate using the malicious-secure OT protocol in [62]. In particular, [62]'s OT protocol can be instantiated under either a classical (i.e., Diffie-Hellman-type with curve25519) or post-quantum (i.e., lattice-type with Kyber512) assumption. Switching between these instantiations yields classical or post-quantum security to the full 2PC-Gold. We consider both in our benchmark.

**Hardware.** Our experiments were executed on two AWS EC2 `m5.large` machines[10] that respectively implemented $\mathcal{P}_s$ and $\mathcal{P}_c$. Each party ran single-threaded. We configured different network settings via Linux `tc` command:

- **WAN-like:** 25Mbps with a 30ms round-trip latency.
- **LAN-like:** 1Gbps with a 2ms round-trip latency.

These configurations are selected to match prior work.

**Metrics.** We report the following metrics:

- **Communication:** the total communication.
- **Computation:** the total execution time of $\mathcal{P}_s$ and $\mathcal{P}_c$ respectively. Note that the end-to-end (E2E) execution time is exactly the $\mathcal{P}_c$'s execution time.

**Optimization Parameter $\phi$.** Recall that our malicious 2PC-Gold can leverage the generalized LPZK technique [82] (Section 2.5) to balance communication and computation. We set $\phi = 4$. See [83] for related experiments.

9. A random integer in $[2^{384}]$ is $< p$ with overwhelming probability.
10. Intel Xeon Platinum 8259CL @ 2.50GHz, 2 vCPUs, 8GiB Memory

## 7.2. Performance of Our Protocols

**Overall Performance.** Table 1a (resp. Table 1b) tabulates the results with classical (resp. post-quantum) OTs.

For the batched test cases, we set the evaluation number $n$ large enough to enable VOLE extension and use up the extended VOLE correlations, seeking the best amortization. In particular, $n \approx 10^7$ (resp. $\approx 3 \times 10^5$) for half-malicious (resp. malicious) protocols.[11] This is purely for benchmarking. While this $n \approx 10^7$ (or $\approx 3 \times 10^5$) is determined by the LPN parameter in [80], one can adjust their VOLE extension parameters (e.g., using the LPN estimator [60]) to smaller values (e.g., $10^5$) with *little effect* on the amortized cost.

The performance of our batched variant is almost identical between classical and post-quantum instantiations. This is because only $\lambda = 128$ base OTs need to be generated accordingly, and the (amortized) cost difference is negligible.

**Fine-Grained Analysis.** We performed fine-grained analysis over the communication of our non-batched protocols, tabulated in Table 2. Here, $\Pi_{\mathsf{offline}}$ (resp. $\Pi_{\mathsf{online}}$) denotes the offline (resp. online) phase of our protocols (see Figure 5) in the VOLE-hybrid model; the communication of OTs includes the generation of base OTs and OT extension. See the full version [83] for more discussions.

Almost all communication of our protocols is used to generate VOLE correlations. As our protocols are black-box in VOLE, it is valuable to study how to more efficiently generate a small amount of VOLE correlations, even classically. Indeed, even improving the cost of generating (base) PQ OTs (e.g., by using [32]) may significantly increase the performance of our non-batched PQ performance.

**Selecting $n$ for Cost Amortization.** In our batched setting, we need a large enough $n$ to effectively amortize the communication cost. For example, in our malicious PQ batched setting, the amortized cost comes from two components: (1) a per OPRF evaluation cost that we can estimate to be about 1.9KB using Table 1b (where $n \approx 3 \times 10^5$); (2) the generation of around $1,800$ base VOLE correlations (in each direction, depending *solely* on the LPN parameters), necessary to initiate VOLE extensions using the *primal LPN*. We can use Table 2 to estimate the latter as follows:

$$\overbrace{(430+333)\text{KB}}^{\text{Base PQ OTs in Table 2}} + \overbrace{(175+28)\text{KB}}^{\text{VOLE setup in Table 2}} + \overbrace{3600 \cdot 3\text{KB}}^{3,600 \text{ more VOLEs}}$$

which is approximately $12,000$KB. We can then approximate the amortized cost for any $n$ to: $1.9\text{KB} + 12,000\text{KB}/n$. For example, for $n = 10^3$, the cost is about 14KB; for $n = 10^5$, it is about 2KB.

We instantiated the LPN-based VOLE extension with a smaller $n$ to verify the above estimation; see Table 3.

11. Note, $\mathcal{P}_s$ and $\mathcal{P}_c$ each only needs to save the base VOLE correlations to support, e.g., day-to-day invocations. For example, for $n = 10^7$, about 5000 base correlations suffice, requiring about 1MB physical memory.

| Variant | Security $\mathcal{P}_s$-$\mathcal{P}_c$ | Comm. | WAN, Time $\mathcal{P}_s$ | WAN, Time $\mathcal{P}_c$ | LAN, Time $\mathcal{P}_s$ | LAN, Time $\mathcal{P}_c$ |
|---|---|---|---|---|---|---|
| Non-batched | ◐-● | 46KB | 338ms | 368ms | 158ms | 160ms |
| (single-input) | ●-● | 242KB | 936ms | 966ms | 506ms | 508ms |
| Batched | ◐-● | 99B | 61µs | 87µs | 30µs | 56µs |
| (amortized) | ●-● | 1.9KB | 1.6ms | 1.6ms | 1ms | 1ms |

(a) Classical

| Variant | Security $\mathcal{P}_s$-$\mathcal{P}_c$ | Comm. | WAN, Time $\mathcal{P}_s$ | WAN, Time $\mathcal{P}_c$ | LAN, Time $\mathcal{P}_s$ | LAN, Time $\mathcal{P}_c$ |
|---|---|---|---|---|---|---|
| Non-batched | ◐-● | 774KB | 538ms | 568ms | 161ms | 163ms |
| (single-input) | ●-● | 970KB | 1.1s | 1.1s | 507ms | 510ms |
| Batched | ◐-● | 100B | 61µs | 87µs | 30µs | 57µs |
| (amortized) | ●-● | 1.9KB | 1.6ms | 1.6ms | 1ms | 1ms |

(b) Post-quantum

TABLE 1: Performance of our protocols with classical and post-quantum OTs. The overall security is inherited from the OT. ◐ denotes the semi-honest security whereas ● denotes the malicious security. The costs of batched protocols correspond to the costs of a single OPRF evaluation when amortized as part of an $n$-batched evaluation with $n \approx 10^7$ for half-malicious and $n \approx 3 \times 10^5$ for malicious (smaller $n$ are also possible, see text and Table 3).

| Security $\mathcal{P}_s$-$\mathcal{P}_c$ | | Offline Comm. (B): OT $\mathcal{P}_s \to \mathcal{P}_c$ | OT $\mathcal{P}_c \to \mathcal{P}_s$ | VOLE $\mathcal{P}_s \to \mathcal{P}_c$ | VOLE $\mathcal{P}_c \to \mathcal{P}_s$ | $\Pi_{\mathsf{offline}}$ $\mathcal{P}_s \to \mathcal{P}_c$ | $\Pi_{\mathsf{offline}}$ $\mathcal{P}_c \to \mathcal{P}_s$ | Online Comm. (B): $\Pi_{\mathsf{online}}$ $\mathcal{P}_s \to \mathcal{P}_c$ | $\Pi_{\mathsf{online}}$ $\mathcal{P}_c \to \mathcal{P}_s$ |
|---|---|---|---|---|---|---|---|---|---|
| ◐-● | Classical | 13,805 | 8,368 | 16 | 25,552 | – | – | 48 | 48 |
| | PQ | 439,757 | 327,800 | | | | | | |
| ●-● | Classical | 14,917 | 22,109 | 179,792 | 28,592 | 2,496 | 192 | 144 | 48 |
| | PQ | 440,869 | 341,541 | | | | | | |

TABLE 2: Fine-grained communication analysis of our non-batched protocols. ◐ denotes the semi-honest security whereas ● denotes the malicious security. Note that the difference between classical and PQ instantiations lies only in OTs.

| Variant | Security $\mathcal{P}_s$-$\mathcal{P}_c$ | Comm. | WAN, Time $\mathcal{P}_s$ | WAN, Time $\mathcal{P}_c$ | LAN, Time $\mathcal{P}_s$ | LAN, Time $\mathcal{P}_c$ |
|---|---|---|---|---|---|---|
| Batched | ◐-● | 175B | 181µs | 216µs | 120µs | 146µs |
| (amortized) | ●-● | 17.1KB | 24.4ms | 24.4ms | 19.4ms | 19.4ms |

TABLE 3: Performance of our PQ protocols with moderate size $n$. ◐ denotes the semi-honest security whereas ● denotes the malicious security. The costs of batched protocols correspond to the costs of a single OPRF evaluation when amortized as part of an $n$-batched evaluation with $n = 10^5$ for half-malicious and $n = 10^3$ for malicious.

An even smaller cost of (2) can be achieved by using VOLE extension protocols that rely on the *dual LPN* [18], which requires fewer base VOLE correlations.

### 7.3. Comparison with Prior Work

We compare our protocols concretely with prior post-quantum OPRFs. Most of them do not have a (public) implementation, but we try our best to compare with them:

- **Isogeny-based:** The state-of-the-art *semi-honest* (both client and server are semi-honest) OPRF based on isogenies is OPUS [48]. OPUS focuses on the non-batched setting, and it is unclear how it can be optimized/amortized in the batched setting. When $\lambda = 128$, the communication cost of OPUS is $\approx$ 24KB. This outperforms our half-malicious *non-batched* PQ protocol. (Note that in our half-malicious model, the client is fully malicious, and only the server is semi-honest.)
  However, OPUS requires 258 rounds while ours requires 3 rounds in the VOLE-hybrid model (with Fiat-Shamir), and the generation of VOLE correlations can be finished in $\leq$5 rounds. More importantly, OPUS is extremely computationally intensive. We tested the open-sourced OPUS implementation on our machines, and it required over 13s E2E for each evaluation in the WAN. Hence, ours is over $20\times$ better in terms of E2E time and has stronger security. Our improvement is over $50\times$ in LAN and over $100000\times$ in the batched setting.
  The state-of-the-art malicious isogeny-based OPRF is [7], which requires 8.7MB communication. Our communication is over $9\times$ smaller. Note, [7]'s protocol is a theoretical result, and no implementation is available. Similar to OPUS, it is very heavy in computation.
- **Lattice-based:** The state-of-the-art malicious OPRF based on lattices is [4], [33], which requires $\approx$200KB in the batched amortized setting. Ours is $\approx 100\times$ better. Both [4] and [33] need 2 rounds.
  The state-of-the-art semi-honest OPRF based on lattices is [49], which requires 23KB in the batched amortized setting. Ours is $\approx 235\times$ better. [49] needs 6 rounds.
  On the other hand, [33] requires $\approx$400KB in the non-batched malicious setting, which is $\approx 2\times$ better than our corresponding protocol in communication. However, [33] relies on SNARG, which is computationally expensive.
- **"Crypto-Dark-Matter" [16]:** The state-of-the-art semi-honest OPRF based on the "Crypto-Dark-Matter" is [1], which requires 119 Bytes of communication in the batched amortized setting. Ours achieves 100 Bytes, and in addition, we provide security against a malicious client. [1] needs 2 rounds in the VOLE-hybrid model, same as ours. On the other hand, in general, "Crypto-Dark-Matter" PRF constructions have a lower computational overhead. It should also be noted that the parameters used in [1] are overly aggressive, as analyzed by [5].
- **Legendre-based:** The state-of-the-art malicious OPRF based on regular Legendre PRFs is [13] which focuses on non-batched setting and estimates a 911KB communication (without implementation). This protocol needs 9 rounds [33].

In comparison, our *real-world tested* communication is 970KB. However, [13] estimates a 296KB cost of OTs, relying on [6] to generate base OTs. By using these OTs, our communication can get down to 502KB—1.8× better. More significantly, our per-evaluation communication reduces to just 1.9KB when amortized over batched evaluations. In contrast, [13]'s protocol does not support batched amortization. This is inherent in their approach because they use (non-standard) VOLE to commit the evaluation input $x$ by revealing $x - \Delta$, where $\Delta$ is the scalar from the VOLE. Thus, for each new $x$, new VOLE correlations (with a new $\Delta$) must be generated. See also Section 1.2. Note that our protocols rely solely on standard VOLE correlations in a black-box manner. [13]'s protocol requires a customized VOLE functionality. Hence, our protocols are more friendly to future improvements of VOLE.

## Acknowledgement

We thank our reviewers for their thoughtful feedback. We also thank Ward Beullens, Julia Hesse, Yuval Ishai, Peter Rindal, and Lawrence Roy for insightful discussions.

## References

[1] N. Alamati, G.-V. Policharla, S. Raghuraman, and P. Rindal, "Improved alternating-moduli PRFs and post-quantum signatures," in *CRYPTO 2024, Part VIII*, ser. LNCS, L. Reyzin and D. Stebila, Eds., vol. 14927. Springer, Cham, Aug. 2024, pp. 274–308.

[2] M. R. Albrecht, A. Davidson, A. Deo, and D. Gardham, "Crypto dark matter on the torus - oblivious PRFs from shallow PRFs and TFHE," in *EUROCRYPT 2024, Part VI*, ser. LNCS, M. Joye and G. Leander, Eds., vol. 14656. Springer, Cham, May 2024, pp. 447–476.

[3] M. R. Albrecht, A. Davidson, A. Deo, and N. P. Smart, "Round-optimal verifiable oblivious pseudorandom functions from ideal lattices," in *PKC 2021, Part II*, ser. LNCS, J. Garay, Ed., vol. 12711. Springer, Cham, May 2021, pp. 261–289.

[4] M. R. Albrecht and K. D. Gür, "Verifiable oblivious pseudorandom functions from lattices: Practical-ish and thresholdisable," in *ASIACRYPT 2024, Part IV*, ser. LNCS, K.-M. Chung and Y. Sasaki, Eds., vol. 15487. Springer, Singapore, Dec. 2024, pp. 205–237.

[5] I. M. Ayala and H. Raddum, "Zeroed out: Cryptanalysis of weak PRFs in alternating moduli," Cryptology ePrint Archive, Report 2024/2055, 2024. [Online]. Available: https://eprint.iacr.org/2024/2055

[6] S. Badrinarayanan, D. Masny, and P. Mukherjee, "Efficient and tight oblivious transfer from PKE with tight multi-user security," in *ACNS 22International Conference on Applied Cryptography and Network Security*, ser. LNCS, G. Ateniese and D. Venturi, Eds., vol. 13269. Springer, Cham, Jun. 2022, pp. 626–642.

[7] A. Basso, "A post-quantum round-optimal oblivious PRF from isogenies," Cryptology ePrint Archive, Report 2023/225, 2023. [Online]. Available: https://eprint.iacr.org/2023/225

[8] A. Basso, P. Kutas, S.-P. Merz, C. Petit, and A. Sanso, "Cryptanalysis of an oblivious PRF from supersingular isogenies," in *ASIACRYPT 2021, Part I*, ser. LNCS, M. Tibouchi and H. Wang, Eds., vol. 13090. Springer, Cham, Dec. 2021, pp. 160–184.

[9] C. Baum, L. Braun, C. Delpech de Saint Guilhem, M. Klooß, E. Orsini, L. Roy, and P. Scholl, "Publicly verifiable zero-knowledge and post-quantum signatures from VOLE-in-the-head," in *CRYPTO 2023, Part V*, ser. LNCS, H. Handschuh and A. Lysyanskaya, Eds., vol. 14085. Springer, Cham, Aug. 2023, pp. 581–615.

[10] D. Beaver, "Precomputing oblivious transfer," in *CRYPTO'95*, ser. LNCS, D. Coppersmith, Ed., vol. 963. Springer, Berlin, Heidelberg, Aug. 1995, pp. 97–109.

[11] R. Bendlin, I. Damgård, C. Orlandi, and S. Zakarias, "Semi-homomorphic encryption and multiparty computation," in *EUROCRYPT 2011*, ser. LNCS, K. G. Paterson, Ed., vol. 6632. Springer, Berlin, Heidelberg, May 2011, pp. 169–188.

[12] W. Beullens, T. Beyne, A. Udovenko, and G. Vitto, "Cryptanalysis of the Legendre PRF and generalizations," *IACR Trans. Symm. Cryptol.*, vol. 2020, no. 1, pp. 313–330, 2020.

[13] W. Beullens, L. Dodgson, S. Faller, and J. Hesse, "The 2Hash OPRF framework and efficient post-quantum instantiations," Cryptology ePrint Archive, Report 2024/450, 2024. [Online]. Available: https://eprint.iacr.org/2024/450

[14] A. Blum, M. L. Furst, M. J. Kearns, and R. J. Lipton, "Cryptographic primitives based on hard learning problems," in *CRYPTO'93*, ser. LNCS, D. R. Stinson, Ed., vol. 773. Springer, Berlin, Heidelberg, Aug. 1994, pp. 278–291.

[15] D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry, "Random oracles in a quantum world," in *ASIACRYPT 2011*, ser. LNCS, D. H. Lee and X. Wang, Eds., vol. 7073. Springer, Berlin, Heidelberg, Dec. 2011, pp. 41–69.

[16] D. Boneh, Y. Ishai, A. Passelègue, A. Sahai, and D. J. Wu, "Exploring crypto dark matter: New simple PRF candidates and their applications," in *TCC 2018, Part II*, ser. LNCS, A. Beimel and S. Dziembowski, Eds., vol. 11240. Springer, Cham, Nov. 2018, pp. 699–729.

[17] D. Boneh, D. Kogan, and K. Woo, "Oblivious pseudorandom functions from isogenies," in *ASIACRYPT 2020, Part II*, ser. LNCS, S. Moriai and H. Wang, Eds., vol. 12492. Springer, Cham, Dec. 2020, pp. 520–550.

[18] E. Boyle, G. Couteau, N. Gilboa, and Y. Ishai, "Compressing vector OLE," in *ACM CCS 2018*, D. Lie, M. Mannan, M. Backes, and X. Wang, Eds. ACM Press, Oct. 2018, pp. 896–912.

[19] R. Canetti, "Universally composable security: A new paradigm for cryptographic protocols," in *42nd FOCS*. IEEE Computer Society Press, Oct. 2001, pp. 136–145.

[20] S. Casacuberta, J. Hesse, and A. Lehmann, "SoK: Oblivious pseudorandom functions," in *2022 IEEE European Symposium on Security and Privacy*. IEEE Computer Society Press, Jun. 2022, pp. 625–646.

[21] M. Chase and P. Miao, "Private set intersection in the internet setting from lightweight oblivious PRF," in *CRYPTO 2020, Part III*, ser. LNCS, D. Micciancio and T. Ristenpart, Eds., vol. 12172. Springer, Cham, Aug. 2020, pp. 34–63.

[22] H. Corrigan-Gibbs and D. J. Wu, "Legendre sequences are pseudorandom under the quadratic-residuosity assumption," Cryptology ePrint Archive, Report 2024/1252, 2024. [Online]. Available: https://eprint.iacr.org/2024/1252

[23] ——, "The one-wayness of jacobi signatures," in *CRYPTO 2024, Part V*, ser. LNCS, L. Reyzin and D. Stebila, Eds., vol. 14924. Springer, Cham, Aug. 2024, pp. 3–13.

[24] I. Damgård, "On the randomness of Legendre and Jacobi sequences," in *CRYPTO'88*, ser. LNCS, S. Goldwasser, Ed., vol. 403. Springer, New York, Aug. 1990, pp. 163–172.

[25] H. Davenport, "On the distribution of quadratic residues (mod p)," *Journal of the London Mathematical Society*, vol. 1, no. 1, pp. 49–54, 1931.

[26] H. Davenport and P. Erdös, "The distribution of quadratic and higher residues," *Publ. Math. Debrecen*, vol. 2, no. 3–4, pp. 252–65, 1952.

[27] A. Davidson, I. Goldberg, N. Sullivan, G. Tankersley, and F. Valsorda, "Privacy pass: Bypassing internet challenges anonymously," *PoPETs*, vol. 2018, no. 3, pp. 164–180, Jul. 2018.

[28] C. Ding, T. Hesseseth, and W. Shan, "On the linear complexity of legendre sequences," *IEEE Transactions on Information Theory*, vol. 44, no. 3, pp. 1276–1278, 1998.

[29] I. Dinur, S. Goldfeder, T. Halevi, Y. Ishai, M. Kelkar, V. Sharma, and G. Zaverucha, "MPC-friendly symmetric cryptography from alternating moduli: Candidates, protocols, and applications," in *CRYPTO 2021, Part IV*, ser. LNCS, T. Malkin and C. Peikert, Eds., vol. 12828. Virtual Event: Springer, Cham, Aug. 2021, pp. 517–547.

[30] S. Dittmer, Y. Ishai, and R. Ostrovsky, "Line-point zero knowledge and its applications," in *2nd Conference on Information-Theoretic Cryptography*, 2021.

[31] Y. Dodis and A. Yampolskiy, "A verifiable random function with short proofs and keys," in *PKC 2005*, ser. LNCS, S. Vaudenay, Ed., vol. 3386. Springer, Berlin, Heidelberg, Jan. 2005, pp. 416–431.

[32] S. Dong, H. Cui, K. Zhang, K. Yang, and Y. Yu, "A simple post-quantum oblivious transfer protocol from mod-LWR," Cryptology ePrint Archive, Paper 2024/1116, 2024. [Online]. Available: https://eprint.iacr.org/2024/1116

[33] M. F. Esgin, R. Steinfeld, E. Tairi, and J. Xu, "LeOPaRd: Towards practical post-quantum oblivious PRFs via interactive lattice problems," Cryptology ePrint Archive, Paper 2024/1615, 2024. [Online]. Available: https://eprint.iacr.org/2024/1615

[34] S. H. Faller, A. Ottenhues, and J. Ottenhues, "Composable oblivious pseudo-random functions via garbled circuits," in *LATINCRYPT 2023*, ser. LNCS, A. Aly and M. Tibouchi, Eds., vol. 14168. Springer, Cham, Oct. 2023, pp. 249–270.

[35] U. Feige, J. Kilian, and M. Naor, "A minimal model for secure computation (extended abstract)," in *26th ACM STOC*. ACM Press, May 1994, pp. 554–563.

[36] A. Fiat and A. Shamir, "How to prove yourself: Practical solutions to identification and signature problems," in *CRYPTO'86*, ser. LNCS, A. M. Odlyzko, Ed., vol. 263. Springer, Berlin, Heidelberg, Aug. 1987, pp. 186–194.

[37] M. J. Freedman, Y. Ishai, B. Pinkas, and O. Reingold, "Keyword search and oblivious pseudorandom functions," in *TCC 2005*, ser. LNCS, J. Kilian, Ed., vol. 3378. Springer, Berlin, Heidelberg, Feb. 2005, pp. 303–324.

[38] P. Frixons and A. Schrottenloher, "Quantum security of the legendre PRF," Cryptology ePrint Archive, Report 2021/149, 2021. [Online]. Available: https://eprint.iacr.org/2021/149

[39] O. Goldreich, S. Goldwasser, and S. Micali, "How to construct random functions," *Journal of the ACM*, vol. 33, no. 4, pp. 792–807, Oct. 1986.

[40] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof-systems (extended abstract)," in *17th ACM STOC*. ACM Press, May 1985, pp. 291–304.

[41] T. Granlund and the GMP development team, *GNU MP: The GNU Multiple Precision Arithmetic Library*, 6th ed., 2020, https://gmplib.org.

[42] L. Grassi, C. Rechberger, D. Rotaru, P. Scholl, and N. P. Smart, "MPC-friendly symmetric key primitives," in *ACM CCS 2016*, E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi, Eds. ACM Press, Oct. 2016, pp. 430–443.

[43] M. Grassl, B. Langenberg, M. Roetteler, and R. Steinwandt, "Applying grover's algorithm to aes: quantum resource estimates," in *International Workshop on Post-Quantum Cryptography*. Springer, 2016, pp. 29–43.

[44] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *28th ACM STOC*. ACM Press, May 1996, pp. 212–219.

[45] K. Gyarmati, C. Mauduit, and A. Sárközy, "The cross-correlation measure for families of binary sequences," in *Applied Algebra and Number Theory*, G. Larcher, F. Pillichshammer, A. Winterhof, and C. Xing, Eds. Cambridge University Press, 2014, no. Theory, pp. 126–143.

[46] C. Hazay, D. Heath, V. Kolesnikov, M. Venkitasubramaniam, and Y. Yang, "LogRobin++: Optimizing proofs of disjunctive statements in VOLE-based ZK," Cryptology ePrint Archive, Paper 2024/1427, 2024. [Online]. Available: https://eprint.iacr.org/2024/1427

[47] C. Hazay and Y. Yang, "Toward malicious constant-rate 2PC via arithmetic garbling," in *EUROCRYPT 2024, Part V*, ser. LNCS, M. Joye and G. Leander, Eds., vol. 14655. Springer, Cham, May 2024, pp. 401–431.

[48] L. Heimberger, T. Hennerbichler, F. Meisingseth, S. Ramacher, and C. Rechberger, "OPRFs from isogenies: Designs and analysis," in *ASIACCS 24*, J. Zhou, T. Q. S. Quek, D. Gao, and A. A. Cárdenas, Eds. ACM Press, Jul. 2024.

[49] L. Heimberger, D. Kales, R. Lolato, O. Mir, S. Ramacher, and C. Rechberger, "Leap: A fast, lattice-based OPRF with application to private set intersection," Cryptology ePrint Archive, Paper 2025/333, 2025. [Online]. Available: https://eprint.iacr.org/2025/333

[50] S. Jarecki, A. Kiayias, and H. Krawczyk, "Round-optimal password-protected secret sharing and T-PAKE in the password-only model," in *ASIACRYPT 2014, Part II*, ser. LNCS, P. Sarkar and T. Iwata, Eds., vol. 8874. Springer, Berlin, Heidelberg, Dec. 2014, pp. 233–253.

[51] S. Jarecki, A. Kiayias, H. Krawczyk, and J. Xu, "Highly-efficient and composable password-protected secret sharing (or: How to protect your bitcoin wallet online)," in *2016 IEEE European Symposium on Security and Privacy*. IEEE Computer Society Press, Mar. 2016, pp. 276–291.

[52] S. Jarecki, H. Krawczyk, and J. K. Resch, "Updatable oblivious key management for storage systems," in *ACM CCS 2019*, L. Cavallaro, J. Kinder, X. Wang, and J. Katz, Eds. ACM Press, Nov. 2019, pp. 379–393.

[53] S. Jarecki, H. Krawczyk, and J. Xu, "OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks," in *EUROCRYPT 2018, Part III*, ser. LNCS, J. B. Nielsen and V. Rijmen, Eds., vol. 10822. Springer, Cham, Apr. / May 2018, pp. 456–486.

[54] N. Kaluderovic, N. Cheng, and K. Mitrokotsa, "A post-quantum distributed OPRF from the legendre PRF," Cryptology ePrint Archive, Report 2024/544, 2024. [Online]. Available: https://eprint.iacr.org/2024/544

[55] N. Kaluđerović, T. Kleinjung, and D. Kostić, "Cryptanalysis of the generalised legendre pseudorandom function," *Open Book Series*, vol. 4, no. 1, pp. 267–282, 2020.

[56] N. Kaluđerović, T. Kleinjung, and D. Kostic, "Improved key recovery on the legendre PRF," Cryptology ePrint Archive, Report 2020/098, 2020. [Online]. Available: https://eprint.iacr.org/2020/098

[57] D. Khovratovich, "Key recovery attacks on the Legendre PRFs within the birthday bound," Cryptology ePrint Archive, Report 2019/862, 2019. [Online]. Available: https://eprint.iacr.org/2019/862

[58] V. Kolesnikov, R. Kumaresan, M. Rosulek, and N. Trieu, "Efficient batched oblivious PRF with applications to private set intersection," in *ACM CCS 2016*, E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi, Eds. ACM Press, Oct. 2016, pp. 818–829.

[59] B. Kreuter, T. Lepoint, M. Orrù, and M. Raykova, "Anonymous tokens with private metadata bit," in *CRYPTO 2020, Part I*, ser. LNCS, D. Micciancio and T. Ristenpart, Eds., vol. 12170. Springer, Cham, Aug. 2020, pp. 308–336.

[60] H. Liu, X. Wang, K. Yang, and Y. Yu, "The hardness of LPN over any integer ring and field for PCG applications," in *EUROCRYPT 2024, Part VI*, ser. LNCS, M. Joye and G. Leander, Eds., vol. 14656. Springer, Cham, May 2024, pp. 149–179.

[61] L. Maino and C. Martindale, "An attack on SIDH with arbitrary starting curve," Cryptology ePrint Archive, Report 2022/1026, 2022. [Online]. Available: https://eprint.iacr.org/2022/1026

[62] D. Masny and P. Rindal, "Endemic oblivious transfer," in *ACM CCS 2019*, L. Cavallaro, J. Kinder, X. Wang, and J. Katz, Eds. ACM Press, Nov. 2019, pp. 309–326.

[63] C. Mauduit and A. Sárközy, "On finite pseudorandom binary sequences i: Measure of pseudorandomness, the legendre symbol," *Acta Arithmetica*, vol. 82, no. 4, pp. 365–377, 1997.

[64] A. May and F. Zweydinger, "Legendre PRF (multiple) key attacks and the power of preprocessing," in *CSF 2022 Computer Security Foundations Symposium*. IEEE Computer Society Press, Aug. 2022, pp. 428–438.

[65] M. Naor and O. Reingold, "Number-theoretic constructions of efficient pseudo-random functions," in *38th FOCS*. IEEE Computer Society Press, Oct. 1997, pp. 458–467.

[66] J. B. Nielsen, P. S. Nordholt, C. Orlandi, and S. S. Burra, "A new approach to practical active-secure two-party computation," in *CRYPTO 2012*, ser. LNCS, R. Safavi-Naini and R. Canetti, Eds., vol. 7417. Springer, Berlin, Heidelberg, Aug. 2012, pp. 681–700.

[67] R. Peralta, "On the distribution of quadratic residues and nonresidues modulo a prime number," *Mathematics of Computation*, vol. 58, no. 197, pp. 433–440, 1992.

[68] S. Pohlig and M. Hellman, "An improved algorithm for computing logarithms over gf (p) and its cryptographic significance (corresp.)," *IEEE Transactions on information Theory*, vol. 24, no. 1, pp. 106–110, 1978.

[69] M. O. Rabin, "How to exchange secrets with oblivious transfer," 1981.

[70] P. Rindal and L. Roy, "libOTe: an efficient, portable, and easy to use Oblivious Transfer Library," https://github.com/osu-crypto/libOTe.

[71] L. Roy, "SoftSpokenOT: Quieter OT extension from small-field silent VOLE in the minicrypt model," in *CRYPTO 2022, Part I*, ser. LNCS, Y. Dodis and T. Shrimpton, Eds., vol. 13507. Springer, Cham, Aug. 2022, pp. 657–687.

[72] A. Russell and I. E. Shparlinski, "Classical and quantum function reconstruction via character evaluation," *Journal of Complexity*, vol. 20, no. 2-3, pp. 404–422, 2004.

[73] J. T. Schwartz, "Fast probabilistic algorithms for verification of polynomial identities," *Journal of the ACM (JACM)*, vol. 27, no. 4, pp. 701–717, 1980.

[74] I. A. Seres, M. Horváth, and P. Burcsi, "The legendre pseudorandom function as a multivariate quadratic cryptosystem: security and applications," *Applicable Algebra in Engineering, Communication and Computing*, pp. 1–31, 2023.

[75] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *35th FOCS*. IEEE Computer Society Press, Nov. 1994, pp. 124–134.

[76] ——, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.

[77] Y. Sun, J. Katz, M. Raykova, P. Schoppmann, and X. Wang, "Actively secure private set intersection in the client-server setting," in *ACM CCS 2024*, B. Luo, X. Liao, J. Xu, E. Kirda, and D. Lie, Eds. ACM Press, Oct. 2024, pp. 1478–1492.

[78] V. Tóth, "Collision and avalanche effect in families of pseudorandom binary sequences," *Periodica Mathematica Hungarica*, vol. 55, pp. 185–196, 2007.

[79] W. van Dam and S. Hallgren, "Efficient quantum algorithms for shifted quadratic character problems," *arXiv preprint quant-ph/0011067*, 2000.

[80] X. Wang, A. J. Malozemoff, and J. Katz, "EMP-toolkit: Efficient MultiParty computation toolkit," https://github.com/emp-toolkit, 2016.

[81] C. Weng, K. Yang, J. Katz, and X. Wang, "Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits," in *2021 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 2021, pp. 1074–1091.

[82] K. Yang, P. Sarkar, C. Weng, and X. Wang, "QuickSilver: Efficient and affordable zero-knowledge proofs for circuits and polynomials over any field," in *ACM CCS 2021*, G. Vigna and E. Shi, Eds. ACM Press, Nov. 2021, pp. 2986–3001.

[83] Y. Yang, F. Benhamouda, S. Halevi, H. Krawczyk, and T. Rabin, "Gold OPRF: Post-quantum oblivious power residue PRF," Cryptology ePrint Archive, Paper 2024/1955, 2024. [Online]. Available: https://eprint.iacr.org/2024/1955

[84] A. C.-C. Yao, "How to generate and exchange secrets (extended abstract)," in *27th FOCS*. IEEE Computer Society Press, Oct. 1986, pp. 162–167.

[85] C. Zalka, "Grover's quantum searching algorithm is optimal," *Physical Review A*, vol. 60, no. 4, p. 2746, 1999.

[86] R. Zippel, "Probabilistic algorithms for sparse polynomials," in *International symposium on symbolic and algebraic manipulation*. Springer, 1979, pp. 216–226.

# Appendix A. Projected Overhead for Additional Properties

In this section, we discuss the overhead that will occur if additional properties are added to our implementation. We focus solely on the non-batched setting, as this overhead would be amortized in the batched setting.

**Server-specified Key.** This only incurs an additional $\mathbb{F}_p$ element (48 Bytes) to be transferred in the online phase.

**Randomized $\Delta^{(s)}$.** This only incurs an additional $\mathbb{F}_p$ element (48 Bytes) to be transferred in the offline phase. Recall that this can eliminate the one-bit leakage if we only need a uniformly sampled key; see Section 3.4.

**Avoiding Leakage for a Server-specified Key.** In the case of a server-specified key, we can eliminate the one-bit leakage by using a key with two leading zeros, as discussed in Section 3.4. This can be done by requiring (malicious) $\mathcal{P}_s$ to bit-decompose $[k]_{\Delta^{(c)}}$. To improve efficiency, $\mathcal{P}_s$ can decompose $k$ into 95 4-bit trunks and 2 1-bit trunks. Hence, it requires 97 VOLE correlations (each cost $\approx$ 3KB) and 97 derandomization over correlations (each cost 48 Bytes, i.e., sending $z - u$ to convert a random IT-MAC $[u]_{\Delta^{(c)}}$ to $[z]_{\Delta^{(c)}}$). In total, this gives a $\approx$ 296KB overhead. We also need to ensure that $\mathcal{P}_s$ uses values within the ranges $[0, 16)$ (and $[0, 2)$). This can be *freely* incorporated with the power-raising check as degree-16 (and 2) polynomials.

**Public Verifiability.** Recall that Gold over 7 public random inputs can be used as the public verification information (i.e., $\mathsf{VK}_7$ in Equation (1)) to achieve public verifiability. To ensure the key used is indeed in line with published $\mathsf{VK}_7$, each public input (i.e., $\ell_{i\in[7]}$) requires $\frac{\lambda}{\phi} + 1 = 33$ more VOLE correlations (each cost $\approx$ 3KB) and $\frac{\lambda}{\phi} = 32$ derandomization over the intermediate result (each cost 48 Bytes) and 1 field element for sending $r_i^{2^\lambda} \cdot (k + \ell_i)$; see Section 6.1. Hence, the total overhead for achieving public verifiability is $\approx$ 703KB.

**UC-Gold.** To achieve UC-Gold, we need to verify $\mathsf{VK}_6$ (which does not need to be public now) and also deploy leakage-free 2PC-Gold; see Section 6.1. Hence, the total overhead for achieving UC-Gold is $\approx$ 899KB ($\approx$ 603KB for verification of $\mathsf{VK}_6$ and $\approx$ 296KB for verification that the key has two leading zeros to remove leakage). This overhead is independent of $n$ and can be amortized.

## Appendix B. Meta-Review

The following meta-review was prepared by the program committee for the 2025 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

### B.1. Summary

This paper investigates the problem of constructing post-quantum oblivious PRFs, and proposes a simple OPRF protocol based on VOLE correlations and the classic power-residue PRF. The authors name their construction Gold, or the opposite of DLog. This protocol comes in two flavors that achieve either semi-honest or malicious security against the server (both achieve malicious security against the client), and extensions are presented to give it private or public verifiability, and to achieve security in the UC model. The scheme requires two rounds if the server is assumed to be semi-honest, or five if the server is malicious. It is distinguished from prior works that construct OPRF protocols from VOLE and the Legendre PRF in that it consumes standard VOLE correlations in a black-box way, which enables amortization in the batched evaluation setting. The authors implement their scheme and provide comparative benchmarks, demonstrating a significant performance improvement over most prior works, including VOLE/Legendre-based prior works in the batched setting, and a modest performance improvement relative to VOLE/Legendre-based prior works in the one-shot setting.

### B.2. Scientific Contributions

- Provides a Valuable Step Forward in an Established Field

### B.3. Reasons for Acceptance

The problem of post-quantum-secure (V)OPRFs is well-motivated by practice, and interesting theoretically. Relative to prior works constructing OPRFs from VOLE and the Legendre PRF, this work's black box use of the standard VOLE correlation enables dramatic concrete efficiency improvements in the batched/amortized setting, and modularizes the problem in such a way that performance improvements for the construction presented follow directly from future performance improvements in standard VOLE. Furthermore, the authors provide an explicit method for adding verifiability to their scheme, which was missing from comparable prior constructions. Relative to alternative approaches based upon lattices, this work demonstrates a concrete efficiency improvement of multiple orders of magnitude, which is significant for practice.

### B.4. Noteworthy Concerns

1) Care is required when using the 1-bit-leakage model in combination with sensitive data.
2) The hardness assumption on which the result relies is not so well studied as those used by prior constructions.