

Crypto Dark Matter on the Torus

Oblivious PRFs from shallow PRFs and TFHE

Martin R. Albrecht¹, Alex Davidson², Amit Deo³, and Daniel Gardham⁴

¹ King’s College London and SandboxAQ

² NOVA LINCS & DI, FCT, Universidade NOVA de Lisboa

³ Crypto Quantique and Zama

⁴ University of Surrey

Abstract. Partially Oblivious Pseudorandom Functions (POPRFs) are 2-party protocols that allow a client to learn pseudorandom function (PRF) evaluations on inputs of its choice from a server. The client submits two inputs, one public and one private. The security properties ensure that the server cannot learn the private input and the client cannot learn more than one evaluation per POPRF query. POPRFs have many applications including password-based key exchange and privacy-preserving authentication mechanisms. However, most constructions are based on classical assumptions, and those with post-quantum security suffer from large efficiency drawbacks.

In this work, we construct a novel POPRF from lattice assumptions and the “Crypto Dark Matter” PRF candidate (TCC’18) in the random oracle model. At a conceptual level, our scheme exploits the alignment of this family of PRF candidates, relying on mixed modulus computations, and programmable bootstrapping in the torus fully homomorphic encryption scheme (TFHE). We show that our construction achieves malicious client security based on circuit-private FHE, and client privacy from the semantic security of the FHE scheme. We further explore a heuristic approach to extend our scheme to support verifiability based on the difficulty of computing cheating circuits in low depth. This would yield a verifiable (P)OPRF. We provide a proof-of-concept implementation and benchmarks of our construction using the `tfhe-rs` software library. For the core online OPRF functionality, we require amortised 5.0kB communication per evaluation and a one-time per-client setup communication of 16.8MB.

Keywords: oblivious PRF, lattices, FHE

1 Introduction

Oblivious pseudorandom functions allow two parties to compute a pseudorandom function (PRF) $z := F_k(x)$ together: a server supplying a key k and a user supplying a private input x . The server does not learn x or z and the user does not learn k . If the user can be convinced that z is correct (i.e. that evaluation

is performed under the correct key), then the function is “verifiable oblivious” (VOPRF), otherwise it is only “oblivious” (OPRF). Both can be efficiently realised from Diffie-Hellman (DH) and may be used in many cryptographic applications. Example applications include anonymous credentials (e.g. Cloudflare’s [PrivacyPass](#) [DGS⁺18]) and Private Set Intersection (PSI) enabling e.g. privacy-preserving contact look-up on chat platforms [CHLR18].

The obliviousness property can be too strong in many applications where it is sufficient or even necessary to only hide part of the client’s input. In this case, the public and private inputs are separated by requiring an additional public input t , called the *tag*. Then we say that we have a *Partially* Oblivious PRF (POPRF). POPRFs are typically used in protocols where a server may wish to rate-limit OPRF evaluations made by a client. Such example protocols include Password-Authenticated Key Exchanges (e.g. [OPAQUE](#) [JKX18], which is in the process of Internet Engineering Task Force (IETF) standardisation) and the Pythia PRF service [ECS⁺15]. This latter work also proposed a bilinear pairing-based construction of a *Verifiable* POPRF (VPOPRF), which is the natural inclusion of both properties: some of the input is revealed to the server and the client is able to check the correct evaluation of its full input.

Despite the wide use of (VP)OPRFs, most constructions are based on classical assumptions, such as DH, RSA or even pairing-based assumptions. The latest in this line of research is a recent VPOPRF construction based on a novel DH-like assumption [TCR⁺22] and DH-based OPRFs are currently being [standardised](#) by the IETF. Their vulnerability to quantum adversaries makes it desirable to find post-quantum solutions. However, known candidates are much less efficient.

Given fully homomorphic encryption (FHE), there is a natural (P)OPRF candidate. The client FHE encrypts input x and sends it with tag t . The server then evaluates the PRF homomorphically or “blindly” using a key derived from t and its own secret key. Finally, the client decrypts the resulting ciphertext to obtain the PRF output. The first challenge with this approach is performance, PRFs tend to have sufficiently deep circuits that FHE schemes struggle to evaluate them efficiently. Even special purpose PRFs such as the LowMC construction [ARS⁺15] require depth ten or more, making them somewhat impractical. More generally, in a binary circuit model we expect to require depth $\Theta(\log \lambda)$ to obtain a PRF resisting attacks with complexity $2^{\Theta(\lambda)}$.

Yet, if we expand our circuit model to arithmetic circuits with both mod p and mod q gates for $p \neq q$ both primes, shallow proposals exist [BIP⁺18,DGH⁺21]. In particular, the (weak) PRF candidate in [BIP⁺18] is

$$z := \sum (\mathbf{A} \cdot \mathbf{x} \bmod 2) \bmod 3$$

where arithmetic operations are over the Integers and \mathbf{A} is the secret key. The same work also contains a proposal to “upgrade” this weak PRF, defined for uniformly random inputs \mathbf{x} , to a full PRF, taking any \mathbf{x} . Furthermore, the works [BIP⁺18,DGH⁺21] already provide oblivious PRF candidates based on this PRF and MPC, but with non-optimal round complexity. Thus, a natural

question to ask is if we can construct a round-optimal (or, 2 message) POPRF based on this PRF candidate using the FHE-based paradigm mentioned above.

1.1 Contributions

Our starting point is the observation that the computational model in [BIP⁺18] aligns well with that of the TFHE encryption scheme [CGGI20] and its “programmable bootstrapping” technique [MP20,Joy21]. Programmable bootstrapping allows us to realise arbitrary, not necessarily low degree, small look-up tables and thus function evaluations on (natively) single inputs. Thus, it is well positioned to realise the required gates.⁵ Indeed, FHE schemes natively compute plaintexts modulo some $P \in \mathbb{Z}$ and programmable bootstrapping allows us to switch between these plaintext moduli, e.g. from mod P_1 to mod P_2 . This implies a weak PRF with two levels of bootstrapping only.⁶ We believe this simple observation and conceptual contribution will have applications beyond this work. We further hope that by giving another application domain for the PRF candidate from [BIP⁺18] – it is not just MPC-friendly but also FHE-friendly – we encourage further cryptanalysis on it.

After some preliminaries in Section 2 we specify our POPRF candidate in Section 3. As is typical with FHE-based schemes, we require the involved parties – here the client – to prove that its inputs are well-formed. We also make use of the protected encoded-input PRF (PEI-PRF) paradigm from [BIP⁺18] where the client performs some computations not dependent on secret key material and then submits the output together with a NIZK proof of well-formedness to the server for processing.

We prove our construction secure in the random oracle model in Section 4. We show that our construction meets the security definitions from [TCR⁺22]: pseudorandomness even in the presence of malicious clients (POPRF security) and privacy for clients. This latter property has two flavours based on the capabilities of the adversary, POPRIV1 (which we achieve) captures security against an honest-but-curious server whereas POPRIV2 ensures security even when the server is malicious. Here, the client maintains privacy by detecting malicious behaviour of the server. POPRF security for the server essentially rests on a variant of circuit-privacy obtained from TFHE bootstrapping and client NIZK. The NIZK is made online extractable in the POPRF proof using a trapdoor and thus avoids any rewinding issues outlined in e.g. [SG98], and similarly mitigates the problem of rewinding for post-quantum security, cf. [Unr12]. POPRIV1 security for the client against a semi-honest server essentially relies on the IND-CPA security of TFHE.

Initially, we focus on oblivious rather than verifiable oblivious PRFs. This is motivated by the presumed high cost associated with zero-knowledge proofs

⁵ The security of the PRF candidate in [BIP⁺18] rests on the absence of any low-degree polynomial interpolating it, ruling out efficient implementations using FHE schemes that only provide additions and multiplications.

⁶ We require two, not one, levels because we appeal to the circuit-privacy properties of a final bootstrapping operation.

for performing FHE computations. In Section 5, we explore a different approach to adding verifiability to our OPRF, inspired by and based on a discussion in [ADDS21]. The idea here is that the server commits to a set of evaluation “check” points and that the client can use the oblivious nature of the PRF to request PRF evaluations of these points to catch a cheating server. However, achieving security of this “cut-and-choose” approach in this setting is non-trivial as the server may still obliviously run a cheating circuit that agrees on those check points but diverges elsewhere.

We explore the feasibility of such a cheating circuit using direct cryptanalysis. In more detail, inspired by the heuristic approach in [CHLR18] for achieving malicious security – forcing the server to compute a deep circuit in FHE parameters supporting only shallow circuits – we explore cheating circuits in bootstrapping depth two. While we were unable to find such a cheating circuit, and conjecture that one does not exist, we stress that this part of our work is highly speculative. Under the heuristic assumption that our construction is verifiable, in Appendix C we then show that it also satisfies POPRIV2. We hope that our work encourages further exploration of such strategies, as these will have applications elsewhere to upgrade FHE-based schemes to malicious security.⁷

We present our proof-of-concept Rust implementation in Section 6, and provide benchmarks that illustrate how quickly the core OPRF functionality (without (verifying the) NIZK proofs) can be run. Our implementation is open-source⁸ and makes use of Zama’s `tfhe-rs` library for implementing TFHE. While the public key material sent by the client to the server is large (16.8MB) this cost can be amortised by reusing the same material for several evaluations. Individual PRF evaluations can then cost about 30.4kB or as little as 5.0kB when amortising client NIZK proofs across several OPRF queries. In terms of runtimes, client online functions run in 36ms on one core and server online functions are expected to run in 123ms on 64 cores. We discuss the (significantly worse) performance of our VOPRF in Section 5.

In Appendix B we then estimate costs of the required non-interactive zero-knowledge proofs. First, we use [BS22] (with a zero-knowledge “shim”) to proof well-formedness of the bootstrapping key material. Second, as previously mentioned, our OPRF construction makes use of the PEI-PRF paradigm [BIP⁺18] which must be made compatible with our zero-knowledge proofs. Here we rely on and specialise the proofs in [LNP22]. We also show this is extendable to the stronger property of *non-unique encoding* required in our verifiable OPRF.

1.2 Related Work

Oblivious PRFs and variants thereof are an active area of research. A survey of constructions, variants and applications was given in [CHL22]. In this work we are interested in plausibly post-quantum and round-optimal constructions.

⁷ We note that a hybrid approach of ours and proving correct evaluation would be to prove that only a limited-depth circuit was evaluated.

⁸ <https://anonymous.4open.science/r/oprf-fhe-C726>

The first candidate construction was given in [ADDS21], which built a verifiable oblivious PRF from lattice assumptions following the blueprint of Diffie-Hellman constructions with additive blinding (a construction for multiplicative blinding is given in an appendix of the full version of [ADDS21]). The work provides both semi-honest and malicious secure candidates with the latter being significantly more expensive. We stress that in the former, both parties are semi-honest.

In [BKW20] two candidate constructions from isogenies were proposed. One, a VOPRF related to SIDH, was unfortunately shown to not be secure [BKM⁺21]. The other, an OPRF related to CSIDH, achieves sub megabyte communication in a malicious setting assuming the security of group-action decisional Diffie-Hellman. In [Bas23] a fixed-and-improved SIDH-based candidate was proposed and in [HMR23] an improved CSIDH-based candidate is presented. Either of these rely on trusted setups. In [SHB21] an OPRF based on the Legendre PRF is proposed based on solving sparse multi-variate quadratic systems of equations. In [DGH⁺21], which also builds on [BIP⁺18], an MPC-based OPRF is proposed that is secure against semi-honest adversaries. It achieves much smaller communication complexity compared to all other post-quantum candidates, but in a preprocessing model where correlated randomness is available to the parties. A protocol computing this correlated randomness, e.g. [BCG⁺22], would add two rounds (or more) and thus make the overall protocol not round-optimal. The question of upgrading security to full malicious security is left as an open problem in [DGH⁺21].⁹

We give a summary comparison of our construction with prior work in Table 1. The only 2-round construction without preprocessing or trusted-setup in Table 1 is that from [ADDS21], where our construction compares favourably by offering stronger security at smaller size. In particular, even in a semi-honest setting, our construction outperforms that from [ADDS21] in terms of bandwidth for $L = 16$ queries.¹⁰

1.3 Open Problems

The most immediate open problem posed by this work is to provide a full, end-to-end, implementation of our OPRF candidate. As we will discuss in Section 6, the `tfhe-rs` API does not enable us to obtain full performance in terms of server running time. Furthermore, we did not implement the NIZK proofs attesting well-formedness of inputs.

Another pressing open problem is to refine our understanding of the security of the PRF candidate from [BIP⁺18]. In particular, our parameter choices may prove to be too aggressive, and we hope that our work inspires cryptanalysis.

Our verifiability approach throws up a range of interesting avenues to explore for VOPRFs but also for verifiable homomorphic computation, more generally.

⁹ We note that [DGH⁺21] may also serve as an indication for the best-case performance of OPRFs built from generic MPC techniques.

¹⁰ We note that while the large sizes for achieving malicious security in [ADDS21] can be avoided using improved NIZKs, the semi-honest base size of 2MB per query stems from requiring $q \approx 2^{256}$ for statistical correctness and security arguments.

Table 1. Post-quantum (P)OPRF candidates in the literature

work	assumption	r	communication cost	flavour	model
[ADDS21]	R(LWE) & SIS	2	$\approx 2\text{MB}$	plain	semi-honest, QROM
[SHB21]	Legendre PRF	3	$\approx \lambda \cdot 13\text{K}$	plain	semi-honest, <i>pp</i> , ROM
[BKW20]	CSIDH	3	424KB	plain	malicious client
[Bas23]	SIDH	2	3.0MB	plain	malicious, <i>ts</i> , ROM
[HMR23]	CSIDH	2	21KB	plain	semi-honest, <i>ts</i>
[HMR23]	CSIDH	4	35KB	plain	malicious client, <i>ts</i>
[HMR23]	CSIDH	258	25KB	plain	semi-honest
[DGH ⁺ 21]	[BIP ⁺ 18]	2	80B	plain	semi-honest, <i>pp</i>
Section 3	lattices, [BIP ⁺ 18]	2	16.8MB + 41.42KB + 0.5kB + 26.7kB + 3.2kB	plain	malicious client, ROM
Section 3	lattices, [BIP ⁺ 18]	2	16.8MB + 41.42KB + 0.5kB + 1.3kB + 3.2kB	plain	malicious client, ROM $L = 64$, per query

The column “r” gives the number of rounds. ROM is the random oracle model, QROM the quantum random oracle model, “pp” stands for “preprocessing”, and “ts” for “trusted setup”. When reporting on our work, the summands are: pk size, pk proof size, client message size, client message proof size, server message size. Our client message proofs can be amortised to e.g. $83.4\text{kB}/64 = 1.3\text{kB}$ per query, when amortising over $L = 64$ queries.

First, our OPRF construction relies on programmable bootstrapping. This restricts the choice of FHE scheme we might instantiate our protocol with, but also gives the server the choice which function to evaluate, something our application does not require. That is, we may not need to rely on *evaluator programmable bootstrapping* if it is possible for the client to define the non-linear functions available to a server (*encrypter programmable bootstrapping*). This would enable reasoning about malicious server security more easily.

Related works, e.g. [CHLR18], have also used similar assumptions as our work over the hardness of computing deep circuits in low FHE depth. There is growing evidence that such assumptions allow for new, interesting or more efficient constructions of cryptographic primitives. However, the hardness of these computational problems needs to be better understood.

Finally, we estimate that our VOPRF is orders of magnitude less efficient than our OPRF candidate and we discuss plausible avenues for alleviating that difference in Section 5. A more direct approach would be to construct a NIZK for correct bootstrapping evaluation, which would have applications beyond this work.

2 Preliminaries

We use $\lfloor \cdot \rfloor$, $\lceil \cdot \rceil$ and $\lceil \cdot \rceil$ to denote the standard floor, ceiling and rounding to the nearest integer functions (rounding down in the case of a tie). We denote the Integers by \mathbb{Z} and for any positive $p \in \mathbb{Z}$, the integers modulo p are denoted by \mathbb{Z}_p . We typically use representatives of \mathbb{Z}_p in $\{-p/2, \dots, (p/2) - 1\}$ if p is even and $\{-\lfloor p/2 \rfloor, \dots, \lfloor p/2 \rfloor\}$ if p is odd, but we will also consider \mathbb{Z}_p as $\{0, 1, \dots, p - 1\}$.

Since it will always be clear from context or stated explicitly which representation we use, this does not create ambiguity. The p -adic decomposition of an integer $x \geq 0$ is a tuple $(x_i)_{0 \leq i < \lceil \log_p(x) \rceil}$ with $0 \leq x_i < p$ such that $x = \sum p^i \cdot x_i$. We denote the set S_m to be the permutation group of m elements.

Let $\mathbb{Z}[X]$ denote the polynomial ring in the variable X whose coefficients belong to \mathbb{Z} . We also denote power-of-two cyclotomic rings $\mathcal{R} := \mathbb{Z}[X]/(X^d + 1)$ where d is a power-of-two, and $\mathcal{R}_q := \mathcal{R}/(q\mathcal{R})$ for any integer “modulus” q . Bold letters denote vectors and upper case letters denote matrices. Abusing notation we write (\mathbf{x}, \mathbf{y}) for the concatenation of the vectors \mathbf{x} and \mathbf{y} . We extend this notation to scalars, too. Additionally, $\|\cdot\|$ and $\|\cdot\|_\infty$ denote standard Euclidean and infinity norms respectively.

For a distribution D , we write $x \leftarrow \$ D$ to denote that x is sampled according to the distribution D . An example of a distribution is the discrete Gaussian distribution over \mathbb{Z} with parameter $\sigma > 0$ denoted as $D_{\mathbb{Z}, \sigma}$. This distribution has its probability mass function proportional to the Gaussian function $\rho_\sigma(x) := \exp(-\pi x^2/\sigma^2)$. We use λ to denote the security parameter. We use the standard asymptotic notation ($\Omega, \mathcal{O}, \omega$ etc.) and use $\text{negl}(\lambda)$ to denote a negligible function, i.e. a function that is $\lambda^{\omega(1)}$. Further, we write $\text{poly}(\lambda)$ to denote a polynomial function i.e. a function that is $\mathcal{O}(n^c)$ for some constant c . An algorithm is said to be polynomially bounded if it terminates after $\text{poly}(\lambda)$ steps and uses $\text{poly}(\lambda)$ -sized memory. Two distribution ensembles $D_1(1^\lambda)$ and $D_2(1^\lambda)$ are said to be *computationally* indistinguishable if for any probabilistic polynomially bounded algorithm \mathcal{A} , $\text{Adv}(\mathcal{A}) := \|\Pr[1 \leftarrow \$ \mathcal{A}_X(1^\lambda)] - \Pr[1 \leftarrow \$ \mathcal{A}_Y(1^\lambda)]\| \leq \text{negl}(\lambda)$. In such a case we write $D_1(1^\lambda) \approx_c D_2(1^\lambda)$. The distribution ensembles are said to be *statistically* indistinguishable if the same holds for all unbounded algorithms, in which case we write $D_1(1^\lambda) \approx_s D_2(1^\lambda)$.

For a keyspace \mathcal{K} , input space \mathcal{X} and output space \mathcal{Z} , a PRF is a function $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Z}$ with a pseudorandomness property. Rather than writing $F(k, x)$ for $k \in \mathcal{K}$ and $x \in \mathcal{X}$, we write $F_k(x)$. The pseudorandomness property of a PRF requires that over a secret and random choice of $k \leftarrow \$ \mathcal{K}$, the single input function $F_k(\cdot)$ is computationally indistinguishable from a uniformly random function. Note here that the dependence of the parameters on λ is present, but is not explicitly written for simplicity. We also use the standard cryptographic notion of a (non-interactive) zero-knowledge proof/argument. For more details on these standard cryptographic notions, see e.g. [Gol04].

2.1 Random Oracle Model

We will prove security by modelling hash functions as random oracles. Since our schemes will make use of more than one hash function, it will be useful to have a general abstraction for the use of ideal primitives, following the treatment in [TCR⁺22]. A random oracle RO specifies algorithms RO.Init and RO.Eval. The initialisation algorithm has syntax $st_{\text{RO}} \leftarrow \$ \text{RO.Init}(1^\lambda)$. The stateful evaluation algorithm has syntax $y \leftarrow \$ \text{RO.Eval}(x, st_{\text{RO}})$. We sometimes use A^{RO} as shorthand for giving algorithm A oracle access to $\text{RO.Eval}(\cdot, st_{\text{RO}})$. We combine access to multiple random oracles $\text{RO} = \text{RO}_0 \times \dots \times \text{RO}_{m-1}$ in the obvious way. We may

arbitrarily label our random oracles to aid readability e.g. RO_{key} to denote a random oracle applied to some “key”.

2.2 (Verifiable) (Partial) Oblivious Pseudorandom Functions

We adopt the notation and definitions for oblivious pseudorandom functions from [TCR⁺22]. An OPRF is a protocol between two parties: a server S who holds a private key and a client who wants to obtain evaluations of F_k on inputs of its choice. We write $z := F_k(x)$. We say that an OPRF is a partial OPRF (OPRF) if part of the client’s input is given to the server. In this case, we write $z := F_k(t, x)$ where t is in the clear and x is hidden from S . When C can verify that the PRF was evaluated correctly we speak of a verifiable OPRF (VOPRF) or VPOPRF when the protocol also supports partially known inputs t .

Definition 1 (Partial Oblivious PRF [TCR⁺22]). *A partial oblivious PRF (POPRF) \mathcal{F} is a tuple of PPT algorithms*

$$(\mathcal{F}.\text{Setup}, \mathcal{F}.\text{KeyGen}, \mathcal{F}.\text{Request}, \mathcal{F}.\text{BlindEval}, \mathcal{F}.\text{Finalise}, \mathcal{F}.\text{Eval})$$

The setup and key generation algorithm generate public parameter pp and a public/secret key pair (pk, sk) . Oblivious evaluation is carried out as an interactive protocol between C and S , here presented as algorithms $\mathcal{F}.\text{Request}$, $\mathcal{F}.\text{BlindEval}$, $\mathcal{F}.\text{Finalise}$ working as follows:

1. *First, C runs the algorithm $\mathcal{F}.\text{Request}_{\text{pp}}^{\text{RO}}(\text{pk}, t, x)$ taking a public key pk , a tag or public input t and a private input x . It outputs a local state st and a request message req , which is sent to the server.*
2. *S runs $\mathcal{F}.\text{BlindEval}_{\text{pp}}^{\text{RO}}(\text{sk}, t, \text{req})$ taking as input a secret key sk , a tag t and the request message req . It produces a response message rep sent back to C .*
3. *Finally, C runs $\mathcal{F}.\text{Finalise}(\text{rep}, st)$ which takes the response message and its previously constructed state st and either outputs a PRF evaluation or \perp if rep is rejected.*

The unblinded evaluation algorithm $\mathcal{F}.\text{Eval}$ is deterministic and takes as input a secret key sk , an input pair (t, x) and outputs a PRF evaluation z .

We also define sets $\mathcal{F}.\text{SK}$, $\mathcal{F}.\text{PK}$, $\mathcal{F}.\text{T}$, $\mathcal{F}.\text{X}$ and $\mathcal{F}.\text{Out}$ representing the secret key, public key, tag, private input, and output space, respectively. We define the input space $\mathcal{F}.\text{In} = \mathcal{F}.\text{T} \times \mathcal{F}.\text{X}$. We assume efficient algorithms for sampling and membership queries on these sets.

Remark 1. Fixing t , e.g. $t = \perp$, recovers the definition of an OPRF.

Remark 2. In Figure 1, the oracle $\text{Prim}(x)$ captures access to the random oracle used in the POPRF construction. For $b = 0$ (the case where the adversary interacts with a simulator and truly random function) the simulator may only use a limited number of random function queries to simulate the random oracle accessed via $\text{Prim}(x)$.

Game $\text{POPRF}_{\mathcal{F},\mathcal{S},\text{RO}}^{A,b}(\lambda)$	Oracle $\text{Eval}(t,x)$	Oracle $\text{BlindEval}(t,req)$
$q_{s,t}, q_t \leftarrow 0,0$	$z_0 \leftarrow \text{RO}_{\text{Fn}}.\text{Eval}((t,x),st_{\text{Fn}})$	$q_t \leftarrow q_t + 1$
$st_{\text{Fn}} \leftarrow \mathcal{S}.\text{RO}_{\text{Fn}}.\text{Init}(1^\lambda) \quad // \mathcal{F}.\text{In} \rightarrow \mathcal{F}.\text{Out}$	$z_1 \leftarrow \mathcal{F}.\text{Eval}_{\text{pp}_1}^{\text{RO}}(sk,t,x)$	$(rep_0, st_S) \leftarrow \mathcal{S}.\text{BlindEval}^{\text{LimitEval}}(t,req,st_S)$
$st_{\text{RO}} \leftarrow \mathcal{R}.\text{Init}(1^\lambda)$	return z_b	$rep_1 \leftarrow \mathcal{F}.\text{BlindEval}_{\text{pp}_1}^{\text{RO}}(sk,t,req)$
$\text{pp}_1 \leftarrow \mathcal{F}.\text{Setup}(1^\lambda)$		return rep_b
$(st_S, \text{pk}_0, \text{pp}_0) \leftarrow \mathcal{S}.\text{Init}(\text{pp}_1)$	Oracle $\text{LimitEval}(t,x)$	
$(sk, \text{pk}_1) \leftarrow \mathcal{F}.\text{KeyGen}_{\text{pp}_1}^{\text{RO}}(1^\lambda)$	$q_{t,s} \leftarrow q_{t,s} + 1$	Oracle $\text{Prim}(x)$
$b' \leftarrow \mathcal{A}^{\text{Eval,BlindEval,Prim}}(\text{pp}_b, \text{pk}_b)$	if $q_{t,s} \leq q_t$ then	$(h_0, st_S) \leftarrow \mathcal{S}.\text{Eval}^{\text{LimitEval}}(x, st_S)$
return b'	return $\text{Eval}(t,x)$	$h_1 \leftarrow \mathcal{R}.\text{Eval}(x, st_{\text{RO}})$
	return \perp	return h_b

Fig. 1. Pseudorandomness against malicious clients.

We adapt the correctness notion from [TCR⁺22], permitting a small failure probability.

Definition 2 (POPRF Correctness (adapted from [TCR⁺22])). A partial oblivious PRF (POPRF)

$$(\mathcal{F}.\text{Setup}, \mathcal{F}.\text{KeyGen}, \mathcal{F}.\text{Request}, \mathcal{F}.\text{BlindEval}, \mathcal{F}.\text{Finalise}, \mathcal{F}.\text{Eval})$$

is correct if

$$\Pr \left[z = \mathcal{F}.\text{Eval}_{\text{pp}}^{\text{RO}}(sk, t, x) \mid \begin{array}{l} \text{pp} \leftarrow \mathcal{F}.\text{Setup}(1^\lambda) \\ (\text{pk}, \text{sk}) \leftarrow \mathcal{F}.\text{KeyGen}_{\text{pp}}^{\text{RO}}(1^\lambda) \\ (st, req) \leftarrow \mathcal{F}.\text{Request}_{\text{pp}}^{\text{RO}}(\text{pk}, t, x) \\ rep \leftarrow \mathcal{F}.\text{BlindEval}_{\text{pp}}^{\text{RO}}(sk, t, req) \\ z \leftarrow \mathcal{F}.\text{Finalise}_{\text{pp}}^{\text{RO}}(rep, st) \end{array} \right] = 1 - \text{negl}(\lambda).$$

We target the same pseudorandomness guarantees against malicious clients as [TCR⁺22].

Definition 3 (Pseudorandomness (POPRF) [TCR⁺22]). We say a partial oblivious PRF \mathcal{F} is pseudorandom if for all PPT adversaries \mathcal{A} , there exists a PPT simulator \mathcal{S} such that the following advantage is $\text{negl}(\lambda)$:

$$\text{Adv}_{\mathcal{F},\mathcal{S},\text{RO},\mathcal{A}}^{\text{po-prf}}(\lambda) = \left| \Pr \left[\text{POPRF}_{\mathcal{F},\mathcal{S},\text{RO}}^{A,1}(\lambda) \Rightarrow 1 \right] - \Pr \left[\text{POPRF}_{\mathcal{F},\mathcal{S},\text{RO}}^{A,0}(\lambda) \Rightarrow 1 \right] \right|.$$

Remark 3. The intuition of this definition is that it requires the simulator to explain a random output (defined via RO_{Fn}) as an evaluation point of the PRF. The simulator provides its own public key and public parameters, but it gets at most one query to $\text{RO}_{\text{Fn}}()$ per BlindEval query that it has to simulate. The simulator queries RO_{Fn} through calls to LimitEval , where the check $q_{t,s} \leq q_t$ enforces the number of queries per BlindEval query and tag t . This implies that BlindEval and Eval queries essentially leak nothing beyond the evaluation at this exact point to the client.

<p>Game $\text{POPRIV1}_{\mathcal{F},\text{RO}}^{\mathcal{A},b}(\lambda)$</p> <hr/> <p>$\text{pp} \leftarrow \mathcal{F}.\text{Setup}(1^\lambda)$ $(\text{pk}, \text{sk}) \leftarrow \mathcal{F}.\text{KeyGen}_{\text{pp}}^{\text{RO}}(1^\lambda)$ $b' \leftarrow \mathcal{A}^{\text{Run,RO}}(\text{pp}, \text{pk}, \text{sk})$ return b'</p> <hr/> <p>Oracle $\text{Run}(t, x_0, x_1)$</p> <hr/> <p>for $j \in \{0, 1\}$ do $(st_j, req_j) \leftarrow \mathcal{F}.\text{Request}_{\text{pp}}^{\text{RO}}(\text{pk}, t, x_j)$ $rep_j \leftarrow \mathcal{F}.\text{BlindEval}_{\text{pp}}^{\text{RO}}(\text{sk}, t, req_j)$ $z_j \leftarrow \mathcal{F}.\text{Finalise}_{\text{pp}}^{\text{RO}}(rep_j, st_j)$ $\tau_0 \leftarrow (req_b, rep_b, z_0)$ $\tau_1 \leftarrow (req_{1-b}, rep_{1-b}, z_1)$ return (τ_0, τ_1)</p>	<p>Game $\text{POPRIV2}_{\mathcal{F},\text{RO}}^{\mathcal{A},b}(\lambda)$</p> <hr/> <p>$\text{pp} \leftarrow \mathcal{F}.\text{Setup}(1^\lambda)$ $i \leftarrow 0$ $b' \leftarrow \mathcal{A}^{\text{Request,Finalise,RO}}(\text{pp})$ return b'</p> <hr/> <p>Oracle $\text{Request}(\text{pk}, t, x_0, x_1)$</p> <hr/> <p>$i \leftarrow i + 1$ $(st_{i,0}, req_0) \leftarrow \mathcal{F}.\text{Request}_{\text{pp}}^{\text{RO}}(\text{pk}, t, x_0)$ $(st_{i,1}, req_1) \leftarrow \mathcal{F}.\text{Request}_{\text{pp}}^{\text{RO}}(\text{pk}, t, x_1)$ return (req_b, req_{1-b})</p> <hr/> <p>Oracle $\text{Finalise}(j, rep_0, rep_1)$</p> <hr/> <p>if $j > i$ then return \perp $z_b \leftarrow \mathcal{F}.\text{Finalise}_{\text{pp}}^{\text{RO}}(rep_0, st_{j,b})$ $z_{1-b} \leftarrow \mathcal{F}.\text{Finalise}_{\text{pp}}^{\text{RO}}(rep_1, st_{j,1-b})$ if $z_0 = \perp$ or $z_1 = \perp$ then return \perp return (z_0, z_1)</p>
--	--

Fig. 2. Request privacy against honest-but-curious servers (left) and against malicious servers (right).

Moreover, the simulator is restricted in that LimitEval oracle will error if more queries are made to it than the number of BlindEval queries (on t) at any point in the game. Meaningful relaxations of this definition are discussed in [TCR⁺22] but for completeness we opt for the full definition.

Definition 4 (Request Privacy (POPRIV) [TCR⁺22]). *We say a partial oblivious PRF \mathcal{F} has request privacy against honest-but-curious and malicious servers respectively if for all PPT adversary \mathcal{A} the following advantage is $\text{negl}(\lambda)$ for $k = 1$ and $k = 2$ respectively:*

$$\text{Adv}_{\mathcal{F},\text{S,RO},\mathcal{A}}^{\text{po-priv}^k}(\lambda) = \left| \Pr \left[\text{POPRIV } k_{\mathcal{F},\text{RO}}^{\mathcal{A},1}(\lambda) \Rightarrow 1 \right] - \Pr \left[\text{POPRIV } k_{\mathcal{F},\text{RO}}^{\mathcal{A},0}(\lambda) \Rightarrow 1 \right] \right|.$$

2.3 Hard Lattice Problems

We will rely on both the M -SIS and the M -LWE problems. Instantiating these over $\mathcal{R} = \mathbb{Z}$ recovers the SIS and LWE problems respectively. Further, instantiating these over some ring of integer of some number field and with $n = 1$, recovers the Ring-SIS and Ring-LWE problems respectively.

Definition 5 (M -SIS, adapted from [LS15]). *Let $\mathcal{R}, q, n, \ell, \beta$ depend on λ . The Module-SIS (or M -SIS) problem, denoted $M\text{-SIS}_{\mathcal{R},q,n,\ell,\beta^*}$, is: Given a uniform $\mathbf{A} \leftarrow \mathcal{R}_q^{n \times \ell}$ find some $\mathbf{u} \neq \mathbf{0} \in \mathcal{R}^\ell$ such that $\|\mathbf{u}\| \leq \beta^*$ and $\mathbf{A} \cdot \mathbf{u} \equiv \mathbf{0} \pmod{q}$.*

Definition 6 (M -LWE, adapted from [LS15]). *Let \mathcal{R}, q, n, m depend on λ and let χ_s, χ_e be distributions over \mathcal{R}_q . Denote by $M\text{-LWE}_{\mathcal{R},q,n,m,\chi_s,\chi_e}$ the*

probability distribution on $\mathcal{R}_q^{m \times n} \times \mathcal{R}_q^m$ obtained by sampling the coordinates of the matrix $\mathbf{A} \in \mathcal{R}_q^{m \times n}$ independently and uniformly over \mathcal{R}_q , sampling the coordinates of $\mathbf{s} \in \mathcal{R}_q^n$, $\mathbf{e} \in \mathcal{R}^m$ independently from χ_s and χ_e respectively, setting $\mathbf{b} := \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \bmod q$ and outputting (\mathbf{A}, \mathbf{b}) . The M -LWE problem is to distinguish the uniform distribution over $\mathcal{R}_q^{m \times n} \times \mathcal{R}_q^m$ from M -LWE $_{\mathcal{R}_q, n, m, \chi_s, \chi_e}$.

2.4 Matrix NTRU Trapdoors

The original formulation [HPS96] of the NTRU problem considers rings of integers of number fields or polynomial rings, but a matrix version is implicit and considered for cryptanalysis in the literature.

Definition 7. Given integers n, p, q, β where p and q are coprime, the matrix-NTRU assumption (denoted $\text{mat-NTRU}_{n, p, q, \beta}$) states that no PPT algorithm can distinguish between \mathbf{A} and \mathbf{B} where

- $\mathbf{A} \leftarrow_{\$} \mathbb{Z}_q^{n \times n}$
- $\mathbf{B} = p^{-1} \cdot \mathbf{G}^{-1} \cdot \mathbf{F} \bmod q$ with

$$\mathbf{F} \leftarrow_{\$} \{0, \pm 1, \dots, \pm \beta\}^{n \times n}, \mathbf{G} \leftarrow_{\$} \{0, \pm 1, \dots, \pm \beta\}^{n \times n} \cap (\mathbb{Z}_q^{n \times n})^*$$

where $(\mathbb{Z}_q^{n \times n})^*$ denotes the set of invertible $(n \times n)$ matrices over \mathbb{Z}_q .

We will use the matrix-NTRU assumption to define a trapdoor. In what follows, we assume an odd q and an even p that is coprime to q . In particular, we define the following algorithms:

$\text{NTRUTrapGen}(n, q, p, \beta)$: Sample

$$\mathbf{F} \leftarrow_{\$} \{0, \pm 1, \dots, \pm \beta\}^{n \times n}, \mathbf{G} \leftarrow_{\$} \{0, \pm 1, \dots, \pm \beta\}^{n \times n} \cap (\mathbb{Z}_q^{n \times n})^*$$

and output public information $\text{pp} := (p^{-1} \cdot \mathbf{G}^{-1} \cdot \mathbf{F} \bmod q, q)$ and a trapdoor $\tau := (\mathbf{F}, \mathbf{G}, p)$.

$\text{NTRUDec}(\mathbf{c}, \tau)$: For $\mathbf{c} \in \mathbb{Z}_q^n$, $\tau := (\mathbf{F}, \mathbf{G}, p)$, compute $\mathbf{c}_1 = p \cdot \mathbf{G} \cdot \mathbf{c} \bmod q$, $\mathbf{c}_2 = \mathbf{c}_1 \bmod (p/2)$, $\mathbf{c}_3 = \mathbf{c} - p^{-1} \cdot \mathbf{G}^{-1} \cdot \mathbf{c}_2 \bmod q$. Finally, compute and output $\mathbf{m}' := \lfloor \frac{2}{q-1} \cdot \mathbf{c}_3 \rfloor$ where the multiplication and rounding is done over the rationals.

The trapdoor functionality is summarised in the lemma below.

Lemma 1. Suppose that p, q are coprime where p is even and q is odd. Suppose also that $\beta \cdot \beta'_s \cdot n < p/4$, and that $\beta'_s, \beta'_e \in \mathbb{R}$ satisfies $\beta \cdot n \cdot (\beta'_s + p \cdot (2\beta'_e + 1)/2) < q/2$. Sample $(\text{pp} := (\mathbf{B}, q), \tau) \leftarrow_{\$} \text{NTRUTrapGen}(n, q, p, \beta)$. Then:

1. \mathbf{B} is indistinguishable from uniform over $\mathbb{Z}_q^{n \times n}$ if the $\text{mat-NTRU}_{n, p, q, \beta}$ assumption holds.
2. If $\mathbf{c} = \mathbf{B} \cdot \mathbf{s} + \mathbf{e} + \lfloor q/2 \rfloor \cdot \mathbf{m} \bmod q$ where $\mathbf{m} \in \mathbb{Z}_2^n$, $(\|\mathbf{s}\|_\infty \leq \beta'_s \vee \|\mathbf{s}\|_2 \leq \beta'_s \cdot \sqrt{n})$ and $(\|\mathbf{e}\|_\infty \leq \beta'_e \vee \|\mathbf{e}\|_e \leq \beta'_e \cdot \sqrt{n})$, then $\text{NTRUDec}(\mathbf{c}, \tau) = \mathbf{m}$.

Proof. For the first part, simply note that distinguishing \mathbf{B} from uniform is exactly the matrix-NTRU problem for (n, p, q, β) . For the second part, reusing the same notation from the description of $\text{NTRUDec}(\mathbf{c}, \tau := (\mathbf{F}, \mathbf{G}, p))$ gives

$$\begin{aligned} \mathbf{c}_1 &= \mathbf{F} \cdot \mathbf{s} + p \cdot \mathbf{G} \cdot (\mathbf{e} + ((q-1)/2) \cdot \mathbf{m}) \bmod q \\ &= \mathbf{F} \cdot \mathbf{s} + (p/2) \cdot \mathbf{G} \cdot (2\mathbf{e} - \mathbf{m}) \bmod q \\ &= \mathbf{F} \cdot \mathbf{s} + (p/2) \cdot \mathbf{G} \cdot (2\mathbf{e} - \mathbf{m}) \end{aligned}$$

over \mathbb{Z} because $\|\mathbf{F} \cdot \mathbf{s} + (p/2) \cdot \mathbf{G} \cdot (2\mathbf{e} - \mathbf{m})\|_\infty < q/2$. We then have $\mathbf{c}_2 = \mathbf{F} \cdot \mathbf{s} \bmod p/2 = \mathbf{F} \cdot \mathbf{s}$ over \mathbb{Z} because $\|\mathbf{F} \cdot \mathbf{s}\|_\infty < p/4$. Next, $\mathbf{c}_3 = \mathbf{e} + ((q-1)/2) \cdot \mathbf{m}$. Note that the conditions in the lemma statement imply that $2\beta'_e \cdot \sqrt{n} < (q-1)/2$. This gives the final output $\lfloor \mathbf{m} + \frac{2}{q-1} \cdot \mathbf{e} \rfloor = \mathbf{m}$ because $\|\frac{2\mathbf{e}}{q-1}\|_\infty < 1/2$. \square

Choosing Parameters. Looking ahead, we will instantiate this trapdoor for $q \approx 2^{32}$ and $n = 2^{11}$. So, we require $\beta \cdot \beta'_s < p/(4n)$ and, say, $\beta \cdot \beta'_e < q/(4n \cdot p)$. Picking $p = 2^{16}$, we get $\log(\beta) + \log(\beta'_s) < 16 - 2 - 11 = 3$ and $\log(\beta) + \log(\beta'_e) < 32 - 2 - 11 - 16 = 3$. Picking $\beta \approx 2^2$ and $\beta'_s = \beta'_e \approx 2$ we obtain an NTRU instance requiring BKZ block size 333 to solve (using the (overstretched) NTRU estimator [DvW21]) and an LWE instance requiring BKZ block size 594 to solve (using the lattice estimator [APS15]). According to the cost model from [MAT22] this costs about 2^{132} classical operations.¹¹

2.5 Homomorphic Encryption and TFHE

Fully homomorphic encryption (FHE) allows to perform computations on plaintexts by performing operations on ciphertexts. In slightly more detail, an FHE scheme consists of four algorithms: FHE.KeyGen , FHE.Enc , FHE.Eval , FHE.Dec . The key generation, encryption and decryption algorithms all work similarly to normal public key encryption. Together, they provide privacy (i.e. IND-CPA security) and decryption correctness. The interesting part of FHE is its homomorphic property. Assume that \mathcal{M} is the message space, e.g. $\mathcal{M} := \mathbb{Z}_P$. The homomorphic property is enabled by the FHE.Eval function which takes as input a public key pk , an arbitrary function $f : \mathcal{M}^k \rightarrow \mathcal{M}$, a sequence of ciphertexts $(c_i)_{i \in [k]}$ encrypting plaintexts $(m_i)_{i \in [k]}$, and outputs a ciphertext $c' \leftarrow \text{FHE.Eval}(\text{pk}, f, (c_0, \dots, c_{k-1}))$. The homomorphic property ensures that c' is an encryption of $f(m_0, \dots, m_{k-1})$. Intuitively, FHE allows to perform arbitrary computation on encrypted data without having to decrypt. Importantly, the privacy of the plaintext is maintained. In addition, an FHE scheme may also maintain the privacy of the evaluated computation (see below).

FHE was first realised by Gentry [Gen09]. A considerable amount of influential follow-up research provides the basis of most practically feasible schemes [FV12, BGV11, GSW13, CKKS17]. We will be focusing on an extension of the

¹¹ We note that quantum algorithms offer only marginal, i.e. less than square-root, speedups here [AGPS20].

third of these works known as TFHE [CGGI20] because its programmable bootstrapping technique lends itself well to our construction. For a good summary of TFHE, see the guide [Joy21].

Programmable bootstrapping. A crucial ingredient of any FHE scheme is a bootstrapping procedure. Essentially, homomorphic evaluation increases ciphertext noise, meaning that after a prescribed number of evaluations, a ciphertext becomes so noisy that it cannot be decrypted correctly. Bootstrapping provides a method of resetting the size of the noise in a ciphertext to allow for correct decryption using some bootstrapping key material. Note that the bootstrapping operation can either produce a ciphertext encrypted under the original key or a new one depending on the bootstrapping key material used.¹²

In TFHE we have access to look-up tables from \mathbb{Z}_d to \mathbb{Z}_P which we will denote by $\text{LUT}_{f(x)}()$ when realising $f(\cdot) : \mathbb{Z}_d \rightarrow \mathcal{M}$. For example we may write $\text{LUT}_{x \bmod 3}()$ for the table interpreting entries in \mathbb{Z}_d as integer $\in \{0, \dots, d-1\}$ and returning the result modulo 3 and treating $\{0, 1, 2\}$ as elements of the plaintext space. TFHE generalises bootstrapping by applying the look-up table to the plaintext at the same time as resetting the size of the noise. Here, d is the degree of a cyclotomic ring $\mathcal{R} = \mathbb{Z}[X]/(X^d + 1)$ and \mathbb{Z}_P is the plaintext space. There is a slight problem here in that the look-up table does not take plaintext-space inputs. However, this is overcome by using appropriate approximations between \mathbb{Z}_n and \mathbb{Z}_Q where n is a plain LWE dimension and Q the ciphertext modulus [CGGI20, Joy21]. We note that this functionality also implies the ability to map from a plaintext space \mathbb{Z}_{P_0} (which is interpreted as an element in \mathbb{Z}_d) to a plaintext space \mathbb{Z}_{P_1} with $P_0 \neq P_1$.¹³

Ax-hiding. In order to prevent an attacker from discovering details of the circuit, circuit privacy ensures that the resulting evaluated ciphertext does not leak anything about the circuit (beyond the evaluation of the circuit on a particular point). We relax this notion to allow the client to learn some structure of the circuit, but not the secret input of the server. In our setting we will not be interested in hiding the circuit but only part of the input (the secret key \mathbf{A}). In particular, we do not aim to hide the lookup tables which are public (essentially: mod p and mod q operations). We call this property **Ax-hiding**, and give the formal definition below. We note, however, that all techniques we are aware of for achieving **Ax-hiding** (see below), including heuristic approaches, seem to also achieve circuit privacy. We still believe formalising the security goal more precisely is a useful approach.

Definition 8. *Let LUT be an arbitrary lookup table. Let $(\mathbf{A}, \mathbf{x}_1), (\mathbf{B}, \mathbf{x}_2)$ have matching dimensions, matching base rings and satisfy $\text{LUT}(\mathbf{A} \cdot \mathbf{x}_1) = \text{LUT}(\mathbf{B} \cdot \mathbf{x}_2)$. We say an FHE scheme is **Ax-hiding** if for any such $(\mathbf{A}, \mathbf{x}_1), (\mathbf{B}, \mathbf{x}_2)$, any*

¹² This allows to restrict the number of sequential bootstrappings that can be performed, a fact we will rely on below.

¹³ This is accomplished by picking the “test polynomial” [Joy21, Sec. 5.3] appropriately.

$(FHE.pk, FHE.sk) \leftarrow \$ FHE.KeyGen(),$ any ct_1, ct_2 such that $FHE.Dec(sk, ct_1) = x_1$ and $FHE.Dec(sk, ct_2) = x_2,$ given $FHE.sk$ we have

$$FHE.Eval(FHE.pk, LUT(\mathbf{A} \cdot \star), ct_1) \approx_s FHE.Eval(FHE.pk, LUT(\mathbf{B} \cdot \star), ct_2).$$

A first approach to achieving circuit privacy and thus \mathbf{Ax} -hiding is to rely on noise drowning [Gen09]. An improvement to this technique was proposed by Ducas and Stehlé [DS16]. Instead of applying super-polynomial noise flooding in one step, Ducas and Stehlé proceed by iteratively applying the following κ times: perform bootstrapping and add some modest noise flooding (close to the size of the ciphertext noise). For FHEW [DM15] and TFHE, it is estimated that κ may be chosen between 8 and 16. Note that all of the extra bootstrapping is done once at the end of the homomorphic evaluation to “sanitize” the ciphertext and remove all remnants of the circuit that was applied.¹⁴

An alternative, also achieving full circuit privacy, is to use the work of Klucznik [Klu22] that applies a randomised bootstrapping procedure. Heuristically, plugging the bounds of [Klu22, Theorem 1] into [Klu22, Lemma 7] suggests the noise grows modestly by a rough factor of $\sqrt{\lambda} \cdot \sqrt{B_{br}}$ where B_{br} is an infinity norm bound on the bootstrapping key noise. We consider this the most promising approach. We give a third, heuristic approach in Appendix A.

Malicious security. Finally, note that the definition above is required to hold for “any” $FHE.pk, FHE.sk, ct_1$ and ct_2 generated honestly. More precisely, this means for *any* possible (i.e. valid) outputs of the appropriate FHE algorithms. For example, the word “any” for an error term distributed as a discrete Gaussian would mean any value satisfying some appropriate bound. Intuitively, malicious circuit privacy (or in our case \mathbf{Ax} -hiding) requires that an adversary cannot learn anything about the circuit (or (\mathbf{A}, \mathbf{x})) even in the presence of maliciously generated keys and ciphertexts. Thus if we can ensure the well-formedness of keys and ciphertexts, then a semi-honest circuit private FHE scheme is also secure against malicious adversaries. In the random oracle model, we can achieve this with NIZKs that show well-formedness of the keys and that the ciphertext is a valid ciphertext under that public key. Then, $FHE.Eval$ could explicitly check that the proof verifies and abort otherwise. As noted in [OPP14], the circuit privacy of FHE needs to hold even if the error distributions of the ciphertext are not generated according some canonical distribution, and are simply in the support of valid key/ciphertext pairs. For simplicity, we effectively ignore this fact and note that this caveat has been discussed in prior works, such as in [OPP14]. We note that the techniques from [BdMW16] achieve this enhanced version of circuit privacy.

¹⁴ This adds about 3 seconds of running time to $F_{\text{poprf}}.BlindEval$ using our implementation from Section 6. We note, however, that this approach is incompatible with our heuristic VOPRF construction.

Table 2. PRF Parameters

	$\lambda = 128$	Explanation		$\lambda = 128$	Explanation
p	2	modulus of \mathbf{x}, \mathbf{A}	n_p, m_p	256	dimensions of \mathbf{A}
q	3	modulus of $\mathbf{z}, \mathbf{G}_{\text{inp}}, \mathbf{G}_{\text{out}}$	n	128	dim. of $\mathbf{x} \pmod{p}$
n_q	192	rows of \mathbf{G}_{inp}	m	128	dim. of $\mathbf{z} \pmod{q}$

2.6 Crypto Dark Matter PRF

Let p, q be two primes where $p < q$. We now describe the ‘‘Crypto Dark Matter’’ PRF candidate [BIP⁺18,DGH⁺21]. It is built from the following *weak* PRF proposal $F_{\text{weak}} : \mathbb{Z}_p^{m_p \times n_p} \times \mathbb{Z}_p^{n_p} \rightarrow \mathbb{Z}_q$ where

$$F_{\text{weak}}(\mathbf{A}, \mathbf{x}) = \sum_{j=0}^{m_p-1} (\mathbf{A} \cdot \mathbf{x} \pmod{p})_j \pmod{q}.$$

Here \mathbf{A} is the secret key, \mathbf{x} is the input and $(\mathbf{A} \cdot \mathbf{x} \pmod{p})_j$ denotes the j -th component of $\mathbf{A} \cdot \mathbf{x} \pmod{p}$. In order to describe the strong PRF construction, we introduce a fixed public matrix $\mathbf{G}_{\text{inp}} \in \mathbb{Z}_q^{n_q \times n}$ and a p -adic decomposition operation $\text{decomp} : \mathbb{Z}_q^{n_q} \rightarrow \mathbb{Z}_p^{[\log_p(q)] \cdot n_q}$ where $[\log_p(q)] \cdot n_q = n_p$. The *strong* PRF candidate is $F_{\text{one}} : \mathbb{Z}_p^{m_p \times n_p} \times \mathbb{Z}_p^n \rightarrow \mathbb{Z}_q$ where

$$F_{\text{one}}(\mathbf{A}, \mathbf{x}) := F_{\text{weak}}(\mathbf{A}, \text{decomp}(\mathbf{G}_{\text{inp}} \cdot \mathbf{x} \pmod{q})).$$

In order to extend the small output of the above PRF constructions, the authors of [BIP⁺18] introduce another matrix $\mathbf{G}_{\text{out}} \in \mathbb{Z}_q^{m \times m_p}$ (with $m < m_p$) which is the generating matrix of some linear code. Then full PRF is $F_{\text{strong}} : \mathbb{Z}_p^{m_p \times n_p} \times \mathbb{Z}_p^n \rightarrow \mathbb{Z}_q^m$ where

$$F_{\text{strong}}(\mathbf{A}, \mathbf{x}) := \mathbf{G}_{\text{out}} \cdot (\mathbf{A} \cdot \text{decomp}(\mathbf{G}_{\text{inp}} \cdot \mathbf{x} \pmod{q}) \pmod{p}) \pmod{q}.$$

Given access to gates implementing mod p and mod q this PRF candidate can be implemented in a depth 3 arithmetic circuit. We give an example implementation in Appendix D and in the [attachment](#).¹⁵

Note that $\text{decomp}(\mathbf{G}_{\text{inp}} \cdot \mathbf{x} \pmod{q}) \in \mathbb{Z}_p^{n_p}$ does not depend on the PRF key. Thus, in an OPRF construction it could be precomputed and submitted by the client knowing \mathbf{x} . However, in this case, we must enforce that the client is doing this honestly via a zero-knowledge proof π that $\mathbf{y} := \text{decomp}(\mathbf{G}_{\text{inp}} \cdot \mathbf{x} \pmod{q})$ is well formed. Specifically, following [BIP⁺18], if $\mathbf{H}_{\text{inp}} \in \mathbb{Z}_q^{(n_q-n) \times n_q}$ is the parity check matrix of \mathbf{G}_{inp} and $\mathbf{G}_{\text{gadget}} := (p^{[\log_p(q)]-1}, \dots, 1) \otimes \mathbf{I}_{n_p}$ we may check

$$\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} \equiv \mathbf{0} \pmod{q}.$$

¹⁵ If the reader’s PDF viewer does not support PDF attachments (e.g. Preview on MacOS does not), then e.g. [pdfdetach](#) can be used to extract these files.

Note that as stated this does not enforce $\mathbf{x} \in \mathbb{Z}_p^n$ but $\mathbf{x} \in \mathbb{Z}_q^n$. Since it is unclear if this has a security implication, we may avoid this issue relying on a comment made in [BIP⁺18] that we may, wlog, replace \mathbf{G}_{inp} with a matrix in systematic (or row echelon) form. That is, writing $\mathbf{G}_{\text{inp}} = [\mathbf{I} \mid \mathbf{A}]^T \in \mathbb{Z}_q^{n_p \times n}$, $\mathbf{y} = (\mathbf{y}_0, \mathbf{y}_1) \in \mathbb{Z}_q^{n_p}$, the *protected encoded-input* PRF is defined as

$$F_{\text{pei}}(\mathbf{A}, \mathbf{y}) := \begin{cases} \mathbf{G}_{\text{out}} \cdot (\mathbf{A} \cdot \mathbf{y} \bmod p) \bmod q & \text{if } \mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} \equiv \mathbf{0} \bmod q \\ & \text{and } \mathbf{y}_0 \in \{0, 1\}^n \\ \perp & \text{otherwise.} \end{cases}$$

Note that with this definition of \mathbf{G}_{inp} the “most significant bits” of \mathbf{y}_0 will always be zero, so there is no point in extracting those when running `decomp`. Thus, we adapt `decomp` to simply return the first n output values in $\{0, 1\}$ and to perform the full decomposition on the remaining $(n_q - n)$ entries. We thus obtain $n_p := n + \lceil \log_p(q) \rceil \cdot (n_q - n)$. A similar strategy is discussed in Remark 7.13 of the full version of [BIP⁺18].

Security Analysis. The initial work [BIP⁺18] provided some initial cryptanalysis and relations to known hard problems to substantiate the security claims made therein. When \mathbf{A} is chosen to be a circulant rather than a random matrix, the scheme has been shown to have degraded security [CCKK21] contrary to the expectation stated in [BIP⁺18]. The same work [CCKK21] also proposes a fix. Further cryptanalysis was performed in [DGH⁺21], supporting the initial claims of concrete security. Our choices for $\lambda = 128$ (classically) are aggressive, especially for a post-quantum construction. This is, on the one hand, to encourage cryptanalysis. On the other hand, known cryptanalytic algorithms against the proposals in [BIP⁺18, DGH⁺21] require exponential memory in addition to exponential time, a setting where Grover-like square-root speed-ups are less plausible, cf. [AGPS20] (which, however, treats the Euclidean distance rather than Hamming distance).

3 OPRF Candidate

We wish to design an (P)OPRF where the server homomorphically evaluates the PRF using its secret key and uses some form of circuit private FHE to protect its key. We use the same notation as in Section 2.6.

3.1 Extending the PEI PRF

Here, we first observe that the PRFs defined in Section 2.6 trivially fail to achieve pseudorandomness as they map $\mathbf{0} \in \mathbb{Z}_p^n \rightarrow \mathbf{0} \in \mathbb{Z}_q^m$ which holds with $\text{negl}(\lambda)$ probability for a random function. We thus define

$$F_{\text{strong}}(\mathbf{A}', \mathbf{x}) := \mathbf{G}_{\text{out}} \cdot (\mathbf{A} \cdot (\text{decomp}(\mathbf{G}_{\text{inp}} \cdot \mathbf{x} \bmod q), 1) \bmod p) \bmod q.$$

and

$$F_{\text{pei}}(\mathbf{A}', \mathbf{y}) := \begin{cases} \mathbf{G}_{\text{out}} \cdot (\mathbf{A}' \cdot (\mathbf{y}, 1) \bmod p) \bmod q & \text{if } \mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} \equiv \mathbf{0} \bmod q \\ \perp & \text{otherwise.} \end{cases}$$

for $\mathbf{A}' \in \mathbb{Z}_p^{m_p \times (n_p+1)}$, i.e. extended by one column. Furthermore, we wish to support an additional input $\mathbf{t} \in \mathbb{Z}_p^n$ to be submitted in the clear. For this, we deploy the standard technique of using a key derivation function to derive a fresh key per tag \mathbf{t} [CHL22,JKR18]. In particular, let $\text{RO}_{\text{key}} : \mathbb{Z}_p^{m_p \times n_p} \times \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p^{m_p \times (n_p+1)}$ be a random oracle, we then define our PRF candidate $F_{\mathbf{A}}^{\text{RO}_{\text{key}}}(\mathbf{t}, \mathbf{x})$ in Algorithm 1. Clearly, if $F_{\text{pei}}(\mathbf{A}', \mathbf{y})$ is a PRF then $F_{\mathbf{A}}^{\text{RO}_{\text{key}}}(\cdot, \cdot)$ is a PRF with input (\mathbf{t}, \mathbf{x}) as $\mathbf{A}_{\mathbf{t}}$ in Algorithm 1 is simply a fresh $F_{\text{pei}}(\cdot)$ key for each distinct value of \mathbf{t} .

Algorithm 1 $F_{\mathbf{A}}^{\text{RO}_{\text{key}}}(\mathbf{t}, \mathbf{x})$

Input: $\mathbf{A} \in \mathbb{Z}_p^{m_p \times n_p}$, $\mathbf{x} \in \mathbb{Z}_p^n$, $\mathbf{t} \in \mathbb{Z}_p^n$

Output: $F_{\mathbf{A}}(\mathbf{t}, \mathbf{x})$

$\mathbf{A}_{\mathbf{t}} \leftarrow \text{RO}_{\text{key}}(\mathbf{A}, \mathbf{t})$

$\mathbf{y} \leftarrow \text{decomp}(\mathbf{G}_{\text{inp}} \cdot \mathbf{x} \bmod q)$

$\mathbf{z} \leftarrow \mathbf{G}_{\text{out}} \cdot (\mathbf{A}_{\mathbf{t}} \cdot (\mathbf{y}, 1) \bmod p) \bmod q$

return \mathbf{z}

3.2 TFHE-based Instantiation

Our main TFHE-based instantiation is given in Figure 3. An important alteration to the FHE key generation is that $\text{FHE.KeyGen}^{(\text{pp})}$ will output a commitment to secret key $\text{FHE.sk} = \mathbf{s} \in \mathbb{Z}_2^e$ in addition to the standard public key FHE.pk . In particular, $\text{FHE.KeyGen}^{(\text{pp})}(\cdot)$ begins by running $(\text{FHE.pk}, \text{FHE.sk}) \leftarrow \text{FHE.KeyGen}(\cdot)$ and then adds the commitment \mathbf{b}_{pk} to FHE.pk . This commitment takes the form $\mathbf{b}_{\text{pk}} = \mathbf{A}_{\text{pp}} \cdot \mathbf{r} + \mathbf{e} + \lfloor Q/2 \rfloor \cdot (\mathbf{s}, \mathbf{0}) \in \mathbb{Z}_Q^N$ where Q is the ciphertext modulus, $\mathbf{r}, \mathbf{e} \leftarrow \text{G}(\chi')^N$ where χ' is a discrete Gaussian of standard deviation $\beta' \approx 4$ and $N = 2048$ as in Section 2.4. One can view \mathbf{b}_{pk} as a partial symmetric LWE encryption of the secret key from the (T)LWE encryption scheme within TFHE, so χ' is simply an error distribution. Therefore, using the same LWE assumption from Section 2.4, \mathbf{b}_{pk} is indistinguishable from random and it is easy to check that its presence does not affect the IND-CPA property of TFHE. Furthermore, since \mathbf{b}_{pk} is simply a randomised function of $(\text{FHE.pk}, \text{FHE.sk})$, it can be constructed by an adversarial client. Thus, its advantage against the \mathbf{Ax} -hiding property (or semi-honest circuit privacy) of FHE in which the key also contains \mathbf{b}_{pk} remains unchanged. To summarise, the public key material output by $\text{FHE.KeyGen}^{(\text{pp})}$ is $\text{FHE.pk}^{(\text{pp})} := (\mathbf{b}_{\text{pk}}, \text{FHE.pk})$.

Note that although the server does not need to create encryptions itself, we still use the public key encryption version of TFHE rather than a symmetric key

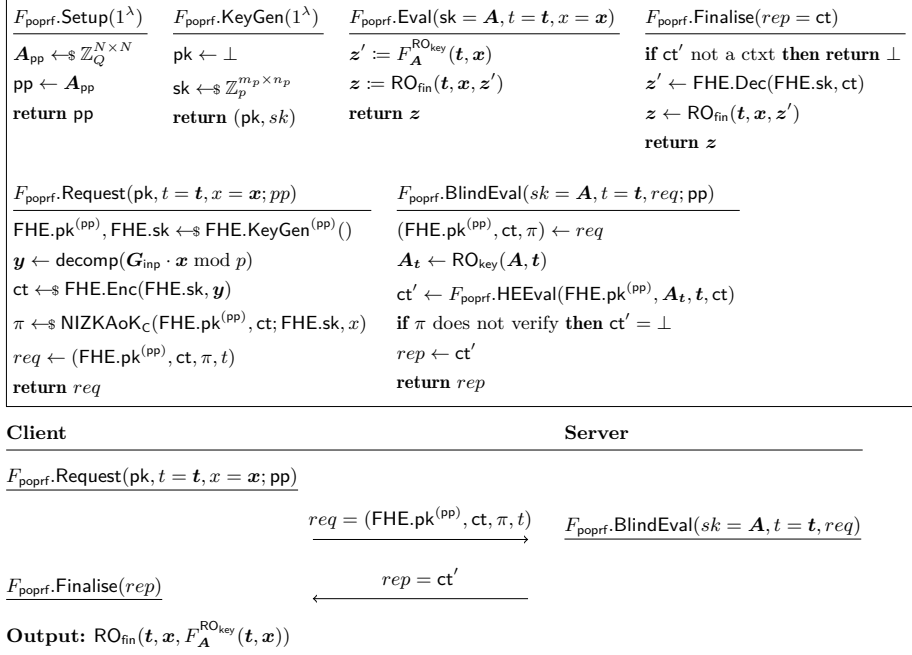


Fig. 3. Main construction.

version. The extra key material in the public key version does not affect efficiency in any noticeable way as the bootstrapping key sizes are the main bottleneck. This can be seen in Appendix B where details of NIZKAoK_C are elaborated and where we fully describe the generation of the bootstrapping keys.

Remark 4. The ciphertext component encrypting 1 from the input of Algorithm 2 may be produced by the server (without knowledge of \mathbf{y}).

Choosing Parameters & Size Estimates. We require that binary-secret LWE is hard for (e, Q, σ) where e is the dimension of the LWE secret, Q the modulus and σ the standard deviation of the noise. We also require M -LWE to be hard for $(\ell, d, Q, \bar{\sigma})$ where d is the ring dimension, ℓ is the module dimension and $\bar{\sigma}$ is the standard deviation of the noise. To pick parameters, we simply follow the `tfhe-rs` TFHE library [CGGI20] and its `(PARAM_MESSAGE_2_CARRY_0)` parameter set, which permits us to set $Q = 2^{32}$ and then fixes $e = 656, \ell = 2, d = 512, \sigma = 2^{17.1}, \bar{\sigma} = 2^{7.5}$. We note that this is likely not optimal as these parameters are meant to support a wide range of computations. We estimate the size of the bootstrapping key (which may be considered an amortisable offline communication cost) in Appendix B as 16.8MB. We also estimate the size of the zero-knowledge proofs accompanying this key as 41.42KB.

Algorithm 2 $F_{\text{poprf}}^{\text{HEEval}}$

Input: pk HE public key \triangleright with ciphertext modulus Q , plaintext modulus P

Input: $\mathbf{A} \in \mathbb{Z}_P^{n_p \times n_p}$, $\mathbf{t} \in \mathbb{Z}_P^n$, $\text{ct} \in \mathcal{C}^{n_p}$ encrypting $\mathbf{y} \in \mathbb{Z}_P^{n_p}$

Output: $\text{ct}' \in \mathcal{C}^m$ encrypting $F_{\mathbf{A}}^{\text{ROkey}}(\mathbf{t}, \mathbf{x})$ or \perp

- 1: $\mathbf{A}_{\mathbf{t}} \leftarrow \text{RO}_{\text{key}}(\mathbf{A}, \mathbf{t})$ in the clear.
 - 2: $\text{ct}' \leftarrow \text{FHE.Enc}(\text{pk}, 1)$ \triangleright can be e.g. $1 \cdot \lfloor Q/P \rfloor$
 - 3: $\text{ct}^{(1)} \leftarrow \text{FHE.Eval}(\text{pk}, \mathbf{A}_{\mathbf{t}} \cdot (\text{ct}, \text{ct}'))$
 - 4: $\text{ct}_i^{(2)} \leftarrow \text{FHE.Eval}(\text{pk}, \text{LUT}_{x \bmod p}(\text{ct}_i^{(1)})) \quad \forall i \in \mathbb{Z}_{m_p}$ $\triangleright \bmod p$, depth ≥ 1
 - 5: $\text{ct}^{(3)} \leftarrow \text{FHE.Eval}(\text{pk}, \mathbf{G}_{\text{out}} \cdot \text{ct}^{(2)})$
 - 6: $\text{ct}'_i \leftarrow \text{FHE.Eval}(\text{pk}, \text{LUT}_{x \bmod q}(\text{ct}_i^{(3)})) \quad \forall i \in \mathbb{Z}_m$ $\triangleright \bmod q$, depth ≥ 1
 - 7: **return** ct'
-

Each request then requires to send LWE encryptions of $n_p = 256$ bits \mathbf{m} using protected encoded inputs, i.e. $(\mathbf{A}', \mathbf{b} := \mathbf{A}' \cdot \mathbf{s} + \mathbf{e} + \lfloor Q/P \rfloor \cdot \mathbf{m})$ where $\mathbf{A}' \in \mathbb{Z}_Q^{n_p \times e}$ and $\mathbf{b} \in \mathbb{Z}_Q^{n_p}$. Here, \mathbf{A}' can be computed from a small seed of 256 bits. For \mathbf{b} we need to transmit $n_p \cdot \log Q$ bits. However, as noted in e.g. [LDK⁺22], given the large size of the noise ($\sigma = 2^{17.1}$) we may drop the least significant, say, 16 bits, reducing the cost of \mathbf{b} to $16 \cdot 256$. In total we have a ciphertext size of 0.5kB. The accompanying zero-knowledge proofs take up 26.7kB but can be amortised to cost about 1.3kB per query when sending 64 queries in one shot (Appendix B).

The message back from the server is $m = 128$ encryptions of the output elements $\in \mathbb{Z}_q$. We can also drop the least significant bits of \mathbf{A}' , since \mathbf{s} is binary. In particular, we may drop, say, the eight least significant bits and we arrive at $e \cdot m \cdot 24$ for \mathbf{A}' . We need $m \log Q$ bits for \mathbf{b} . We can use the same trick of dropping lower order bits for \mathbf{b} again, so we obtain $128 \cdot 16$. In total we get 246.2kB. We note that it is more efficient to pack all return values into a single RLWE sample using techniques from [CDKS21], since the cost of transmitting \mathbf{A}' dominates here. This increases the public key size (since we would include another key switching key) but reduce the size of response to about 3.2kB. For more details on these values, see Appendix B.1. We give our code for estimating parameters in Appendices E and F.

Remark 5. Looking ahead, we will see that our implementation in Section 6 does not achieve these sizes. The reason for this is that (a) \mathbf{A}' is not compressed as above and (b) the API we are using defaults to using M -LWE ciphertexts with dimension $\ell \cdot d = 1024$ per bit.

4 Security Proofs

We first prove the pseudorandomness property against malicious clients in Theorem 1 and then privacy (POPRIV1 only) against servers in Theorem 2.

Theorem 1. *Let FHE denote the TFHE scheme. The construction F_{poprf} from Figure 3 satisfies the POPRF property from Definition 3, with random oracles RO_{fin} and RO_{key} , if:*

- The client zero-knowledge proof is sound
- FHE is \mathbf{Ax} -hiding
- The $\text{mat-NTRU}_{N,P',Q,B}$ assumption holds where Q is the FHE modulus.
- P' is even and coprime to Q such that $Q > BN(\beta' + P' \cdot (2\beta' + 1)/2)$ where β' is the standard deviation used in $\text{FHE.KeyGen}^{(\text{pp})}$
- $F(\cdot, \cdot)$, defined in Section 3, is a PRF with output range super-polynomial in the security parameter¹⁶.

Proof. The following simulator S will be used to prove the result:

$\mathsf{S}.\text{Init}$: Sample $\mathbf{A}^{(\mathsf{S})} \leftarrow_{\mathcal{S}} \mathbb{Z}_p^{m_p \times n_p}$ and $\mathbf{A}_{\text{pp}}^{(\mathsf{S})} \leftarrow_{\mathcal{S}} \mathbb{Z}_Q^{N \times N}$. Set $st_{\mathsf{S}} = \mathbf{A}^{(\mathsf{S})}$, $\text{pp}_0 := \mathbf{A}_{\text{pp}}^{(\mathsf{S})}$ and $\text{pk}_0 = \perp$.

$\mathsf{S}.\text{BlindEval}(t, \text{req}, st_{\mathsf{S}})$: Return $\mathcal{F}.\text{BlindEval}_{\text{pp}_0}^{\text{RO}_{\text{key}}}(\mathbf{A}^{(\mathsf{S})}, t, \text{req})$. Note that this algorithm does not need to make any calls to LimitEval .

$\mathsf{S}.\text{Eval}^{\text{LimitEval}}(\mathbf{x}_{\text{in}}, st_{\mathsf{S}})$: If query $\mathbf{x}_{\text{in}} := (t, \mathbf{x}, \mathbf{z})$ appears in st_{S} as a previous query, return the same answer. If $\mathbf{z} \neq F_{\mathbf{A}^{(\mathsf{S})}}(t, \mathbf{x})$ return a uniformly random h and store $(\mathbf{x}_{\text{in}}, h)$ in st_{S} . If $\mathbf{z} = F_{\mathbf{A}^{(\mathsf{S})}}(t, \mathbf{x})$, query $\text{LimitEval}(\mathbf{x}_{\text{in}})$ and return its answer \mathbf{h}' , storing $(\mathbf{x}_{\text{in}}, \mathbf{h}')$ in st_{S} .

Define G_0 to be the $\text{POPRF}_{\mathcal{F}, \text{Sim}, H}^{\mathbf{A}, b=0}$ game and G_1 to be the $\text{POPRF}_{\mathcal{F}, \text{Sim}, H}^{\mathbf{A}, b=1}$ game. Note that \mathcal{A} has oracle access to Eval , BlindEval and Prim . These three oracles behave (jointly) identically in G_0 and G_1 as long as S does not get \perp in a LimitEval query when answering $\mathsf{S}.\text{Eval}$. Therefore, the distinguishing advantage between G_0 and G_1 is at most the probability that after making q BlindEval queries, \mathcal{A} has managed to submit $q + 1$ distinct tuples of the form $(t_j, \mathbf{x}_j, F_{\mathbf{A}^{(\mathsf{S})}}(t_j, \mathbf{x}_j))$ to Prim in G_0 . Denote this event E . To complete the proof, we bound $\Pr[\mathsf{E}]$ using Lemma 2. \square

Lemma 2. *Assume that all of the conditions in Theorem 1 hold. Then $\Pr[\mathsf{E}]$ is negligible.*

Proof (Of Lemma 2). Note that every adversarial input to the BlindEv oracle is either well-formed, or answered with $\text{rep} = \perp$ due to the soundness of the client zero-knowledge proof. Therefore, we know that with overwhelming probability, the ciphertext and $(\text{FHE.pk}^{(\text{pp})}, \text{FHE.sk})$ that the malicious client uses is of the correct form in the event that the response is not \perp . We now describe hybrid games \mathcal{H}_i and events E_i for $i \in \{1, 2, 3\}$:

\mathcal{H}_1 : Here S is changed to S_1 , where S_1 is identical to S apart from that pp_0 is replaced by $\text{pp}'_0 = \mathbf{A}'$ where $(\mathbf{A}', \tau) \leftarrow_{\mathcal{S}} \text{NTRUTrapGen}(N, Q, P', B)$ and τ is added to the initial st_{S} (but is unused throughout). Here, $P > 4N$ is a power of 2 such that $Q > BN(\beta' + P' \cdot (2\beta' + 1)/2)$ with β' the standard deviation of \mathbf{r} and \mathbf{e} used to produce the commitment part \mathbf{b}_{pk} of $\text{FHE.pk}^{(\text{pp})}$ (see Section 3.2). E_1 is defined as the event that after q BlindEval queries, \mathcal{A} has managed to submit $q + 1$ distinct pairs of the form $(t_j, \mathbf{x}_j, F_{\mathbf{A}^{(\mathsf{S})}}(t_j, \mathbf{x}_j))$ to Prim .

¹⁶ e.g. 2^λ works.

\mathcal{H}_2 : S_1 is modified to S_2 , where S_2 is identical to S_1 apart from that $S_2.\text{BlindEval}$ is modified in the following way. On input

$$req = (\text{FHE.pk}^{(\text{pp})}) := (\mathbf{b}_{\text{pk}}, \text{FHE.pk}), \text{ct}, \pi, t),$$

S_2 checks π . If the proof verifies then it sets $\mathbf{s} := \text{NTRUDec}(\mathbf{b}_{\text{pk}}, \tau)$, $\mathbf{y} := \text{FHE.Dec}(\mathbf{s}, \text{ct})$, computes \mathbf{x} from \mathbf{y} using \mathbf{G}'_{inp} in systematic form, and computes $\mathbf{z} = F_{\mathbf{A}^{(\text{s})}}(\mathbf{t}, \mathbf{x})$ as in the original construction. It then samples a matrix $\mathbf{B} \in \mathbb{Z}_2^{m_p \times n_p}$ and further computes $\mathbf{w} = \mathbf{B} \cdot \mathbf{y} \bmod p$. Then, it is easy to compute some $\mathbf{G}'_{\text{out}} \in \mathbb{Z}_q^{m \times m_p}$ such that $\mathbf{z} = \mathbf{G}'_{\text{out}} \cdot \mathbf{w} \bmod q$. Finally, define F' to be the function F but using \mathbf{G}'_{out} instead of \mathbf{G}_{out} . It returns $\text{FHE.Eval}(\text{FHE.pk}, F'_{\mathbf{B}^{(\text{s})}}(\mathbf{t}, \cdot), \text{ct})$ in response to req . E_2 is defined as the event that after q BlindEval queries, \mathcal{A} has managed to submit $q + 1$ distinct pairs of the form $(\mathbf{t}_j, \mathbf{x}_j, F_{\mathbf{A}^{(\text{s})}}(\mathbf{t}_j, \mathbf{x}_j))$ to Prim.

\mathcal{H}_3 : S_3 is S_2 apart from $S_3.\text{BlindEval}$ and $S_3.\text{Eval}$. In these two functions, S_3 replaces all computations of $F_{\mathbf{A}^{(\text{s})}}(\cdot, \cdot)$ by a truly random function \mathbf{G} . Additionally, S_3 does not sample $\mathbf{A}^{(\text{s})}$ during $S_3.\text{Init}$. E_3 is defined as the event that after q BlindEval queries, \mathcal{A} has managed to submit $q + 1$ distinct pairs of the form $(\mathbf{t}_j, \mathbf{x}_j, \mathbf{G}(\mathbf{t}_j, \mathbf{x}_j))$ to Prim.

We first show that $|\Pr[\mathbf{E}] - \Pr[\mathbf{E}_1]| = \text{negl}(\lambda)$ by the $\text{mat-NTRU}_{N, P', Q, B}$ assumption. In order to do so, we build a matrix NTRU distinguisher $\mathcal{B}_{\text{mat-NTRU}}$ that implements \mathbf{S} for the POPRF adversary \mathcal{A} using its $\text{mat-NTRU}_{N, P', Q, B}$ challenge as \mathbf{A}_{pp} . Clearly, if the challenge is a uniform matrix, $\mathcal{B}_{\text{mat-NTRU}}$ perfectly recreates \mathbf{S} whereas if the challenge is non-uniform, it perfectly simulates S_1 from \mathcal{A} 's perspective. Therefore, if $|\Pr[\mathbf{E}] - \Pr[\mathbf{E}_i]|$ was not negligible, then $\mathcal{B}_{\text{mat-NTRU}}$ would be able to distinguish and break the $\text{mat-NTRU}_{N, P, Q}$ assumption by testing whether \mathcal{A} manages to submit $q + 1$ distinct tuples of the form $(\mathbf{t}_j, \mathbf{x}_j, F_{\mathbf{A}^{(\text{s})}}(\mathbf{t}_j, \mathbf{x}_j))$ to Prim given just q BlindEval queries.

Next, we show that $|\Pr[\mathbf{E}_1] - \Pr[\mathbf{E}_2]| = \text{negl}(\lambda)$ by the \mathbf{Ax} -hiding property of TFHE. To do so, we consider a sequence of hybrid events $\mathbf{E}_{1, i}$ where simulator $S_{1, i}$ answers the first i calls to BlindEval as in S_1 and the remainder as in S_2 . Clearly, $\mathbf{E}_1 = \mathbf{E}_{1, 0}$ and $\mathbf{E}_{1, q_t} = \mathbf{E}_2$ if q_t is a polynomial upper bound on the number of BlindEval queries. Suppose there is an index $i^* \in [q_t]$ such that $|\Pr[\mathbf{E}_{1, i^*}] - \Pr[\mathbf{E}_{1, i^*+1}]|$ is non-negligible. Note that if the $(i^* + 1)$ -th request did not have a correctly verifying proof, then $\Pr[\mathbf{E}_{1, i^*}] - \Pr[\mathbf{E}_{1, i^*+1}] = 0$. Therefore, we assume that the proof in the $(i^* + 1)$ -th query verifies and all ciphertexts/keys in the request are well-formed. This tells us that the NTRUDec correctly recovers the FHE secret key which in turn implies that the \mathbf{y} and \mathbf{x} recovered are the correct ones. Therefore, if $|\Pr[\mathbf{E}_{1, i^*}] - \Pr[\mathbf{E}_{1, i^*+1}]|$ is non-negligible, there would exist a $\text{FHE.pk}, \text{FHE.sk}, \text{ct} = \text{FHE.Enc}(\text{FHE.pk}, \mathbf{x})$ (in particular, the key-pair and ciphertext associated to \mathcal{A} 's $(i^* + 1)$ -th BlindEval query) that allows an efficient distinguisher \mathcal{B} between

$$\text{FHE.Eval}(\text{FHE.pk}, F_{\mathbf{A}^{(\text{s})}}(\mathbf{t}, \cdot), \text{ct}) \tag{1}$$

and

$$\text{FHE.Eval}(\text{FHE.pk}, F'_{\mathbf{B}^{(\text{s})}}(\mathbf{t}, \cdot), \text{ct}) \tag{2}$$

for the value of \mathbf{t} from the (i^*+1) -th query. Note that by construction $F'_{\mathcal{B}(s)}(\mathbf{t}, \mathbf{x}) = F_{\mathcal{A}(s)}(\mathbf{t}, \mathbf{x})$ yet $F'_{\mathcal{B}(s)}$ contains no secret information.

Essentially, \mathcal{B} implements S_{1,i^*} for \mathcal{A} apart from the $(i^* + 1)$ -th BlindEval query. Suppose \mathcal{B} is given $\text{FHE.pk}, \mathbf{t}$ and ct in the $(i^* + 1)$ -th query. Suppose also that it is given a challenge ciphertext ct' that is either of the form (1) or (2), and that it uses ct' to respond to the $(i^* + 1)$ -th BlindEval query. If ct' is of the first form, then \mathcal{B} implements S_{1,i^*} and if it is of the latter form, \mathcal{B} implements S_{1,i^*+1} . Note that, the forms (1) and (2) can be rewritten as the $\text{FHE.Eval}(\text{FHE.pk}, \text{LUT}(\mathbf{G}_{\text{out}} \cdot \star), \tilde{\text{ct}}_1)$ and $\text{FHE.Eval}(\text{FHE.pk}, \text{LUT}(\mathbf{G}'_{\text{out}} \cdot \star), \tilde{\text{ct}}_2)$ respectively, where $\tilde{\text{ct}}_i$ are the intermediate ciphertexts that encrypt the vectors $\tilde{\mathbf{w}}_i$ of intermediate evaluations in the two cases. Therefore, distinguishing between ct' of form (1) and (2) is precisely breaking the \mathbf{Ax} -hiding property. Assume \mathcal{B} decides its output by checking whether \mathcal{A} manages to submit $q + 1$ distinct pairs of the form $(\mathbf{t}_j, \mathbf{x}_j, F_{\mathcal{A}(s)}(\mathbf{t}_j, \mathbf{x}_j))$ to Prim given just q BlindEval queries. If $|\Pr[\mathbf{E}_{1,i^*}] - \Pr[\mathbf{E}_{1,i^*+1}]|$ is non-negligible, then \mathcal{B} can distinguish the two cases for ct' and thus break \mathbf{Ax} -hiding.

Next, we show $|\Pr[\mathbf{E}_2] - \Pr[\mathbf{E}_3]| = \text{negl}(\lambda)$. Suppose that this was not the case. Then we construct an algorithm \mathcal{B} that distinguishes $F_{\mathcal{A}(s)}(\cdot, \cdot)$ from uniform random. \mathcal{B} interacts with a PRF challenger, and uses it to implement S_2 by querying the PRF challenger on input (\mathbf{t}, \mathbf{x}) instead of computing $F_{\mathcal{A}(s)}(\mathbf{t}, \mathbf{x})$. In doing this, if the PRF challenger is returning uniform values, \mathcal{B} simulates S_3 for \mathcal{A} . Otherwise, \mathcal{B} perfectly simulates S_2 for \mathcal{A} . Therefore, if $|\Pr[\mathbf{E}_2] - \Pr[\mathbf{E}_3]|$ was not negligible, then \mathcal{B} could distinguish F from a PRF by checking whether \mathcal{A} manages to submit $q + 1$ distinct pairs of the form $(\mathbf{t}_j, \mathbf{x}_j, \mathbf{z}_j)$ to Prim such that \mathbf{z}_j agrees with the PRF oracle on input $(\mathbf{t}_j, \mathbf{x}_j)$ given just q BlindEval queries. To complete the proof, we use the next lemma. \square

Lemma 3. $\Pr[\mathbf{E}_3] = \text{negl}(\lambda)$ assuming the PRF output space is super-polynomial in the security parameter.

Proof. Recall that when considering event \mathbf{E}_3 , the adversary \mathcal{A} has oracle access to Eval, BlindEval and Prim. Further, these oracles are implemented by S_3 using a truly random function \mathbf{G} instead of $F_{\mathcal{A}(s)}$. \mathcal{A} has access to \mathbf{G} outputs through BlindEval queries. In particular, one \mathbf{G} output is computed for every distinct BlindEval query. The oracles Eval and Prim can be used by \mathcal{A} to see whether a query is not of the form $(\mathbf{t}_j, \mathbf{x}_j, \mathbf{G}(\mathbf{t}_j, \mathbf{x}_j))$ through consistency. This is the full extent to which \mathcal{A} has access to \mathbf{G} . Therefore, from the perspective of \mathcal{A} , \mathbf{G} is a random oracle with an additional oracle that tells if a guessed output is *incorrect*. Given this setup, an adversary has a negligible probability of producing $q + 1$ outputs of \mathbf{G} if it only knows q evaluations of \mathbf{G} if the output space of \mathbf{G} is super-polynomial. Therefore, $\Pr[\mathbf{E}_3] = \text{negl}(\lambda)$. \square

Theorem 2. Let FHE denote the TFHE scheme. The construction \mathcal{F} from Figure 3 satisfies the POPRIV1 property if the following hold:

- FHE is IND-CPA.
- The client proof is zero-knowledge.

- The $\text{LWE}_{Q,N,N,\chi',\chi'}$ assumption holds where $\chi' = D_{\mathbb{Z},\beta'}$ is the error distribution used in $\text{FHE.KeyGen}^{(pp)}$.

Proof. Let G_0 and G_1 denote the $\text{POPRIV1}_{\mathcal{F},H}^{A,b}$ game for $b = 0$ and $b = 1$ respectively. Furthermore, let \bar{G}_0 and \bar{G}_1 be the G_0 and G_1 modified so that all Trans oracle queries have their zero-knowledge proofs in the requests replaced by *simulated* zero-knowledge proofs. Clearly, $G_0 \approx_c \bar{G}_0$ and $G_1 \approx_c \bar{G}_1$ by the zero-knowledge property of the client proofs. Next, we let G'_0 and G'_1 be the same as \bar{G}_0 and \bar{G}_1 apart from the way all public keys of the form $\text{FHE.pk}^{(pp)} := (\mathbf{b}_{\text{pk}}, \text{FHE.pk})$ are sampled. In particular, FHE.pk will remain the same, however, \mathbf{b}_{pk} will be replaced by uniform random values \mathbf{u}_{pk} . Note that $\mathbf{b}_{\text{pk}} = \mathbf{A}_{pp} \cdot \mathbf{r} + \mathbf{e} + \lfloor Q/2 \rfloor \cdot \mathbf{s}$ is an LWE encryption. As we argue next, $\bar{G}_0 \approx_c G'_0$ and $\bar{G}_1 \approx_c G'_1$ assuming the $\text{LWE}_{Q,N,N,\chi',\chi'}$ assumption holds. Suppose we want to prove $\bar{G}_0 \approx_c G'_0$. This can be formally argued by building a distinguisher \mathcal{B} between an M multi-secret LWE challenge of the form $(\mathbf{A}, \mathbf{B}) \in \mathbb{Z}_Q^{N \times N} \times \mathbb{Z}_Q^{N \times M}$ where $\mathbf{B} \leftarrow_{\$} \mathbb{Z}_Q^{N \times M}$ or $\mathbf{B} = \mathbf{A} \cdot \mathbf{R} + \mathbf{E}$ where $\mathbf{R} \leftarrow_{\$} (\chi')^{N \times M}$, $\mathbf{E} \leftarrow_{\$} (\chi')^{N \times M}$. Denoting \mathbf{b}_i as the i -th column of \mathbf{B} , \mathbf{b}_i is either uniform or $\mathbf{b}_i = \mathbf{A} \cdot \mathbf{r}_i + \mathbf{e}_i$. Therefore, the distinguisher \mathcal{B} can simulate \bar{G}_0 or G'_0 for an adversary \mathcal{A} by setting $\mathbf{A}_{pp} = \mathbf{A}$ and running all algorithms as specified in \bar{G}_0 apart from $\text{FHE.KeyGen}^{(pp)}$. To run the i -th instance of $\text{FHE.KeyGen}^{(pp)}$, \mathcal{B} samples $\text{FHE.pk}, \text{FHE.sk} \leftarrow_{\$} \text{FHE.KeyGen}$ and then sets $\text{FHE.pk}^{(pp)} := (\mathbf{b}_i + \lfloor Q/2 \rfloor \text{FHE.sk})$. Clearly, if \mathbf{b}_i is not uniform, it perfectly simulates \bar{G}_0 for \mathcal{A} . Otherwise, it perfectly simulates G'_0 as shifting a uniform value does not affect uniformity. Therefore, \mathcal{B} would have the same advantage in distinguishing its multi-secret LWE problem as \mathcal{A} has in distinguishing \bar{G}_0 and G'_0 . By a standard hybrid, the multi-secret LWE assumption holds (with a polynomial number of secrets) if the plain single secret LWE assumption holds. Therefore, $\bar{G}_0 \approx_c G'_0$ and similarly, $\bar{G}_1 \approx_c G'_1$.

Now we show that $G'_0 \approx_c G'_1$ using the IND-CPA property of FHE. We will use a sequence of hybrids $G'_{0,i}$ where $i \in [q_R]$ for polynomial q_R that bounds the number of Run queries. In $G'_{0,i}$, all Run queries after the i -th one are answered as in G'_0 . All prior Run queries are answered according to the following description:

- Set $req' = (\text{FHE.pk}'^{(pp)} := (\mathbf{b}'_{\text{pk}}, \text{FHE.pk}'), \text{ct}', t, \pi')$ where $\text{FHE.pk}'$ is a freshly sampled public key with a uniform \mathbf{b}'_{pk} , $\text{ct}' = \text{FHE.Enc}(\text{FHE.pk}', x_1)$ and π' is a simulated zero-knowledge proof.
- Compute $rep' \leftarrow_{\$} \mathcal{F}.\text{BlindEval}(\text{sk}, req')$ and set $y_0 = \mathcal{F}.\text{Eval}(\text{sk}, t, x_0)$
- Set $\tau'_0 = (req', rep', y_0)$ and τ_1 as in G'_0 .
- Return (τ'_0, τ_1)

We now show that for every i , $G'_{0,i} \approx_c G'_{0,i+1}$. In order to show this, we build an IND-CPA adversary \mathcal{B} that distinguishes the IND-CPA game for FHE with the same advantage that \mathcal{A} has in distinguishing $G'_{0,i}$ from $G'_{0,i+1}$. \mathcal{B} acts as the POPRIV1 challenger for \mathcal{A} , sampling $\text{pp} \leftarrow_{\$} \mathcal{F}.\text{Setup}$, $(\text{pk}, \text{sk}) \leftarrow_{\$} \mathcal{F}.\text{KeyGen}$ and initialising the random oracles RO and RO_{key} honestly. This allows \mathcal{B} to answer the $(i+1)$ -th Run query onwards as in G_0 . For the first $(i-1)$ queries, the Run queries are answered by \mathcal{B} as in $G'_{0,i}$ i.e. as described above. However, for the i -th

Run query \mathcal{A} makes, denoted $(t^{(i)}, x_0^{(i)}, x_1^{(i)})$, \mathcal{B} asks its IND-CPA challenger for a public key FHE.pk^* , queries the IND-CPA challenge oracle to encrypt either $x_0^{(i)}$ or $x_1^{(i)}$ receiving ct^* in response. \mathcal{B} then sets $\text{req}^* = (\text{FHE.pk}^*, \text{ct}^*, t^{(i)}, \pi^*)$ where π^* is a simulated proof. Finally, \mathcal{B} responds to the i -th Run query by setting $\tau'_0 = ((\text{req}^*, \mathcal{F}.\text{BlindEval}(\text{sk}, \text{req}^*)), \mathcal{F}.\text{Eval}(\text{sk}, t, x_0))$ and returning (τ'_0, τ_1) to \mathcal{A} where τ_1 is computed in the same way as G'_0 . If \mathcal{B} received an encryption of $x_0^{(i)}$, it perfectly simulates the game $G'_{0,i}$ for \mathcal{A} . Otherwise it perfectly simulates $G'_{0,i+1}$. Therefore, if \mathcal{A} distinguishes $G'_{0,i}$ from $G'_{0,i+1}$, then \mathcal{B} also distinguishes the IND-CPA game with the same advantage. This allows us to conclude that $G'_0 \approx_c G'_{0,q}$. We can run a symmetric argument, changing the way τ_1 is computed in G'_1 (i.e. by encrypting x_0 instead of x_1 in the request message) to show $G'_{0,q} \approx_c G'_1$ by the IND-CPA property of FHE. Putting everything together and noting that there are a polynomial number of hybrid experiments, we have that $G_0 \approx_c G_1$ as required. \square

5 Verifiability

In this section we aim to leverage the oblivious nature of the OPRF to extend our POPRF construction to achieve verifiability. We base our technique on the heuristic trick informally discussed in [ADDS21, Sec. 3.2] but with some modifications. In particular, we identify a blind-evaluation attack on this verifiability strategy in our context, the mitigation for which requires enlarging certain parameters. We then use cryptanalysis to study the security of our protocol, i.e. our construction does not reduce to a known (or even new but clean) hard problem. We view our analysis as an exploration into achieving secure protocols from bounded depth circuits which we hope has applications beyond this work.

We now describe our verifiability procedure, based on our OPRF presented in Figure 3 using the cut-and-choose method from [ADDS21, Sec. 3.2]. Intuitively, the client, \mathcal{C} , sends the server, \mathcal{S} , a set of known answer checkpoints amongst genuine OPRF queries. The client checks if the “check points” match the known answer values, if evaluations on the same points produce the same outputs and if evaluations on different points produce different outputs. \mathcal{C} may then assume that \mathcal{S} computed the (P)OPRF correctly provided the size of the checkpoint set is large enough. In more detail, let $\gamma = \alpha + \beta$ be the number of points \mathcal{C} submits to \mathcal{S} .

1. For some fixed \mathbf{t} , \mathcal{S} commits to κ points $\mathbf{z}_k^* := F_{\mathcal{A}}(\mathbf{t}, \mathbf{x}_k^*)$ for $k \in \mathbb{Z}_\kappa$ and publishes them (or sends them to \mathcal{C}); \mathcal{S} attaches a NIZK proof that these are well-formed, i.e. satisfy the relation

$$\mathfrak{R}_{\mathbf{t}} := \{\mathcal{A}, \mathbf{t}, \mathbf{z}_k^*; \mathbf{x}_k^* \text{ for } k \in \mathbb{Z}_\kappa \mid \mathbf{z}_k^* = F_{\mathcal{A}}(\mathbf{t}, \mathbf{x}_k^*)\}$$

2. \mathcal{C} wishes to evaluate α distinct points $\mathbf{x}_i^{(\alpha)}$ for $i \in \mathbb{Z}_\alpha$. It samples $\mathbf{x}_j^{(\beta)} \leftarrow_{\mathcal{S}} \{\mathbf{x}_k^*\}_{k \in \mathbb{Z}_\kappa}$ for $j \in \mathbb{Z}_\beta$. \mathcal{C} and shuffles the indices and submits these queries.

3. C receives (claimed) $\mathbf{z}_i^{(\alpha)} \stackrel{?}{=} F_{\mathbf{A}}(\mathbf{t}, \mathbf{x}_i^{(\alpha)})$ and $\mathbf{z}_j^{(\beta)} \stackrel{?}{=} F_{\mathbf{A}}(\mathbf{t}, \mathbf{x}_j^{(\beta)})$. C rejects if check points do not match, i.e. $\mathbf{z}_j^{(\beta)} \neq \mathbf{z}_k^*$. Otherwise C accepts.

We formalise the security definition with a game in Figure 4. In the first phase, the adversary provides n check points. Since the tag \mathbf{t} remains constant throughout we suppress it here. The experiment samples $m - n$ additional points. It executes the request algorithms and permutes the output according to $\rho \leftarrow_{\$} S_m$. In the second phase, the adversary receives the m requests and it must also output m responses. The adversary wins if all points have been correctly evaluated and there exists at least one additional point which does not correspond to a point evaluated under $F_{\mathbf{A}}$. We say a (P)OPRF is verifiable if the following advantage is negligible in the security parameter λ :

$$\text{Adv}_{\mathcal{F}, \mathcal{S}, \mathcal{A}}^{\text{verif}}(\lambda) = \Pr[\text{VERIF}_{\mathcal{F}, m, n}^{\mathbf{A}}(\lambda) = 1]$$

$\text{VERIF}_{\mathcal{F}, m, n}^{\mathbf{A}, \text{RO}}$
1 : $\text{pp} \leftarrow \mathcal{F}.\text{Setup}(\lambda)$
2 : $(\{\mathbf{x}_i\}_{i=0}^{n-1}, (\text{sk}, \text{pk})) \leftarrow \mathcal{A}; \quad \{\mathbf{x}_i\}_{i=n}^{m-1} \leftarrow_{\$} \mathbb{Z}_p^n$
3 : $\text{req}_i \leftarrow \mathcal{F}.\text{Request}(\text{pk}, \mathbf{t}, \mathbf{x}_i), \quad i \in \mathbb{Z}_m$
4 : $\rho \leftarrow_{\$} S_m$
5 : $\{\text{rep}'_i\}_{i=0}^{m-1} \leftarrow \mathcal{A}^{\text{Request, Finalise, RO}}(\text{req}_{\rho(0)}, \dots, \text{req}_{\rho(m-1)})$
6 : $\mathbf{z}_i \leftarrow \mathcal{F}.\text{Eval}(\text{sk}, \mathbf{t}, \mathbf{x}_i), \quad i \in \mathbb{Z}_m$
7 : if $\forall i \in [0, n-1] \mathcal{F}.\text{Finalise}(\text{rep}'_i) = \mathbf{z}_{\rho(i)} \wedge \exists j \in [n, m-1] \text{ s.t. } \mathcal{F}.\text{Finalise}(\text{rep}'_j) \neq \mathbf{z}_{\rho(j)}$
8 : return 1
9 : else return 0

Fig. 4. Verifiability Experiment for (P)OPRF

5.1 Verifiability from Levelled FHE

The heuristic we use to claim security is inspired by [CHLR18] which argues that evaluating a deep circuit in an FHE scheme supporting only shallow circuits is a hard problem. We pursue the same line of reasoning, albeit with significantly tighter security margins. That is, our assumption is significantly stronger than that in [CHLR18].

First, we will assume that the bootstrapping keys for the FHE scheme provided by C to S do not provide fully homomorphic encryption but restrict to a limited number of levels. More precisely, while our main construction Figure 3 and its implementation in Section 6 make use of *fully* homomorphic encryption (by relying on circular security), to target verifiability we pick parameters such that only Lines 4 and 6 of Algorithm 2 cost a bootstrapping operation, i.e. all linear operations are realised without bootstrapping. We note that this significantly increases the parameters required to evaluate our VOPRF candidate.

Specifically, we observe that our POPRF in Algorithm 2 can then be evaluated in depth two and S only receives two bootstrapping keys, presumably preventing it from computing higher depth circuits.

Secondly, we assume that C provides non-unique encodings of its inputs. For example if its input is $\mathbf{x} \in \mathbb{Z}_2^n$ it may provide them as a vector in e.g. $\{-1, 0, 1\}^n$ equivalent to \mathbf{x} modulo 2. This forces a cheating server to perform normalisation to prevent a cheating POPRF circuit from being computed in depth two.

Third, here we also assume that the native TFHE plaintext modulus P is coprime with $p = 2$. This is because when $p \mid P$, the non-unique encoding approach we utilise here is futile as a unique encoding can be recovered by considering $P/2 \cdot m$ where m is the non-unique encoding of a value $\in \{0, 1\}$.¹⁷

Parameters. We estimate parameters of our VOPRF in Table 3. While the online cost (134.6kB) is perhaps acceptable in some settings, the public key size (4077.8MB) is clearly beyond the realm of practicality. These parameters are dictated by requiring $P \approx n_p$ so that $\mathbf{A} \cdot (\mathbf{y}, 1)$ can be computed over the Integers. We thus base our parameters on `PARAM_MESSAGE_8_CARRY_0` in `tfhe-rs` (but replace $P = 2^8$ with a prime close to it).

Table 3. Post-quantum VOPRF candidates in the literature

work	assumption	r	communication cost	flavour	model
[ADDS21]	R(LWE) & SIS	2	> 128GB	verifiable	malicious, QROM
[Bas23]	SIDH	2	8.7MB	verifiable	malicious, <i>ts</i> , ROM
Section 5	heuristic	2	0.5MB + 4077.8MB + 57.86KB + 0.8kB + 37.6kB + 96.2kB	verifiable	malicious, ROM
Section 5	heuristic	2	0.5MB + 4077.8MB + 57.86KB + 0.8kB + 3.4kB + 96.2kB	verifiable	malicious, ROM $L = 64$, per query

The column “r” gives the number of rounds. ROM is the random oracle model, QROM the quantum random oracle model and “pp” stands for “preprocessing” and “ts” for “trusted setup”. When reporting on our work, the summands are: check point size, pk size, pk proof size, client message size, client message proof size, server message size. Our client message proofs can be amortised to e.g. $217.1\text{kB}/64 = 3.4\text{kB}$ per query when amortising over $L = 64$ queries.

5.2 Cryptanalysis

We explore cheating strategies of a malicious server.

Guessing. Assume the adversary guesses the positions of the check points in order to make C accept an incorrect output. Recall that $\gamma = \alpha + \beta$ is the number of points C submits to S , where there are β such check point positions

¹⁷ We would also expect that P should be somewhat far away from a power of two, but we leave the detailed analysis for future work.

of a possible κ choices. Thus, if we assume semantic security of the underlying homomorphic encryption scheme then the probability of an adversarial server guessing correctly is bounded by the probability it selects the positions of a particular check point, for each of the β check points. We obtain a probability $(1/\kappa)^\beta \cdot 1/\binom{\gamma}{\beta}$ of guessing correctly.

Interpolation. We consider an adversarial S that uses a circuit F'_{pei} , of depth at most 2, to win the verifiability experiment. At a high level, it solves a quadratic system of equations and then uses the second level to do modular reduction. More precisely, it chooses F'_{pei} such that it agrees on the κ points with of the POPRF circuit $F_{\text{pei}}(\mathbf{A}, \mathbf{t}, \cdot)$, but that can differ elsewhere. Since the server knows the checkpoints, this can be trivially done by the server. To prevent such an attack, one would need to publish $\kappa = n + \frac{n}{2}(n+1) + \mathcal{O}(1)$ check points where n is the input/output dimension. This, implies no quadratic polynomial interpolation exists. If we let $\kappa = 128^3$, then the communication cost of checkpoints is approximately an additional 64MB, which can be reduced to 32MB by generating the input values from a seed. If the assumption is weakened such that we only require $\kappa = 128^2$ then this becomes only 0.5MB.

Check and cheat. Finally, we consider that the adversary is able to construct a shallow cheating circuit, which is described in Algorithm 3, with $[\cdot]$ denoting a homomorphic encryption of a value. Intuitively, it homomorphically checks the clients inputs against known answers and then homomorphically selects which output to return. This circuit has depth three. Here we use the non-unique encoding of the input to force a malicious server to perform the initial mod p operation, costing a level.

Algorithm 3 Cheating Circuit

Input: non-unique encoding $[\mathbf{y}] \in \mathbb{Z}_p^{np}$; checkpoints $\{\mathbf{y}_j^*\}$; $H : \mathbb{Z}_p^{np} \rightarrow \mathbb{Z}_q^n$ any function

```

 $[r], [\text{found}] \leftarrow 0, 0$ 
for all  $\mathbf{y}_j^*$  do
   $[d] \leftarrow [\mathbf{y}] - \mathbf{y}_j^*$ 
   $[d_i^{(0)}] \leftarrow \text{LUT}_{x \bmod p}(d_i) \forall i \in \mathbb{Z}_{n_p}$  ▷ mod  $p$ , depth 1
   $[h] \leftarrow \sum_{i \in \mathbb{Z}_{n_p}} d_i^{(0)}$ 
  if  $[h] = 0$  then ▷ CMUX, depth 1
     $[r] \leftarrow [r] + \mathbf{y}_j^*$ ;  $[\text{found}] \leftarrow [\text{found}] + 1$ 
  else
     $[r] \leftarrow [r] + \mathbf{0}$ ;  $[\text{found}] \leftarrow [\text{found}] + 0$ 
  end if
end for
return  $\text{CMUX}_{[\text{found}]}([r], H([\mathbf{y}]))$  ▷ depth 1

```

5.3 Compression Potential

The central reason for the large parameters reported above is that we must force a cheating server to compute a circuit in depth three whereas a honest server only requires depth two. We achieve this by making the client submit non-unique encodings and thus need to pick parameters that $\mathbf{A} \cdot (\mathbf{y}, 1)$ can be computed over \mathbb{Z} , i.e. a plaintext space $\approx \mathbb{Z}_{n_p}$. A first approach to circumvent this issue is by making an even stronger assumption: our cheating circuits all need to homomorphically sum over κ values related to our checkpoints. In our analysis, we assume additions are “free”, i.e. hardly increase the noise, but as we increase the number of check points this assumption breaks down. Indeed, $\kappa \approx n^2$ already suggests a noise growth of n , plausibly too big for our OPRF parameters. This, of course, assumes this noise growth cannot be managed more efficiently by a cheating server and is thus a stronger assumption, as mentioned.

A second approach is to reduce the bootstrapping depth of an honest server. Recall, that we rely on our final bootstrapping to argue **Ax-hiding** of the evaluation. If, instead, we could rely solely on the circuit privacy provided by bootstrapping after computing $\mathbf{A} \cdot (\mathbf{y}, 1)$, we would obtain a depth one circuit for the honest server. Thus, a cheating server in depth two would not break security and we could drop the non-unique encoding and evaluation over \mathbb{Z} blowing up our parameters. We leave exploring these approaches for future work.

6 Proof-of-Concept Implementation

We demonstrate that we can build a simple and efficient implementation of the basic oblivious pseudorandom function given in Figure 3. Our implementation is written in Rust and provided open-source.¹⁸ It makes use of Zama’s `tfhe-rs` FHE library [CGGI20] for writing FHE functionalities using TFHE. The implementation itself requires less than 700 lines of code for instantiating the core functionality (not including tests and examples).¹⁹

Implementation choices. The matrix $\mathbf{G}_{\text{out}} \in \mathbb{Z}_q^{m \times m_p}$ is constructed (as in Section 2.6) by sampling each entry uniformly in \mathbb{Z}_3 . The matrix $\mathbf{G}_{\text{inp}} \in \mathbb{Z}_q^{n_q \times n}$ is sampled with an identity matrix in the top $n \times n$ component, and the remaining $(n_q - n)$ rows as random vectors in \mathbb{Z}_3 (i.e. in systematic form, see Section 2.6). The matrix \mathbf{A} is sampled randomly in $\mathbb{Z}_2^{m_p \times n_p}$. The dimensions of these matrices are set using the values given in Table 2. The matrix \mathbf{A}_t is constructed as the output of a random oracle RO_{key} , where for each column $\mathbf{A}[i]$, the random oracle outputs column $\mathbf{A}_t[i] \in \mathbb{Z}_2^{m_p}$ by evaluating $\text{RO}_{\text{key}}(t, \mathbf{A}[i], \text{"vopr_fhe_rom_key"})$. The random oracle RO_{key} is implemented using SHA-384, while the random oracle RO_{fin} , which is used for finalising the client input (as in [TCR⁺22]), is implemented using SHA-256. Note that in our implementation, all matrices are transposed and therefore all multiplication operations are therefore reversed.

¹⁸ <https://anonymous.4open.science/r/oprf-fhe-C726>

¹⁹ We do not implement the zero-knowledge arguments that are used as part of the main protocol construction, but provide bandwidth estimates for these in Appendix B.

Table 4. Computational runtimes for running OPRF functionality.

(P)OPRF function	Time
C : F_{poprf} .KeyGen	220ms
C : F_{poprf} .Request	36ms
S : F_{poprf} .BlindEval (“slow”)	14.4s
S : F_{poprf} .BlindEval (“full”)	123ms
C : F_{poprf} .Finalise	0.1ms

The client public key includes the size of both the bootstrapping and key switching key used in the implementation of Concrete.

Table 5. Sizes of cryptographic material, RAM usage, and bootstrapping depth per ciphertext.

Data point	Size
C: Secret key	0.02MB
C: Public key	66.8MB
C: Request	2.02MB
S: Response	1.01MB

tfhe-rs implementation and parameters. The Rust implementation of Concrete that we use corresponds to the `tfhe-rs` library (version “0.2.4”).²⁰ with parameter set `PARAM_MESSAGE_2_CARRY_0`. Ideally, we would implement our PRF candidate by first computing with a native TFHE plaintext modulus 2, followed by a bootstrapped key switch to native TFHE plaintext modulus 3. However, as it stands, the `tfhe-rs` API does not support this. We, thus, implement a “slow” variant, where we perform all computations modulo 4 (implied by the parameter set `PARAM_MESSAGE_2_CARRY_0`) where we carefully bootstrap and reduce modulo 3 to avoid overflows. That is, \mathbb{Z}_3 is implemented using the set $\{0, 1, 2\} \in \mathbb{Z}_4$.²¹ We then also give estimates (rather than benchmarks) for a “full” variant, where we avoid these additional bootstrapping operations to work around the API limitation.

Benchmarks. The results that we discuss below were acquired by using a server with 96 Intel Xeon Gold 6252 CPU @ 2.10GHz cores and 768 GB of RAM. Server evaluation is run with parallelisation enabled, meaning that each multiplication of the client input encrypted vector with a matrix column is run in its own separate thread, but we only use 64 threads/cores. Client evaluations use only a single core. Each of the benchmarks was established after running it ten times and taking the average runtime.

Runtimes. As shown in Table 4, key generation (`KeyGen`) is the heaviest client-side: only taking 220ms. Even so, this operation can be regarded as a one-time cost in many applications. For example, considering OPAQUE [JKX18], clients and servers already register persistent identifiers for each other (such as the client-specific OPRF key). Therefore, the client keypair can be registered as part of this process. Similarly for Privacy Pass [DGS⁺18], the issuance phase of the protocol does not discount clients from registering persistent information that they use whenever they make VOPRF evaluations (which could include this key information). As a result, in many applications, clients will generate a single FHE keypair and use that over multiple interactions with the server.

²⁰ <https://docs.rs/tfhe/0.2.4/tfhe/index.html>

²¹ We note that this implementation is incompatible with our VOPRF candidate.

For the online client-side algorithms (`Request` and `Finalise`) that must be run on every invocation of the OPRF, it is clear that the overheads are minimal, and are suitable even for relatively constrained devices. For the online server-side algorithm (`BlindEval`), the runtime is *estimated* as just 123ms, which may be quick enough for certain applications that have a hard requirement to ensure post-quantum security. Previous classical constructions of (P)OPRFs, such as [TCR⁺22] take only a few milliseconds to run the server evaluation algorithm, and so the efficiency gap between our “full” FHE-based approach and previous work is evident, but fairly minimal.

Bandwidth and storage costs. As shown in Table 5, the client public key dominates the bandwidth and storage costs for any given run-through of the protocol, at around 70MB. However, since the client key pair can be used across multiple invocations of the protocol, the bandwidth cost does not occur during subsequent evaluations (the “online phase”). The online phase of the protocol incurs a total communication cost of approximately 3MB, since ciphertexts are stored in M -LWE format in the `tfhe-rs` library. However, as remarked in Section 3.2, much smaller communication costs are possible by carefully selecting the format (LWE instead of M -LWE), compressing and amortising.

Acknowledgements

We thank Christian Weinert for discussing MPC approaches with us; we thank Ilaria Chillotti, Ben Curtis and Jean-Baptiste Orfila for answering questions about TFHE and Concrete; we thank Nicolas Gama for answering questions about TFHE. We thank Ward Beullens and Gregor Seiler for answering questions about LaBRADOR and sharing their size estimation Pari/GP script with us.

The research of Martin Albrecht was supported by EPSRC grants EP/S02-0330/1, EP/S02087X/1 and by the European Union Horizon 2020 Research and Innovation Program Grant 780701. Part of this work was done while Martin Albrecht was at Royal Holloway, University of London, and while Alex Davidson was at Brave Software. The majority of this research was done while Amit Deo was at Crypto Quantique.

References

- ADDS21. Martin R. Albrecht, Alex Davidson, Amit Deo, and Nigel P. Smart. Round-optimal verifiable oblivious pseudorandom functions from ideal lattices. In Garay [Gar21], pages 261–289. Full version available at <https://eprint.iacr.org/2019/1271>. 1.1, 1.2, 10, 1, 5, 3
- AGPS20. Martin R. Albrecht, Vlad Gheorghiu, Eamonn W. Postlethwaite, and John M. Schanck. Estimating quantum speedups for lattice sieves. In Moriai and Wang [MW20], pages 583–613. 11, 2.6
- APS15. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015. 2.4

- ARS⁺15. Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. Ciphers for MPC and FHE. In Oswald and Fischlin [OF15], pages 430–454. 1
- Bas23. Andrea Basso. A post-quantum round-optimal oblivious PRF from isogenies. Cryptology ePrint Archive, Report 2023/225, 2023. <https://eprint.iacr.org/2023/225>. 1.2, 1, 3
- BCG⁺22. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Nicolas Resch, and Peter Scholl. Correlated pseudorandomness from expand-accumulate codes. In Dodis and Shrimpton [DS22], pages 603–633. 1.2
- BdMW16. Florian Bourse, Rafaël del Pino, Michele Minelli, and Hoeteck Wee. FHE circuit privacy almost for free. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 62–89. Springer, Heidelberg, August 2016. 2.5
- BGV11. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. Fully homomorphic encryption without bootstrapping. Cryptology ePrint Archive, Report 2011/277, 2011. <https://eprint.iacr.org/2011/277>. 2.5
- BIP⁺18. Dan Boneh, Yuval Ishai, Alain Passelègue, Amit Sahai, and David J. Wu. Exploring crypto dark matter: New simple PRF candidates and their applications. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 699–729. Springer, Heidelberg, November 2018. Full version available at <https://eprint.iacr.org/2018/1218>. 1, 1.1, 5, 1.2, 1.3, 1, 2.6, 2.6, 2.6
- BKM⁺21. Andrea Basso, Péter Kutas, Simon-Philipp Merz, Christophe Petit, and Antonio Sanso. Cryptanalysis of an oblivious PRF from supersingular isogenies. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 160–184. Springer, Heidelberg, December 2021. 1.2
- BKW20. Dan Boneh, Dmitry Kogan, and Katharine Woo. Oblivious pseudorandom functions from isogenies. In Moriai and Wang [MW20], pages 520–550. 1.2, 1
- BS22. Ward Beullens and Gregor Seiler. LaBRADOR: Compact proofs for R1CS from module-SIS. Cryptology ePrint Archive, Report 2022/1341, 2022. <https://eprint.iacr.org/2022/1341>. 1.1, B, B.1, B.1
- CCKK21. Jung Hee Cheon, Wonhee Cho, Jeong Han Kim, and Jiseung Kim. Adventures in crypto dark matter: Attacks and fixes for weak pseudorandom functions. In Garay [Gar21], pages 739–760. 2.6
- CDKS21. Hao Chen, Wei Dai, Miran Kim, and Yongsoo Song. Efficient homomorphic conversion between (ring) LWE ciphertexts. In Kazue Sako and Nils Ole Tippenhauer, editors, *ACNS 21, Part I*, volume 12726 of *LNCS*, pages 460–479. Springer, Heidelberg, June 2021. 3.2
- CGGI20. Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, January 2020. 1.1, 2.5, 2.5, 3.2, 6, B.1
- CHL22. Sílvia Casacuberta, Julia Hesse, and Anja Lehmann. Sok: Oblivious pseudorandom functions. In *7th IEEE European Symposium on Security and Privacy, EuroS&P 2022*, pages 625–646. IEEE, 2022. 1.2, 3.1
- CHLR18. Hao Chen, Zhicong Huang, Kim Laine, and Peter Rindal. Labeled PSI from fully homomorphic encryption with malicious security. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1223–1237. ACM Press, October 2018. 1, 1.1, 1.3, 5.1

- CKKS17. Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. Homomorphic encryption for arithmetic of approximate numbers. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 409–437. Springer, Heidelberg, December 2017. 2.5
- DGH⁺21. Itai Dinur, Steven Goldfeder, Tzipora Halevi, Yuval Ishai, Mahimna Kelkar, Vivek Sharma, and Greg Zaverucha. MPC-friendly symmetric cryptography from alternating moduli: Candidates, protocols, and applications. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 517–547, Virtual Event, August 2021. Springer, Heidelberg. 1, 1.2, 9, 1, 2.6, 2.6
- DGS⁺18. Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. Privacy pass: Bypassing internet challenges anonymously. *PoPETs*, 2018(3):164–180, July 2018. 1, 6
- DM15. Léo Ducas and Daniele Micciancio. FHEW: Bootstrapping homomorphic encryption in less than a second. In Oswald and Fischlin [OF15], pages 617–640. 2.5
- DS16. Léo Ducas and Damien Stehlé. Sanitization of FHE ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 294–310. Springer, Heidelberg, May 2016. 2.5
- DS22. Yevgeniy Dodis and Thomas Shrimpton, editors. *CRYPTO 2022, Part II*, volume 13508 of *LNCS*. Springer, Heidelberg, August 2022. 6
- DvW21. Léo Ducas and Wessel P. J. van Woerden. NTRU fatigue: How stretched is overstretched? In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 3–32. Springer, Heidelberg, December 2021. 2.4
- ECS⁺15. Adam Everspaugh, Rahul Chatterjee, Samuel Scott, Ari Juels, and Thomas Ristenpart. The pythia PRF service. In Jaeyeon Jung and Thorsten Holz, editors, *USENIX Security 2015*, pages 547–562. USENIX Association, August 2015. 1
- FV12. Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012. <https://eprint.iacr.org/2012/144>. 2.5
- Gar21. Juan Garay, editor. *PKC 2021, Part II*, volume 12711 of *LNCS*. Springer, Heidelberg, May 2021. 6
- Gen09. Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. crypto.stanford.edu/craig. 2.5, 2.5
- Gol04. Oded Goldreich. *Foundations of Cryptography*, volume 2. Cambridge University Press, 2004. 2
- GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, August 2013. 2.5, B.1
- HMR23. Lena Heimberger, Fredrik Meisingseth, and Christian Rechberger. Oprfs from isogenies: Designs and analysis. Cryptology ePrint Archive, Paper 2023/639, 2023. <https://eprint.iacr.org/2023/639>. 1.2, 1
- HPS96. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A new high speed public key cryptosystem, 1996. Draft Distributed at Crypto’96, available at <http://web.securityinnovation.com/hubfs/files/ntru-orig.pdf>. 2.4

- JKR18. Stanislaw Jarecki, Hugo Krawczyk, and Jason Resch. Threshold partially-oblivious PRFs with applications to key management. Cryptology ePrint Archive, Report 2018/733, 2018. <https://eprint.iacr.org/2018/733>. 3.1
- JKX18. Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 456–486. Springer, Heidelberg, April / May 2018. 1, 6
- Joy21. Marc Joye. Guide to fully homomorphic encryption over the [discretized] torus. Cryptology ePrint Archive, Report 2021/1402, 2021. <https://eprint.iacr.org/2021/1402>. 1.1, 2.5, 2.5, 13
- Klu22. Kamil Klucznik. Circuit privacy for fhew/tfhe-style fully homomorphic encryption in practice. Cryptology ePrint Archive, Paper 2022/1459, 2022. <https://eprint.iacr.org/2022/1459>. 2.5, B.1
- LDK⁺22. Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. 3.2, B.1, B.1, B.2
- LNP22. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In Dodis and Shrimpton [DS22], pages 71–101. 1.1, B.1, B.1, B.2, B.2, B.3
- LS15. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, June 2015. 5, 6
- MAT22. MATZOV. Report on the Security of LWE: Improved Dual Lattice Attack. Available at <https://doi.org/10.5281/zenodo.6412487>, April 2022. 2.4
- MP20. Daniele Micciancio and Yuriy Polyakov. Bootstrapping in FHEW-like cryptosystems. Cryptology ePrint Archive, Report 2020/086, 2020. <https://eprint.iacr.org/2020/086>. 1.1
- MW20. Shiho Moriai and Huaxiong Wang, editors. *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*. Springer, Heidelberg, December 2020. 6
- OF15. Elisabeth Oswald and Marc Fischlin, editors. *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*. Springer, Heidelberg, April 2015. 6
- OPP14. Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2014. 2.5
- SG98. Victor Shoup and Rosario Gennaro. Securing threshold cryptosystems against chosen ciphertext attack. In Kaisa Nyberg, editor, *EUROCRYPT'98*, volume 1403 of *LNCS*, pages 1–16. Springer, Heidelberg, May / June 1998. 1.1
- SHB21. István András Seres, Máté Horváth, and Péter Burcsi. The legendre pseudorandom function as a multivariate quadratic cryptosystem: Security and applications. Cryptology ePrint Archive, Report 2021/182, 2021. <https://eprint.iacr.org/2021/182>. 1.2, 1

- TCR⁺22. Nirvan Tyagi, Sofia Celi, Thomas Ristenpart, Nick Sullivan, Stefano Tessaro, and Christopher A. Wood. A fast and simple partially oblivious PRF, with applications. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 674–705. Springer, Heidelberg, May / June 2022. 1, 1.1, 2.1, 2.2, 1, 2.2, 2, 2.2, 3, 3, 4, 6, 6
- Unr12. Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 135–152. Springer, Heidelberg, April 2012. 1.1

A Alternative heuristic approach for \mathbf{Ax} -hiding

We conjecture that the following approach, relying on the form of programmable bootstrapping when implementing an LUT, heuristically achieves \mathbf{Ax} -hiding. Programmable bootstrapping of an LWE ciphertext with look-up table LUT_f consists of three steps: blind rotation, sample extraction and key-switching. In particular, blind rotation begins with an LWE ciphertext $\mathbf{c} = (\mathbf{a}, b = \mathbf{a} \cdot \mathbf{s} + e + \mu \cdot (Q/P))$ encrypting $\mu \in \mathbb{Z}_P$ and a programmed evaluation polynomial $v(X)$ for function f . Note that we do not wish to hide f and thus $v(X)$. Furthermore, let (\mathbf{a}_z, b_z) be some encryption of zero under the same LWE secret key.

We first compute $(\mathbf{a}, b) \leftarrow (\mathbf{a}, b) + (v \cdot \mathbf{a}_z + \mathbf{e}', v \cdot b_z + \mathbf{e}'')$ for some short $v, \mathbf{e}', \mathbf{e}''$ to rerandomise (\mathbf{a}, b) , especially the \mathbf{a} part (using the LWE assumption). Then using the bootstrapping key, blind rotation (homomorphically) multiplies $v(X)$ by X^{k_i} where k_i is either 0 or some value depending on \mathbf{a} or b for $i = 0, \dots, n-1$. Whether k_i is 0 or not depends on the bootstrapping key. In more detail, these blind rotations are realised using CMux gates acting on some accumulator α which is update to $\alpha \leftarrow G^{-1}((X^{a_i} - 1) \cdot \alpha) \cdot (\mathbf{E} + \mu \cdot \mathbf{G}^T) + \alpha$ where \mathbf{G} is a known gadget matrix, $G^{-1}(\cdot)$ is a matching (possibly randomised) decomposition function and \mathbf{E} is some small error. At the end of the blind rotation phase, we have an encryption of $X^{\lfloor (b - \mathbf{a} \cdot \mathbf{s}) \cdot 2n/Q \rfloor} \cdot v(X) = X^{\lfloor (e + \mu \cdot (Q/P)) \cdot 2n/Q \rfloor} \cdot v(X)$. Here $\lfloor (b - \mathbf{a} \cdot \mathbf{s}) \cdot 2n/Q \rfloor$ indicates that $(b - \mathbf{a} \cdot \mathbf{s})$ is mapped from values in $\{0, \dots, Q-1\}$ to values in $\{0, \dots, 2n-1\}$ via rounding. We then perform sample extraction to get an LWE encryption of $f(\lfloor (P/Q) \cdot (b - \mathbf{a} \cdot \mathbf{s}) \rfloor)$.

Now, in our variant, we initialise the accumulator to $v(x)$ (and not $X^b \cdot v(x)$). We then run the CMUXes realising the blind rotation in some random order. Finally, we rotate the result in the accumulator α by X^b . Here, the dependency on (\mathbf{a}, b) is almost removed, using the LWE assumption on $v \cdot \mathbf{a}_z + \mathbf{e}'$. In the final rotation step we do get a dependency based on the original computation, but a key difference is that we do not need a CMUX to perform this rotation. This is then followed by a sample extraction that essentially picks out a single coefficient of the noisy RLWE ciphertext, partially obscuring the rotation. We stress, however, that this argument is heuristic, as b is not independent of \mathbf{a} if the secret key is known.

B Client Non-Interactive Zero Knowledge Proofs

A key part of security against malicious clients is the non-interactive zero knowledge proof that a POPRF request is well-formed. This requires a proof system that can (a) show that the public key material of the FHE scheme is well-formed and (b) show that the accompanying ciphertext is an encryption of a valid input. To make things modular, we discuss how to prove (a) and (b) separately. The overall system can be trivially obtained by combining the two proofs in a straight-forward manner.

For the VOPRF construction, there are two main complications as far as the client-side zero-knowledge proofs are concerned. The first is that the public key material sent by the client must only permit two bootstrapping operations. The second is that PRF inputs must have non-unique encodings. To avoid repetition, we describe the client proofs for (a) in the more complicated context of VOPRFs and derive cost estimates for the OPRF construction as a simpler case.

We provide our scripts computing our size estimates in Appendices E and F. The former is based off a script provided by the authors of [BS22].

B.1 Public Keys

There are two components to the public key material sent by the client: a normal TLWE public key/commitment and a TFHE bootstrapping key. A TFHE bootstrapping key consists of three components: A standard GSW key [GSW13], a standard bootstrapping key and a key-switching key. The latter two are used in TFHE’s (programmable) bootstrapping procedure whereas the first is used to create the bootstrapping key. One could remove the GSW key, but then the zero-knowledge proofs become far more expensive because there are large elements that form part of the witness (namely the uniform elements used in the creation of a private GSW encryption). Note that the bootstrapping key takes the form of GSW ciphertexts over *rings*, i.e. over $\mathcal{R}_Q = \mathbb{Z}_Q[X]/(X^d + 1)$ whereas the key-switching key is a collection of *plain* LWE ciphertexts, i.e. over \mathbb{Z}_Q^e .

To perform normal homomorphic operations, all that is required is the ciphertexts. In particular, we do not really need to use a (T)LWE public key (denoted \mathbf{c}_{pk} below) for this. However, the circuit privacy techniques of Klucznik [Klu22] does use the public key to re-randomise ciphertexts. Therefore, it is important that we include \mathbf{c}_{pk} in the proofs.

In the VOPRF construction, we will only be providing TFHE allowing two bootstrapping operations. After each bootstrapping, the TFHE encryption key changes. Let $\bar{s}^{(\tau)} = (\bar{s}_0^{(\tau)}, \dots, \bar{s}_{e-1}^{(\tau)}) \leftarrow_{\$} \mathbb{Z}_2^e$ denote the TFHE key after the τ -th bootstrapping for $\tau = 0, 1, 2$. Also, take the gadget matrix permitting *approximate* decompositions to be

$$\mathbf{G} = ([Q/B], \dots, [Q/B^\ell]) \otimes \mathbf{I}_2$$

for decomposition parameters B and ℓ . The encryption key used to produce the τ -th bootstrapping key for $\tau \in \{1, 2\}$ will be $\tilde{s}^{(\tau)} = \sum_{i=0}^{d-1} \tilde{s}_i^{(\tau)} \cdot X^i \leftarrow_{\$} \mathcal{R}_2$.

We will denote error distributions for LWE/MLWE assumptions by χ and $\bar{\chi}$ respectively.

Key Material. Overall, the key material sent from the client to the server is:

- **Root TLWE Public Key/Root Secret Key Commitment:**

$$\mathbf{b}_{\text{pk}} = \mathbf{A}_{pp} \cdot \mathbf{r}_{\text{com}} + \mathbf{e}_{\text{com}} + \lfloor Q/2 \rfloor \cdot \begin{bmatrix} \bar{s}^{(0)} \\ \mathbf{0}^{e_{\text{com}} - e} \end{bmatrix} \in \mathbb{Z}_Q^{e_{\text{com}}}$$

$$\mathbf{c}_{\text{pk}} = \mathbf{A}_{\text{pk}} \cdot \bar{s}_0 + \mathbf{e}_{\text{pk}} \in \mathbb{Z}_Q^{e \log Q}$$

where $\mathbf{A}_{pp} \in \mathbb{Z}_Q^{e_{\text{com}} \times e_{\text{com}}}$ is from the public parameters, $\mathbf{A}_{\text{pk}} \leftarrow_{\$} \mathbb{Z}_Q^{e \log Q \times e}$, $\mathbf{e}_{\text{com}}, \mathbf{r}_{\text{com}} \leftarrow_{\$} \chi_{\text{com}}^{e_{\text{com}}}$ and $\mathbf{e}_{\text{pk}} \leftarrow_{\$} \chi^{e \log Q}$.

- **GSW Public Keys for bootstrapping:** For $\tau \in \{1, 2\}$

$$\mathbf{A}_{GSW}^{(\tau)} := \begin{bmatrix} \bar{s}^{(\tau)} & \mathbf{a}^{(\tau)} \\ \bar{s}^{(\tau)} & \mathbf{a}^{(\tau)} + \mathbf{e}^{(\tau)} \end{bmatrix} \in \mathcal{R}_Q^{2 \times 2}$$

where $(\mathbf{a}^\tau)^\top \leftarrow_{\$} \mathcal{R}_Q^2$, $\mathbf{e}^{(\tau)} \leftarrow_{\$} \bar{\chi}^2$.

- **Bootstrapping Keys:** For $i = 0, \dots, e - 1$ and $\tau \in \{1, 2\}$:

$$\text{bsk}[\tau, i] := \mathbf{A}_{GSW}^{(\tau)} \cdot \mathbf{R}_i^{(\tau)} + \mathbf{E}_i^{(\tau)} + \bar{s}_i^{(\tau-1)} \cdot \mathbf{G}$$

where $\mathbf{R}_i^{(\tau)} \leftarrow_{\$} \mathcal{R}_2^{2 \times 2\ell}$, $\mathbf{E}_i^{(\tau)} \leftarrow_{\$} \bar{\chi}^{2 \times 2\ell}$.

- **Key-Switching Keys:** For $i = 0, \dots, d - 1, j = 0, \dots, \tilde{\ell} - 1, \tau \in \{1, 2\}$,

$$\text{ksk}[\tau, i, j] := \left(\mathbf{a}_{i,j}^{(\tau)}, \sum_{k=0}^{e-1} (\mathbf{a}_{i,j}^{(\tau)})_k \cdot \bar{s}_k^{(\tau)} + e_{i,j}^{(\tau)} + \bar{s}_i^{(\tau)} \cdot \lfloor Q/\tilde{B}^j \rfloor \right),$$

where $\mathbf{a}_{i,j}^{(\tau)} \leftarrow_{\$} \mathbb{Z}_Q^e$, $e_{i,j}^{(\tau)} \leftarrow_{\$} \chi$ and $\tilde{B}, \tilde{\ell}$ are decomposition parameters.

In order to set some example parameters, we rely on the parameters from our implementation and only consider the simpler non-verifiable (P)OPRF case. In particular these parameters are (PARAM_MESSAGE_2_CARRY_0) defined in the `tfhe-rs` library [CGGI20] with $e_{\text{com}} = 2^{11}$ and χ_{com} a discrete Gaussian with standard deviation $\sigma_{\text{com}} = 2^2$ (see Section 2.4 for setting the latter). The TFHE scheme uses discretisation over a torus where the resolution of the discretisation allows us to pick modulus $Q \approx 2^{32}$ and the other parameters are $e = 656, d = 1,024, \ell = 2, \tilde{\ell} = 4$. Note that our description of the bootstrapping key is written in terms of R -LWE for simplicity. However, our implementation parameters use a ring dimension of 512 with rank 2, explaining why $d = 1,024$ in our analysis.

If a random seed is expanded to generate all of the random elements and considering the OPRF case, only the root public key, \mathbf{A}_{GSW}^1 , $\text{bsk}[1, i]$ and $\text{ksk}[1, i, j]$ need to be sent (when calculating $\text{ksk}[1, i, j]$, we use $\bar{s}^{(1)} = \bar{s}^{(0)}$). Therefore, the size of the key material is

$$(e_{\text{com}} + e \log Q) \cdot \log Q + 2 d \log Q + (2 \cdot 2 \ell) \cdot e \cdot d \log Q + d \tilde{\ell} \cdot \log Q.$$

which is dominated by $2e \cdot 2\ell \cdot d \log Q$.

Finally, to pack multiple LWE response ciphertexts into a *single* RLWE ciphertext, we need to add some key-switching material to the public keys. These key-switching keys are associated to a total of $\log(d)$ automorphisms and therefore consist in a total of $\log(d) \cdot d \cdot \tilde{\ell}$ noisy LWE products. Using the parameters above, we end up adding 72kB to the public key material. Overall, we get around 16.8MB of public key material sent by the client for the OPRF, which we obtain by also dropping some lower-order bits [LDK⁺22]. The VOPRF parameters `PARAM_MESSAGE_8_CARRY_0` use larger parameters $Q \approx 2^{64}, d = 32768, \ell = 2, \tilde{\ell} = 5$ to ensure that two bootstrappings on the server side suffices for PRF evaluation. Unfortunately, this results in around 4091MB of public key material.

Proofs. Essentially, the well-formedness of the public key material is a large “noisy” linear system where some of the equations are over \mathcal{R}_Q and others are over \mathbb{Z}_Q . Additionally, some of the solution is binary i.e. $\bar{s}^{(\tau)}, \tilde{s}^{(\tau)}$ and \mathbf{R}_i^τ , while the rest of the solution i.e. $\mathbf{r}_{\text{com}}, \mathbf{e}_{\text{com}}, \mathbf{e}_{\text{pk}}, \mathbf{E}_i^{(\tau)}$ and $e_{i,j}^{(\tau)}$ is small.

We give proof sizes both in the zero-knowledge proof system from [LNP22] and for a zero-knowledge variant of [BS22]. We rely on the former for proving ciphertexts, so we chose to also give the costs under this proof system for the FHE public keys.

[LNP22]. In order to estimate the cost of the zero-knowledge proof, we rewrite the statement being proved in a form compatible with [LNP22]. We will have the unknowns $\bar{s}^{(0)}, \bar{s}^{(1)}, \bar{s}^{(2)} \in \mathbb{Z}_2^e, \|\mathbf{r}_{\text{com}}\| \leq \beta_{\text{com}}, \tilde{s}^{(0)}, \tilde{s}^{(1)} \in \mathcal{R}_2,$

$$\{\mathbf{r}_{i,j}^{(\tau)} \in \mathcal{R}_2 : i \in \{0, \dots, e-1\}, j \in \{0, \dots, 2\ell-1\}, \tau \in \{1, 2\}\}$$

where $\mathbf{r}_{i,j}^{(\tau)}$ is the j -th column of $\mathbf{R}_i^{(\tau)}$ and also

$$\{e_{i,j}^{(\tau)} : i \in \{0, \dots, d-1\}, j \in \{0, \dots, \tilde{\ell}-1\}, \tau \in \{1, 2\}\}.$$

We rewrite the statement as follows:

- (Binary $\bar{s}^{(0)}, \bar{s}^{(1)}, \bar{s}^{(2)}, \tilde{s}^{(1)}, \tilde{s}^{(2)}, \{\mathbf{r}_{i,j}^{(\tau)}\}$) Take \mathbf{x} to be the concatenation of coefficients of $\bar{s}^{(0)}, \bar{s}^{(1)}, \bar{s}^{(2)}, \tilde{s}^{(1)}, \tilde{s}^{(2)}$ and all of the $\mathbf{r}_{i,j}$. Then $\langle \mathbf{x}, \mathbf{x} - \mathbf{1} \rangle \equiv 0 \pmod{Q}$ and $\|\mathbf{x}\|$ short enough so that the former holds over \mathbb{Z} too. This implies that these components are binary.
- (**Root public key/Root Secret Key Commitment**)

$$\|\mathbf{r}_{\text{com}}\| \leq \beta_{\text{com}},$$

$$\left\| \mathbf{b}_{\text{pk}} - \mathbf{A}_{\text{pp}} \cdot \mathbf{r}_{\text{com}} - \lfloor Q/2 \rfloor \bar{s}^{(0)} \right\| \leq \beta_{\text{com}},$$

$$\left\| \mathbf{c}_{\text{pk}} - \mathbf{A}_{\text{pk}} \cdot \bar{s}^{(0)} \right\| \leq \beta_{\text{pk}}$$

where β_{com} and β_{pk} are length-bounds on $\mathbf{e}_{\text{com}} \in \mathbb{Z}_Q^{e_{\text{com}}}$ and $\mathbf{e}_{\text{pk}} \in \mathbb{Z}_Q^e$ respectively.

- **(GSW Public Keys)** For $\tau \in \{1, 2\}$

$$\left\| (\mathbf{A}_{GSW}^{(\tau)})^\top \cdot \begin{bmatrix} \tilde{\mathbf{s}}^{(\tau)} \\ -1 \end{bmatrix} \right\| \leq \beta$$

where β is a length bound on error vector $\mathbf{e}^{(\tau)} \in \mathcal{R}_Q^2$.

- **(Bootstrapping)** For $\tau \in \{1, 2\}, i = 0, \dots, e-1, j = 0, \dots, 2\ell-1$,

$$\| \mathbf{A}_{GSW}^{(\tau)} \cdot \mathbf{r}_{i,j}^{(\tau)} + \tilde{\mathbf{s}}_i^{(\tau-1)} \cdot \mathbf{G}_j - \text{bsk}[\tau, i]_j \| \leq \beta'$$

where the index j in $\mathbf{G}_j, \text{bsk}[\tau, i]_j$ denotes the j -th column and β' is a bound on the columns of \mathbf{E}'_i .

- **(Key-Switching)** For $\tau \in \{1, 2\}, j = 0, \dots, \tilde{\ell}-1$,

$$\| (\mathbf{A}_j^{(\tau)}) \cdot \tilde{\mathbf{s}}^{(\tau)} + \tilde{\mathbf{s}}^{(\tau)} \cdot [Q/\tilde{B}^j] - \text{ksk}[\tau, j] \| \leq \beta''$$

where $\mathbf{A}_j^{(\tau)}$ is the matrix whose i -th row is $\mathbf{a}_{i,j}^{(\tau)}$, $\text{ksk}[\tau, j] \in \mathbb{Z}_Q^d$ has i -th entry $\text{ksk}[\tau, i, j]$ and β'' is a bound on the vector whose i -th component is $e_{i,j}^{(\tau)}$.

Following [LNP22], the size of the non-interactive proof is around

$$(n + \tilde{\ell} + 2 \cdot \lceil 256/d \rceil + \lambda + 2) \cdot d \log(Q) + \log(2\kappa + 1) \cdot d + (m_1 + v_e)d \cdot \log(2^{2.57} \mathfrak{s}_1) \\ + m_2 \cdot d \cdot \log(2^{2.57} \mathfrak{s}_2) + 256 \cdot \log(2^{2.57} \mathfrak{s}^{(e)})$$

where

- n is the height of a M -SIS matrix
- d is a ring dimension
- $\tilde{\ell} = 0$ because the entire witness is small
- Q is a modulus
- λ is a repetition rate so that $Q^{-\lambda}$ is negligible
- κ is a bound on the challenge coefficients e.g. $\kappa = 2$
- m_1 is the module dimension of $(\mathbf{r}_{\text{com}}, \tilde{\mathbf{s}}^{(0)}, \tilde{\mathbf{s}}^{(1)}, \tilde{\mathbf{s}}^{(2)}, \tilde{\mathbf{s}}^{(0)}, \tilde{\mathbf{s}}^{(1)}, \{\mathbf{r}_{i,j}^{(\tau)}\})$
- v_e is the module dimension of \mathbf{x}' where \mathbf{x}' consists of binary representation of norms associated to the exact bound statements. Specifically, \mathbf{x}' consists of one ring element per exact bound amounting to $v_e = (3 + 2 + 4e\ell + 2\tilde{\ell})$ ring elements.
- M -SIS $_{\mathcal{R}_q, n, m_1 + v_e + m_2, \beta_{\text{MSIS}}}$ is hard for

$$\beta_{\text{MSIS}} = 8 \cdot 59 \cdot \sqrt{2d(m_1 + v_e)\mathfrak{s}_1^2 + 2dm_2\mathfrak{s}_2^2}.$$

- \mathfrak{s}_1 is a small multiple²² of $\left\| (\mathbf{r}_{\text{com}}, \tilde{\mathbf{s}}^{(0)}, \tilde{\mathbf{s}}^{(1)}, \tilde{\mathbf{s}}^{(2)}\tilde{\mathbf{s}}^{(1)}, \tilde{\mathbf{s}}^{(2)}, \{\mathbf{r}_{i,j}^{(\tau)}\}, \mathbf{x}') \right\|$, i.e.

$$\mathfrak{s}_1 = \gamma_1 \cdot 59 \sqrt{\beta_{\text{com}}^2 d \lceil e_{\text{com}}/d \rceil + d(m_1 - 1 + v_e)}$$

²² The constant γ_1 , as well as the constants γ_2 and $\gamma^{(e)}$ for \mathfrak{s}_2 and $\mathfrak{s}^{(e)}$ appear as part of the repetition rate of the proof.

- \mathfrak{s}_2 is a part of the M -SIS bound (e.g. $\mathfrak{s}_2 = \gamma_2 \cdot 59 \cdot \sqrt{dm_2}$)
- $\mathfrak{s}^{(e)}$ is a small multiple of

$$\frac{\sqrt{(m_1 - 1) \cdot d + 2\beta_{\text{com}}^2 + \beta_{\text{pk}}^2 + 2\beta^2 + 2e \cdot \ell \cdot (\beta')^2 + 2\tilde{\ell}(\beta'')^2 + 2\log(\beta_{\text{com}})}}{\log(\beta_{\text{pk}}) + 2\log(\beta) + 2e\ell\log(\beta') + 2\tilde{\ell}\log(\beta'')}$$

i.e. $\gamma^{(e)}\sqrt{337}$ times the above.

- The M -LWE $_{\mathcal{R}_q, m_2 - n - \bar{\ell} - \lambda - 4, n + \bar{\ell} + \lambda + 4, \mathbb{Z}_3, \mathbb{Z}_3}$ problem is hard.

To offer some concrete estimates, we use our OPRF implementation parameters so that $d = 512, e = 656, Q \approx 2^{32}, \lambda = 4$. In actual fact, the structured LWE assumption used by the parameters is a M -LWE assumption in rank 2. As mentioned before, we effectively set $d = 1,024$ in our analysis. We also have a standard deviation of $\sigma \approx 2^{17}$ for the plain LWE error distributions, $\bar{\sigma} \approx 174$ for the M -LWE (or R -LWE as we are viewing it) error distributions and $\sigma_{\text{com}} = 2$. This implies that we should take $\beta_{\text{com}} = 2 \cdot \sqrt{e_{\text{com}}} = 2 \cdot \sqrt{2048}$, $\beta_{\text{pk}} = 2^{17} \cdot \sqrt{e \log Q} = 2^{44} \cdot \sqrt{656 \cdot 32}$, $\beta = \beta' = 174 \cdot \sqrt{2d} = 174 \cdot \sqrt{2048}$, $\beta'' = 2^{17} \cdot \sqrt{d} = 2^{17} \cdot \sqrt{1024}$. We also have $\ell = 2, \bar{\ell} = 4$. Recall that for our OPRF, the variables $\bar{s}_1, \bar{s}_2, \tilde{s}^{(1)}$ and $\{\mathbf{r}_{i,j}^{(\tau)} : \tau = 2\}$ are not present in the OPRF setting so we update the proof parameters accordingly.

For the M -SIS assumption, we aim for $\delta = 1.0045$ as in [LNP22]. We may then take $\gamma_1 = 40, \gamma_2 = 10, n = 3$, for the M -SIS problem noting that the parameter m_2 dimension has essentially no effect on hardness. Next, the lattice estimator tells us that a module rank of 2 is all that is required to make the M -LWE problem hard. Therefore, we require that $m_2 \geq 2 + n + \bar{\ell} + \lambda + 4 = 11$. With the constants $\gamma_1 = 40, \gamma_2 = \gamma^{(e)} = 10$, we get an approximate proof size of 25.2MB. Additionally, the expected repetition rate of the prover (that is, the expected number of times the prover runs the proof before the rejection sampling algorithms output a success) with these parameters is less than 3. Note that there is a trade-off between the prover's expected running time and the proof size. That is, we can reduce the “ γ ” parameters to get better proof sizes. However, the price to pay is a larger expected repetition rate. For example, taking $\gamma_1 = 10, \gamma_2 = \gamma^{(e)} = 5$ gives an expected repetition rate of less than 9 and a proof size of around 23.2MB.

Dilithium Compression. We can also use the compression techniques from the Dilithium signature scheme [LDK⁺22]. By dropping D bits for certain ring elements, we save $(D - 2.25) \cdot n \cdot d \cdot \log Q$ bits in the proof size. However, even with the overly optimistic $D = 30$, this only saves a few hundred kilobytes so the optimisation does not change the overall proof size in any meaningful way.

Verifiable OPRF. In the verifiable case, our parameters are somewhat speculative and we do not recommend implementing our scheme. However, to give an illustration of performance, we may choose parameters so that the bootstrapping depth of the server evaluation is exactly 2. For this, we may use the

PARAM_MESSAGE_8_CARRY_0 from the `tfhe-rs` library, modifying the plaintext space to $P = 251$. We also re-introduce the variables $\bar{s}_1, \bar{s}_2, \tilde{s}^{(1)}$ and $\{\mathbf{r}_{i,j}^{(\tau)} : \tau = 2\}$ and the associated bounds. The VOPRF parameters are $Q \approx 2^{64}, d = 32768, e = 1017, \sigma \approx 2^{40}, \bar{\sigma} = 4$. Additionally, the decomposition parameters are $\ell = 2$ and $\tilde{\ell} = 5$. The reason for this increase in parameters is to allow all linear operations in the homomorphic PRF evaluation to be performed without any bootstrapping. Using $(m_2, n) = (8, 1)$ and the same repetition rate parameters, the overall proof size becomes around 2763MB or 2572MB for repetition rates less than 3 and 9 respectively. Note that these proof sizes may be compressed by working over a smaller ring size as is done in using the alternative proof system below. For brevity, we choose not to explore this here.

LaBRADOR [BS22]. The recursive LaBRADOR proof system [BS22] can also handle the type of statements we require zero-knowledge proofs for here. LaBRADOR is a recursive proof system achieving a poly-logarithmic size and a very small slackness parameter of approximately 2. This slackness factor should be taken into account when choosing TFHE parameters because the size of the errors in the bootstrapping key may lead to incorrect evaluation and the leaking of secret information. That is, the extracted witness is at most twice as long as a genuine witness. We note that LaBRADOR is not zero-knowledge but as pointed out in [BS22] this can be addressed by adding a thin “zero-knowledge shim”. In what follows, we will assume that this does not change the order of the proof sizes.

In order to use LaBRADOR proofs, we will consider the error terms

$$\mathbf{e}_{\text{com}}, \mathbf{e}_{\text{pk}}, \mathbf{e}^{(\tau)}, \mathbf{E}_i^{(\tau)}, e_{i,j}^{(\tau)}$$

as part of the witness where the super-script τ either ranges over $\{1, 2\}$ or is simply dropped depending on whether we are in the verifiable or non-verifiable case. In particular, whether these terms are over \mathbb{Z}_Q or not, they will be considered as \mathcal{R}_Q elements where at most d coefficients are placed in a single \mathcal{R}_Q element. The error terms add up to

$$\lceil (e_{\text{com}} + e \cdot \log Q)/d \rceil + T_v \cdot (2 + 4e \cdot \ell + \tilde{\ell})$$

ring elements where $T_v = 2$ in the verifiable case and 1 in the non-verifiable case. Adding in the terms coming from the $\bar{s}^{(\cdot)}, \tilde{s}^{(\cdot)}, \mathbf{r}_{\text{com}}$ and $\mathbf{r}_{i,j}^{(\tau)}$, we end up with a witness consisting of

$$\lceil ((2T_v - 1) \cdot e + 2e_{\text{com}} + e \cdot \log Q)/d \rceil + T_v \cdot (3 + 8e \cdot \ell + \tilde{\ell})$$

ring elements. In order to cost LaBRADOR proofs, we also require a Euclidean norm bound on the witness. In the case of our statement, this equates to

$$\frac{\sqrt{(2T_v - 1)e + T_v d + 4T_v e \ell d + \sigma_{\text{com}}^2 \cdot (2e_{\text{com}})}}{\sqrt{\sigma^2 \cdot (e \log Q + T_v \tilde{\ell} d) + \bar{\sigma}^2 T_v \cdot (2d + 4e \ell d)}}$$

if all Gaussians with the same standard deviation are grouped together. Plugging in the aforementioned OPRF parameters, the witness consists of 21069 ring elements and the norm bound on the witness is around 2^{11} . Using the LaBRADOR proof system then yields a proof size of 41.42kB. In order to justify this number, we write our constraints/witness over $\mathbb{Z}_q[X]/(X^d + 1)$ for power-of-two $d > 64$ as constraints/witness over $\mathbb{Z}_q[X]/(X^{64} + 1)$. Next, we divide our 168464 witness ring elements ($\in \mathbb{Z}_q[X]/(X^{64} + 1)$) into 24 module elements, each of rank 7019 (padding the witness to make this splitting exact). Finally, we exhaustively search over five recursion levels to find recursion parameters (3, 3, 2, 1, 3). That is the first recursion uses $\nu = 3$, the second uses $\nu = 3$, the third $\nu = 2$ etc.

For the verifiable OPRF parameters, splitting the witness into 256 module elements yields a LaBRADOR proof size is around 69.85kB at five recursion levels. Note that all of these costs may be optimised further by tweaking the proof parameters.

Remark 6. We note that in our final constructions the preparation of public keys and accompanying zero-knowledge proofs can be a one-time cost, as they can be re-used.

Ciphertext Packing. Finally, we will estimate the cost of the LaBRADOR proofs in the case where we pack multiple LWE ciphertexts into a single RLWE ciphertext to save bandwidth. Essentially, we add $\log(d) \cdot d \cdot \tilde{\ell}$ error terms to the witness, each of which a discrete Gaussian with parameter $\sigma \approx 2^{17}$. There are no extra secret terms as the key-switching keys are associated to automorphisms that essentially shuffle the secret coefficients. With these updated values, the LaBRADOR proofs remains roughly the same size as reported above. Therefore, the additional key material does not contribute significantly to the communication cost of our constructions.

B.2 Protected Encoded Input Proofs (OPRF)

Computation of the function $F_{\text{poprf}}.\text{HEEval}$ (see Algorithm 2) requires a ciphertext encrypting \mathbf{y} that $\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} \equiv \mathbf{0} \pmod{q}$. Therefore, if we want the server to compute this function, the client must prove that its ciphertext is well-formed, i.e. that it encrypts a vector $\mathbf{y} \in \mathbb{Z}_p^{n_p}$ satisfying the requirement. In the following section we consider the case where \mathbf{y} is a non-unique ternary representation of a valid encoded binary input which is useful for verifiability. Recall that our suggested parameters are $p = 2, q = 3$ and $n_p = 256$. An honest server never has to multiply ciphertexts, so a full GSW ciphertext is not necessary and we can assume a TLWE ciphertext of $\mathbf{y} \in \mathbb{Z}_p^{n_p}$ is sent. As before, we let Q denote the ciphertext modulus of TFHE, e the TLWE dimension and e_{com} the commitment dimension. The plaintext modulus will be denoted P . We utilise a public key, or more specifically in our case, the commitment \mathbf{b}_{pk} to the root TFHE secret key $\mathbf{s} \in \mathbb{Z}_2^e$ (this is the key denoted $\mathbf{s}^{(0)}$ in the previous section). Then we prove that the ciphertext is a *secret key* encryption of \mathbf{y} where the key used is consistent with \mathbf{b}_{pk} . This is more efficient than proving knowledge of

the comparatively large amount of encryption randomness required to produce encryptions using the public key method. In what follows, we assume the matrix \mathbf{A} is sampled uniformly from $\mathbb{Z}_Q^{n_p \times e}$. We want a zero knowledge proof for $\mathbf{s} \in \mathbb{Z}_2^e, \mathbf{e} \in \mathbb{Z}_Q^{n_p}, \mathbf{r}_{\text{com}}, \mathbf{e}_{\text{com}} \in \mathbb{Z}_Q^{\text{com}}$ as well as some $\mathbf{y} \in \mathbb{Z}_p^{n_p}$ such that:

- $\mathbf{b}_{\text{pk}} = \mathbf{A}_{pp} \cdot \mathbf{r}_{\text{com}} + \mathbf{e}_{\text{com}} + \lfloor Q/2 \rfloor \cdot \mathbf{s} \bmod Q$
- $\mathbf{C} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{s}^\top \mathbf{A}^\top + \mathbf{e}^\top \end{bmatrix} + \begin{bmatrix} \mathbf{0}^{e \times n_p} \\ \lfloor Q/P \rfloor (\mathbf{y})^\top \end{bmatrix} \bmod Q$
- The entries of \mathbf{y} are all in \mathbb{Z}_p
- $\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} \equiv 0 \bmod q$
- $\mathbf{s} \in \mathbb{Z}_2^e$
- $\|\mathbf{r}_{\text{com}}\|, \|\mathbf{e}_{\text{com}}\| \leq \beta_{\text{com}}$
- $\|\mathbf{e}\| \leq \beta_{\text{ct}}$

Note that $\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \in \mathbb{Z}_q^{(n_q - n) \times n_p}$. If we assume that Q is big enough so that $\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y}$ does not wrap-around modulo Q , then we may replace the final equation by $\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} = q\mathbf{v}$ over the Integers where $\|q\mathbf{v}\|_\infty \leq \|\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}}\|_1 \ll Q$. This implies that the same equation holds modulo Q i.e. $\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} = q\mathbf{v} \bmod Q$ where $\|\mathbf{v}\|_\infty \leq \frac{\|\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}}\|_1}{q} =: B_v \ll Q$.

Recall that $(p, q) = (2, 3)$. Denoting \mathbf{c} as the last row of \mathbf{C} , we want a zero-knowledge proof of $\mathbf{s} \in \mathbb{Z}_2^e, \mathbf{y} \in \mathbb{Z}_p^{n_p}$ and $\mathbf{r}_{\text{com}} \in \mathbb{Z}_Q^{\text{com}}$ such that

- \mathbf{y} and \mathbf{s} have binary entries (these statements are split into a quadratic relation and a bound check)
- $\|\mathbf{r}_{\text{com}}\| \leq \beta_{\text{com}}$
- $\|\mathbf{b}_{\text{pk}} - \mathbf{A}_{pp} \cdot \mathbf{r}_{\text{com}} - \lfloor Q/2 \rfloor \cdot \mathbf{s}\| \leq \beta_{\text{com}}$
- $\|\mathbf{c} - \mathbf{A} \cdot \mathbf{s} - \lfloor Q/P \rfloor \mathbf{y}\| \leq \beta_{\text{ct}}$
- $\|\mathbf{v}\|_\infty := \|q^{-1} \cdot \mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y}\|_\infty \leq B_v$ (approximately)

Note that only $\mathbf{r}_{\text{pk}}, \mathbf{s}$ and \mathbf{y} are part of the witness in the above. In particular, \mathbf{v} is not part of the witness. When expressed as a proof of the form captured by [LNP22], the cost of running a proof is around

$$(n + \bar{\ell} + 2 \cdot \lceil 256/d \rceil + \lambda + 2) \cdot d \cdot \log(Q) + \log(2\kappa + 1) \cdot d + (m_1 + v_e) \cdot d \cdot \log(2^{2.57} \mathfrak{s}_1) \\ + m_2 \cdot d \cdot \log(2^{2.57} \mathfrak{s}_2) + 256 \cdot \left(\log(2^{2.57} \mathfrak{s}^{(e)}) + \log(2^{2.57} \mathfrak{s}^{(d)}) \right)$$

where

- n is the height of a M -SIS matrix
- $\bar{\ell} = 0$ as the entire witness is short
- $d = 128$ is a ring dimension
- Q is a modulus
- λ is a repetition rate so that $Q^{-\lambda}$ is negligible
- κ is a bound on the challenge coefficients e.g. $\kappa = 2$
- m_1 is the module dimension required to store $(\mathbf{r}_{\text{com}}, \mathbf{s}, \mathbf{y})$ i.e.

$$m_1 = \lceil e_{\text{com}}/128 \rceil + \lceil e/128 \rceil + \lceil n_p/128 \rceil$$

- v_e is the module dimension of \mathbf{x}' where \mathbf{x}' is the concatenation of binary representation of the witness norms in the exact norm proofs i.e.

$$v_e = 2\lceil \log \beta_{\text{com}}/128 \rceil + \lceil \beta_{\text{ct}}/128 \rceil = 3$$

- $M\text{-SIS}_{\mathcal{R}_Q, n, m_1 + v_e + m_2, \beta}$ is hard for $\beta = 8 \cdot 59 \cdot \sqrt{2d(m_1 + v_e)\mathfrak{s}_1^2 + 2dm_2\mathfrak{s}_2^2}$.
- \mathfrak{s}_1 is a small multiple²³ of $\|(\mathbf{r}_{\text{pk}}, \mathbf{s}, \mathbf{y}, \mathbf{x}')\|$ where \mathbf{x}' is a binary module element of rank v_e i.e.

$$\mathfrak{s}_1 = \gamma_1 \cdot 59 \sqrt{\beta_{\text{com}}^2 e_{\text{com}} + d(m_1 - \lceil e_{\text{com}}/128 \rceil + v_e)}$$

- \mathfrak{s}_2 is a part of the $M\text{-SIS}$ bound, e.g. $\mathfrak{s}_2 = \gamma_2 \cdot 59 \cdot \sqrt{dm_2}$
- $\mathfrak{s}^{(e)} = \gamma^{(e)} \cdot \sqrt{337} \cdot \sqrt{e + n_p + 2\beta_{\text{com}}^2 + \beta_{\text{ct}}^2 + 2\log(\beta_{\text{com}}) + \log(\beta_{\text{ct}})}$.
- $\mathfrak{s}^{(d)}$ is a small multiple of B_v i.e. $\gamma^{(d)} \cdot \sqrt{337} \cdot B_v$
- The $M\text{-LWE}_{\mathcal{R}_Q, m_2 - n - \bar{\ell} - \lambda - 4, n + \bar{\ell} + \lambda + 4, \mathbb{Z}_3, \mathbb{Z}_3}$ problem is hard.

For the implementation parameter set, we take $Q = 2^{32}$, $\beta_{\text{com}} = 2 \cdot \sqrt{2048}$, $\beta_{\text{ct}} \approx 2^{17} \cdot \sqrt{n_p} = 2^{17} \cdot \sqrt{256}$. We also take $B_v = 128$. To choose parameters so that the $M\text{-SIS}$ problem is hard, we may take $n = 11$ for the moment. Note that the $M\text{-SIS}$ problem does not depend heavily on the number of columns so there is no resulting restriction on m_2 aside from its appearance in the bound term. However, when analysing the $M\text{-LWE}$ problem using the lattice estimator, we find that we must take $m_2 \geq 11 + n + \bar{\ell} + \lambda + 4 = 30$ where $\lambda = 4$. Taking the most efficient choice of $m_2 = 30$, it can be verified that the $M\text{-SIS}$ problem is indeed hard when aiming for a Hermite root factor of $\delta = 1.0045$ as in [LNP22] with $\gamma_1 = 40$, $\gamma_2 = \gamma^{(e)} = \gamma^{(d)} = 10$. This gives an overall proof size of 23.8kB and an expected repetition rate of less than 3. We stress that there is a trade-off between the number of repetitions by the prover and proof-size, so these estimates only give a rough idea of the communication cost. For example, taking $\gamma_1 = 10$, $\gamma_2 = \gamma^{(e)} = \gamma^{(d)} = 5$ allows us to pick $n = 10$, $m_2 = 29$ leading to an expected repetition rate of less than 9 and a 22.1kB proof.

Dilithium Compression. We can compress the proof further by cutting $D > 0$ lower bits of certain full-sized module elements and $\log(\gamma)$ bits of smaller module elements at the expense of introducing some small hint data. This idea was introduced by the Dilithium signature scheme [LDK⁺22]. When applying Dilithium compression, the $M\text{-SIS}$ assumption changes to $M\text{-SIS}_{\mathcal{R}_Q, n, m_1 + v_e + m_2, \beta'}$ where $\beta' = 4 \cdot 59 \sqrt{\beta_1^2 + \beta_2^2}$ for $\beta_1 := 2\mathfrak{s}_1 \sqrt{2(m_1 + v_e)d}$ and $\beta_2 := 2\mathfrak{s}_2 \sqrt{2m_2d} + 2^D \cdot 59\sqrt{nd} + \gamma\sqrt{nd}$. In order to pick new parameters, we stick to the same value of n , choose γ to be the largest such that the $M\text{-SIS}$ problem remains hard for $D = 0$ and $(Q - 1)/\gamma \in \mathbb{Z}$, then pick the largest D such that the $M\text{-SIS}$ problem is hard and $2^{D-1}\kappa d < \gamma$ [LNP22]. For this, we set $Q = 2^{32} - 315 \approx 2^{32}$, $\gamma = 947 \cdot 281 \cdot 5$ and $D = 12$. The proof size decreases by $(D - 2.25) \cdot nd$ bits which equates to 1.7kB. The compression parameters for repetition rate 9 can be chosen as $D = 11$, $\gamma = 957 \cdot 281$ leading to a reduction of 1.4kB. Taking all of this into account, we have a final proof size of around 22.1kB at an expected repetition rate of less than 3 and 26.7kB at a rate of less than 9.

²³ Again, the constants are related to the repetition rate of the proof.

Amortisation Note that the above is concerned only with the case where the client submits a single encrypted PRF input. However, in the interest of amortising the cost of zero-knowledge proofs, we can consider the case where the client wishes to send L ciphertexts at once. The main changes to the proof parameters are that $m_1 = \lceil e_{\text{com}}/128 \rceil + \lceil e/128 \rceil + L \cdot \lceil n_p/128 \rceil$, $v_e = 2 + L$, $\mathbf{s}^{(e)} = \gamma^{(e)} \cdot \sqrt{337} \cdot \sqrt{e + L \cdot n_p + 2\beta_{\text{com}}^2 + L \cdot \beta_{\text{ct}}^2 + 2\log(\beta_{\text{com}}) + L \cdot \log(\beta_{\text{ct}})}$, $s^{(d)} = \gamma^{(d)} \cdot \sqrt{337} \cdot \sqrt{L \cdot B_v^2}$. Other than that, the description of the proof is the same. As an example we can take $L = 64$. Before Dilithium compression, the proof size is around 93.1kB or 85.1kB for the repetition rate 3 and 9 parameters respectively. Note that these estimates are computed using securely chosen $(m_2, n) = (31, 12)$ and $(30, 11)$ for the repetition rate 3 and 9 parameters respectively. Furthermore, Dilithium compression parameters can be chosen as $(D, \gamma) = (13, 947 \cdot 281 \cdot 6)$ and $(12, 947 \cdot 281 \cdot 2)$ to result in a decrease of 2.0kB and 1.7kB respectively. This results in an amortised client online proof size of 1.42kB or 1.3kB per query.

LaBRADOR Proofs. We also tried to use the LaBRADOR proof system for the statements in this section. However, when attempting this, it turns out that the size of the statements for e.g. $L = 64$ render recursion ineffective and the proof sizes are larger than those reported above.

B.3 Non-Unique Encoding Proofs (VOPRF)

We now discuss the costs of the protected encoded input proofs when the encoding is not unique. This is used in our VOPRF candidate. In other words, we now discuss the case where a client ciphertext encrypts an input $\tilde{\mathbf{y}}$ that is equivalent modulo p to an $\mathbf{y} \in \mathbb{Z}_p^{n_p}$ such that $\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} \equiv \mathbf{0} \pmod{q}$. Reusing the notation from Section B.2, we want a zero knowledge proof for $\mathbf{s} \in \mathbb{Z}_2^e$, $\mathbf{e} \in \mathbb{Z}_Q^{n_p}$, $\mathbf{r}_{\text{com}}, \mathbf{e}_{\text{com}} \in \mathbb{Z}_Q^{e_{\text{com}}}$ as well as some $\mathbf{y} \in \mathbb{Z}_p^{n_p}$ and $\tilde{\mathbf{y}} \in \mathbb{Z}_q^{n_p}$ such that:

- $\mathbf{b}_{\text{pk}} = \mathbf{A}_{pp} \cdot \mathbf{r}_{\text{com}} + \mathbf{e}_{\text{com}} + \lfloor Q/2 \rfloor \mathbf{s} \pmod{Q}$
- $\mathbf{C} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{s}^\top \mathbf{A}^\top + \mathbf{e}^\top \end{bmatrix} + \begin{bmatrix} \mathbf{0}^{e \times n_p} \\ \lfloor Q/P \rfloor (\tilde{\mathbf{y}}^{(2)})^\top \end{bmatrix} \pmod{Q}$
- $\mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y} \equiv \mathbf{0} \pmod{q}$
- The entries of \mathbf{y} are all in \mathbb{Z}_p and the entries of $\tilde{\mathbf{y}}$ are in \mathbb{Z}_q
- $\tilde{\mathbf{y}}$ and \mathbf{y} encode the same value modulo p
- $\mathbf{s} \in \mathbb{Z}_2^e$
- $\|\mathbf{r}_{\text{com}}\|, \|\mathbf{e}_{\text{com}}\| \leq \beta_{\text{com}}$
- $\|\mathbf{e}\| \leq \beta_{\text{ct}}$

Once again, we consider $(p, q) = (2, 3)$ concretely. Suppose we restrict \mathbf{y} to be a binary, i.e. \mathbb{Z}_p , vector. Suppose also that we restrict $\tilde{\mathbf{y}}$ to be ternary i.e. a \mathbb{Z}_q vector (assuming integers modulo $q = 3$ are represented in $\{-1, 0, 1\}$). Then one way to provide a non-unique encoding is to require the existence of a short vector \mathbf{w} , such that $\mathbf{y} = \tilde{\mathbf{y}} + 2\mathbf{w}$ over the Integers. Note that on the assumption that \mathbf{y} and $\tilde{\mathbf{y}}$ are binary and ternary respectively, we can prove this by showing that

$\mathbf{y} = \bar{\mathbf{y}} + 2\mathbf{w} \bmod Q$ and providing a very loose range proof on \mathbf{w} . Furthermore, in order to show that $\bar{\mathbf{y}}$ is ternary, we may write it as the difference of two binary vectors e.g. $\bar{\mathbf{y}} = \bar{\mathbf{y}}_\ell - \bar{\mathbf{y}}_r$ and show that $\bar{\mathbf{y}}_\ell$ and $\bar{\mathbf{y}}_r$ are both binary. Taking all of this into account and using the strategy of the previous section, we may prove knowledge of $\mathbf{r}_{\text{com}} \in \mathbb{Z}_Q^{e_{\text{com}}}$, $\mathbf{s} \in \mathbb{Z}_Q^e$, $\mathbf{y}, \bar{\mathbf{y}}_\ell, \bar{\mathbf{y}}_r, \mathbf{w} \in \mathbb{Z}_Q^{n_p}$ such that:

- $\mathbf{y}, \bar{\mathbf{y}}_\ell, \bar{\mathbf{y}}_r$ and \mathbf{s} have binary entries (which are treated as a quadratic relation plus a bound check)
- $\|\mathbf{r}_{\text{com}}\| \leq \beta_{\text{com}}$
- $\|\mathbf{b}_{\text{pk}} - \mathbf{A}_{pp} \cdot \mathbf{r}_{\text{com}} - \lfloor Q/2 \rfloor \mathbf{s}\| \leq \beta_{\text{com}}$
- $\|\mathbf{c} - \mathbf{A} \cdot \mathbf{s} - \lfloor Q/P \rfloor (\bar{\mathbf{y}}_\ell - \bar{\mathbf{y}}_r)\| \leq \beta_{\text{ct}}$
- $\|\mathbf{v}\|_\infty := \|3^{-1} \cdot \mathbf{H}_{\text{inp}} \cdot \mathbf{G}_{\text{gadget}} \cdot \mathbf{y}\|_\infty \leq B_v$ (approximately)
- $\mathbf{y} = (\bar{\mathbf{y}}_\ell - \bar{\mathbf{y}}_r) + 2\mathbf{w} \bmod Q$
- $\|\mathbf{w}\|_\infty \leq 1$ (approximately)

Once again, the cost of running a proof using [LNP22] is around

$$(n + \bar{\ell} + 2 \cdot \lceil 256/d \rceil + \lambda + 2) \cdot d \cdot \log(Q) + \log(2\kappa + 1) \cdot d + (m_1 + v_e) \cdot d \cdot \log(2^{2.57} \mathbf{s}_1) \\ + m_2 \cdot d \cdot \log(2^{2.57} \mathbf{s}_2) + 256 \cdot \left(\log(2^{2.57} \mathbf{s}^{(e)}) + \log(2^{2.57} \mathbf{s}^{(d)}) \right)$$

where

- n is the height of a M -SIS matrix
- $\bar{\ell} = 0$ as the entire witness is short
- $d = 128$ is a ring dimension
- Q is a modulus
- λ is a repetition rate so that $Q^{-\lambda}$ is negligible
- κ is a bound on the challenge coefficients e.g. $\kappa = 2$
- m_1 is the module dimension required to store $(\mathbf{r}_{\text{com}}, \mathbf{s}, \mathbf{y}, \bar{\mathbf{y}}_\ell, \bar{\mathbf{y}}_r, \mathbf{w})$ i.e. $m_1 = \lceil e_{\text{com}}/d \rceil + \lceil e/d \rceil + 4 \cdot \lceil n_p/d \rceil$
- v_e is the module dimension of \mathbf{x}' where \mathbf{x}' contains the binary representation of the vectors in the exact norm proofs (i.e. $v_e = 2 \lceil \log(\beta_{\text{com}})/128 \rceil + \lceil \log(\beta_{\text{ct}})/128 \rceil = 3$)
- $M\text{-SIS}_{\mathcal{R}_Q, n, m_1 + v_e + m_2, \beta}$ is hard for $\beta = 8 \cdot 59 \cdot \sqrt{2d(m_1 + v_e)\mathbf{s}_1^2 + 2dm_2\mathbf{s}_2^2}$.
- $\mathbf{s}_1 = \gamma_1 \cdot 59 \sqrt{\beta_{\text{com}}^2 + d(m_1 - \lceil e_{\text{com}}/128 \rceil + v_e)}$
- $\mathbf{s}_2 = \gamma_2 \cdot 59 \cdot \sqrt{dm_2}$
- $\mathbf{s}^{(d)} = \gamma^{(d)} \cdot \sqrt{337} \cdot \sqrt{B_v^2 + n_p}$
- $\mathbf{s}^{(e)} = \gamma^{(e)} \cdot \sqrt{337} \cdot \sqrt{e + 3n_p + 2\beta_{\text{com}}^2 + \beta_{\text{ct}}^2 + 2\log(\beta_{\text{com}}) + \log(\beta_{\text{ct}})}$
- The $M\text{-LWE}_{\mathcal{R}_Q, m_2 - n - \bar{\ell} - \lambda - 4, n + \bar{\ell} + \lambda + 4, \mathbb{Z}_3, \mathbb{Z}_3}$ problem is hard.

Pending further cryptanalysis, we do not recommend implementing our VO-PRF. However, we may still provide an account of performance by using the `PARAM_MESSAGE_8_CARRY_0` parameters assuming a plaintext space of $P = 251$. The main parameters of interest are $Q \approx 2^{64}$, $e = 1017$, $\sigma \approx 2^{40}$. The commitment parameters σ_{com} and e_{com} are unchanged. For the M -LWE problem to be hard, we now require a module rank of 20. Before Dilithium compression, the proofs for repetition rate 3 and 9 use $(m_2, n) = (32, 6)$ and $(31, 5)$ respectively resulting in proof sizes of 40.8kB and 37.4kB. We will use $Q = 2^{64} - 179$ along with

Dilithium compression parameters $(D, \gamma) = (12, 4051 \cdot 41 \cdot 4)$ and $(11, 4051 \cdot 41 \cdot 2)$ respectively, we obtain proof sizes of 40.8kB or 37.6kB.

For the amortised case where we send $L = 64$ ciphertexts at once. The pre-compressed proofs have sizes of 238.2kB and 217.1kB using $(m_2, n) = (33, 7)$ and $(32, 6)$ respectively. It can be shown that we may use Dilithium compression parameters $(D, \gamma) = (17, 6199 \cdot 101 \cdot 41)$ and $(14, 6199 \cdot 101 \cdot 4)$. Overall this leads to an amortised proof size of 3.7kB or 3.4kB per query.

C Verifiability Construction

In this section, we present our verifiable construction in Figure 5. Note that for simplicity, we have presented a VOPRF, i.e. it is not partial, by setting the tag $\mathbf{t} = \perp$.

C.1 POPRIV2 from Verifiability

We show that our verifiability property implies POPRIV2 under some additional assumptions.

Theorem 3. *Our VOPRF construction satisfies POPRIV2 if*

- F_{vpoprf} satisfies POPRIV1
- F_{vpoprf} is verifiable
- NIZKAoK is sound.

Proof. Let G_0 and G_1 denote the $\text{POPRIV2}_{\mathcal{F}, H}^{A, b}$ game for $b = 0$ and $b = 1$ respectively. We consider the possibility of two events, E_0 and E_1 , that occur in a query to the Finalise oracle:

1. ct_0 is not the output of $F_{\text{vpoprf}}.\text{BlindEval}(\mathbf{A}, \mathbf{t}, \text{req}_0)$ but $F_{\text{vpoprf}}.\text{Verify}(\mathbf{z}_0, \mathbf{z}_0^*, \rho_0) = 1$.
2. ct_1 is not the output of $F_{\text{vpoprf}}.\text{BlindEval}(\mathbf{A}, \mathbf{t}, \text{req}_1)$ but $F_{\text{vpoprf}}.\text{Verify}(\mathbf{z}_1, \mathbf{z}_1^*, \rho_1) = 1$.

We claim that the probability of these events occurring is bound by an adversary against the verifiability property of \mathcal{F} and soundness of NIZKAoK.

To see this, we consider the game G'_0 that is defined as G_0 but with the change that the Request oracle requires the adversary to submit its secret key \mathbf{A} . We further require an additional check when the oracle executes $\mathcal{F}.\text{Request}^{\text{RO}}$, on line 2 we verify that $\hat{\mathbf{z}} = \mathcal{F}.\text{Eval}(\mathbf{A}, \mathbf{t}, \hat{\mathbf{x}})$. The success probability of an adversary between these games is bound by the soundness of NIZKAoK. This follows from the fact that the extra check is ensuring that the relation of the argument of knowledge holds using the witness. Any adversary that wins against G_0 but not against G'_0 has created a proof $\hat{\pi}$ such that $\hat{\mathbf{z}} \neq F_{\text{vpoprf}}.\text{Eval}(\mathbf{A}, \mathbf{t}, \hat{\mathbf{x}})$. A similar argument also bounds the difference in winning probability between G_1 and an analogous game G'_1 . Since we assume that NIZKAoK_{Rt} is sound, we have that $G_0 \approx_c G'_0$ and similarly, $G_1 \approx_c G'_1$.

We then next observe that if the event E_0 or E_1 has occurred in G'_0 or G'_1 , then we break the verifiability property of \mathcal{F} . Once more we initially consider the experiment G'_0 . We construct an adversary $\mathcal{B}_{\text{verif}}$ against G'_0 when ct_0 is not the output of $\mathcal{F}.\text{BlindEval}(\mathcal{A}, \mathbf{t}, \text{req}_0)$ but it is accepted by the Finalise oracle. To do so, it invokes its own verifiability experiment and receives the public parameters pp , which it uses as pp in G'_0 . \mathcal{A} can make queries of two types, Request and Finalise, which $\mathcal{B}_{\text{verif}}$ handles by simply forwarding to its own oracles in the verifiability experiment. Then, $\mathcal{B}_{\text{verif}}$ must guess which of the queries contains the ‘winning’ response. It does so with probability $1/q_R$ for polynomial q_R which bounds the number of Request queries. It receives the checkpoints $\mathbf{x}_i, i \in \mathbb{Z}_n$ and forwards these to its verifiability experiment. It receives back a permuted set of req_i which it sends to \mathcal{A} in response to its Request query. It waits for the corresponding call to the Finalise oracle, at which point it forwards the set rep_i to the experiment G'_0 . It wins if \mathcal{A} was able to provide rep that passed verification but was not the honest output of BlindEval . By assumption, this probability is negligible, and therefore we have $\Pr[G'_0 \Rightarrow 1] \leq \text{negl}(\lambda)$. An analogous argument holds for when the bit $b = 1$ and thus we have shown that probability of event E_0 or E_1 in G_0 or G_1 is negligible.

Then, since we have shown G_i always correctly computes responses rep (with all but negligible probability), we can apply the same argument as we have for POPRIV1. Thus we conclude that the transcript observed by an adversary for POPRIV2 is independent of the challenge bit b , and hence the advantage of the adversary against POPRIV2 is negligible.

By considering the sequence of games, we have shown that $G_0 \approx_c G_1$ and thus we obtain the theorem statement. \square

D PRF Sage Example

```
# -*- coding: utf-8 -*-
from sage.all import set_random_seed, GF, vector, matrix, random_matrix, codes

class WeakPRF:
    """
    Based on Construction 3.1 of:

    - Boneh, D., Ishai, Y., Passelègue, A., Sahai, A., & Wu, D. J. (2018).
      Exploring crypto dark matter: new simple PRF candidates and their
      applications. Cryptology ePrint Archive, Report 2018/1218.
      https://eprint.iacr.org/2018/1218
    """

    def __init__(self, m_p=256, n_p=256, m_bound=128, p=2, q=3, t=None, seed=None):
        # p.40, optimistic, λ=128
        self.m_p, self.n_p = m_p, n_p
        if seed is not None or t is not None:
            set_random_seed(hash((t, seed)))
        self.A = random_matrix(GF(p), m_p, n_p + 1)
        self.q = q

        if m_bound == 1:
            self.Gout = matrix(GF(q), 1, m_p, [1] * m_p)
        else:
```

```

        i = 2
        for i in range(2, m_p // 2):
            C = codes.BCHCode(GF(q), m_p, i)
            if C.dimension() <= m_bound:
                self.Gout = C.generator_matrix()
                break
        else:
            raise RuntimeError

    def __call__(self, x):
        y = self.A * vector(list(x) + [1])
        y = y.lift_centered()
        z = self.Gout * y
        return z

class PRF:
    """
    Based on Construction 7.9 of:

    - Boneh, D., Ishai, Y., Passelègue, A., Sahai, A., & Wu, D. J. (2018).
    Exploring crypto dark matter: new simple PRF candidates and their
    applications. Cryptology ePrint Archive, Report 2018/1218.
    https://eprint.iacr.org/2018/1218
    """

    def __init__(self, m_p=256, n_p=256, n=128, m_bound=128, p=2, q=3, t=None, seed=None):
        if p != 2 or q != 3:
            raise NotImplementedError
        self.F = WeakPRF(m_p=m_p, n_p=n_p, m_bound=m_bound, p=p, q=q, t=t, seed=seed)
        self.n = n
        self.Ginp = identity_matrix(GF(q), n).stack(random_matrix(GF(q), (n_p-n)//2, n))

    def __call__(self, x):
        x = x.lift().change_ring(GF(self.F.q))
        y = self.Ginp * x
        z = []
        for y_ in y[self.n].lift():
            z.append(y_)
        for y_ in y[:self.n].lift():
            z.append(y_ % 2)
            z.append(y_ // 2)
        z = vector(GF(2), self.F.n_p, z)
        return self.F(z)

```

E Size Estimates for LaBRADOR

The script is also [attached](#).

```

# -*- coding: utf-8 -*-
"""
LaBRADOR Pari/GP Code in Sage.
"""

from sage.all import (
    log,
    ceil,
    sqrt,
    vector,
    round,
    floor,
    exp,
    ZZ,
    RR,
    pi,
    exp,
    cached_function,
    cached_method,

```



```

    Infinity,
    get_verbose,
)

def gaussian_entropy(sigma):
    if sigma >= 4:
        a = floor(sigma / 2)
        sigma /= a
    else:
        a = 1

    d = 1 / (2 * sigma**2)
    n = sum(exp(-(i**2) * d) for i in range(-ceil(15 * sigma), 0))
    n = 2 * n + 1
    logn = log(n)
    e = 0
    for i in range(-ceil(15 * sigma), 0):
        f = exp(-(i**2) * d)
        e += f * (log(f) - logn)
    e = (-2 * e + logn) / (n * log(2))

    return float(e + log(a, 2))

def deltaf(b):
    """
    Compute root Hermite factor for block size ``b``.
    """
    small = (
        (2, 1.02190),
        (5, 1.01862),
        (10, 1.01616),
        (15, 1.01485),
        (20, 1.01420),
        (25, 1.01342),
        (28, 1.01331),
        (40, 1.01295),
    )

    if b <= 2:
        return 1.0219
    elif b < 40:
        for i in range(1, len(small)):
            if small[i][0] > b:
                return small[i - 1][1]
    elif b == 40:
        return small[-1][1]
    else:
        return float(b / (2 * pi * exp(1)) * (pi * b) ** (1.0 / b)) ** (
            1.0 / (2 * b - 2.0)
        )

def block_sizef(delta):
    b = 40
    while deltaf(2 * b) > delta:
        b *= 2
    while deltaf(b + 10) > delta:
        b += 10
    while deltaf(b) >= delta:
        b += 1

    return b

def adps16(block_size):
    return block_size * log(sqrt(3.0 / 2.0), 2.0)

```

```

default_costf = adps16

def sis_required_block_size(
    n,
    q,
    b,
    m=None,
    step_size=10,
    start=40,
):
    """
    Hardness of SIS in the  $\ell_2$  norm on  $A \in \mathbb{Z}_q^{n \times m}$ 

    :param n: SIS rank
    :param q: Integer modulus
    :param b: Target  $\ell_2$  norm of the solution.
    :param m: SIS dimension (default:  $n \div \log q$ )
    :param step_size: Search in steps of this size first (default: 10)
    :param start: Start search at this block size.

    """
    if b > q:
        raise ValueError(f"Size bound {b} > modulus {q}.")

    if m is None:
        m = ceil(2 * n * log(q, 2))

    for block_size in range(start, m + step_size, step_size):
        delta = deltax(block_size)
        d = min(sqrt(n * log(q) / log(delta)), m)
        if delta ** (d - 1) * q ** (n / d) < b:
            break

    for block_size in range(block_size - step_size, block_size + 1):
        delta = deltax(block_size)
        d = min(sqrt(n * log(q) / log(delta)), m)
        if delta ** (d - 1) * q ** (n / d) < b:
            return block_size

    return Infinity

@cached_function
def sis_hard_enough(kappa, eta, b, q):
    """
    Return `i` such that for `n = i * eta` and a sufficiently big `m`  $\beta$ SIS on  $\mathbb{Z}_q^{n \times m}$ 
    requires block size  $k$ `.

    """
    if b > q:
        raise ValueError(f"Size bound {b} > modulus {q}.")

    i = 1
    while True:
        n = i * eta
        delta = deltax(kappa - 1)
        d = sqrt(n * log(q) / log(delta))
        if delta ** (d - 1) * q ** (n / d) > b:
            return i
        i += 1

class LaBRADOR:
    def __init__(
        self,
        d: int = 64,

```

```

logq: int = 32,
tau: int = 71,
T: int = 15,
slack: float = sqrt(128 / 30.0),
max_beta: int = 0,
secpair: int = 100,
costf=default_costf,
):
    self.d = d
    self.logq = logq
    self.tau = tau
    self.T = T
    self.slack = slack
    self.max_beta = max_beta
    self.secpair = secpair
    self.costf = default_costf

    block_size = None
    for block_size in range(self.secpair, 2048, 32):
        if self.costf(block_size) >= self.secpair:
            break

    for block_size in range(block_size - 32, block_size + 1):
        if self.costf(block_size) >= self.secpair:
            self.block_size = block_size
            break

def sis_rank(self, beta):
    self.max_beta = max(self.max_beta, beta)

    # we round to a nearby value to allow for caching which improves performance
    # beta = 1.2 ** ceil(log(beta, 1.2))

    try:
        return sis_hard_enough(self.block_size, self.d, ceil(beta), 2**self.logq)
    except ValueError:
        return Infinity

def main(self, n, r, beta, nu, decompose):
    old_beta = vector(beta).norm(2).n()
    # TODO: this does not reflect secpair yet
    size = 256 * gaussian_entropy(old_beta / sqrt(2)) # JL projection
    size += ceil(128 / self.logq) * self.d * self.logq # JL proof

    sigs = [float(beta[i] / sqrt(r[i] * n * self.d)) for i in range(len(r))]
    sigz = sqrt(
        sigs[0] ** 2 * (1 + (r[0] - 1) * self.tau)
        + sum([sigs[i] ** 2 * r[i] * self.tau for i in range(1, len(r))])
    )
    sigh = float(sqrt(2 * n * self.d) * max(sigs) ** 2)

    if decompose:
        t = 2
        b = round(sqrt(sqrt(12) * sigz))
    else:
        t = 1
        b = 1

    t1 = round(self.logq / log(sqrt(12) * sigz / b, 2))
    t1 = max(2, t1)
    t1 = min(14, t1)

    b1 = ceil(2 ** (self.logq / t1))
    t2 = round(log(sqrt(12) * sigh) / log(sqrt(12) * sigz / b))
    t2 = max(1, t2)
    b2 = ceil((sqrt(12) * sigh) ** (1 / t2))

    r = sum(r)

```

```

beta = [0, 0]
beta[0] = float(sigz / float(b) * sqrt(t * n * self.d))
for i in range(16):
    kappa = i + 1
    beta[1] = float(
        sqrt(
            b1**2 / 12.0 * t1 * r * kappa * self.d
            + (b1**2 * t1 + b2**2 * t2) / 12.0 * (r**2 + r) / 2.0 * self.d
        )
    )
    new_beta = vector(beta).norm(2).n()
    if (
        self.sis_rank(
            max(
                6 * self.T * b * self.slack * new_beta,
                2 * b * self.slack * new_beta
                + 4 * self.T * self.slack * old_beta,
            )
        )
        <= kappa
    ):
        break

kappa1 = self.sis_rank(2 * self.slack * new_beta)
size += 2 * kappa1 * self.d * self.logq
# outer commitments
m = t1 * r * kappa + (t1 + t2) * (r**2 + r) / 2
mu = round(m / ceil(n / nu))
mu = max(1, mu)
n = ceil(n / nu)
m = ceil(m / mu)
n = max(n, m)
r = [t * nu, mu]

if get_verbose() >= 1:
    print("Main:")
    print("Commitments: kappa = %d; kappa1 = kappa2 = %.2f" % (kappa, kappa1))
    print("Decomposition bases: b = %d; b1 = %d; b2 = %d" % (b, b1, b2))
    print("Expansion factors: t = %d; t1 = %d; t2 = %d" % (t, t1, t2))
    print("Target relation: n = %d; r = %s; b = %s" % (n, r, b))
    print("Norm balance: %.2f%%" % ((b[1] - b[0]) / max(b[0], b[1]) * 100))

return size, n, r, beta

def tail(self, n, r, beta):
    old_beta = vector(beta).norm(2).n()
    size = 256 * gaussian_entropy(old_beta / sqrt(2)) # JL projection
    size += ceil(128 / self.logq) * self.d * self.logq # JL proof
    size += 128 # challenges

    sigs = [float(beta[i] / sqrt(r[i] * n * self.d)) for i in range(len(r))]
    sigh = float(sqrt(2 * n * self.d) * max(sigs) ** 2)

    t1 = round(self.logq / log(sqrt(12) * sum(sigs) / len(sigs), 2))
    t1 = max(2, t1)
    t1 = min(14, t1)
    b1 = ceil(2 ** (self.logq / t1))
    t2 = round(log(sqrt(12) * sigh) / log(sqrt(12) * sum(sigs) / len(sigs)))
    t2 = max(1, t2)
    b2 = ceil((sqrt(12) * sigh) ** (1 / t2))

    for i in range(16):
        kappa = i + 1
        x = sum(r)
        n2 = x * kappa * t1 + (x**2 + x) / 2 * t2
        r2 = round(n2 / n)
        r2 = max(1, r2)

```

```

sigz = sqrt(
    sigs[0] ** 2 * (1 + (r[0] - 1) * self.tau)
    + sum([sigs[i] ** 2 * r[i] for i in range(1, len(r))]) * self.tau
    + r2 * max(b1, b2) ** 2 / 12.0 * self.tau
)

beta = sigz * sqrt(max(n, ceil(n2 / r2)) * self.d)
if self.sis_rank(6 * self.T * beta) <= kappa:
    break

r = sum(r)
n = max(n, ceil(n2 / r2))

size += r2 * kappa * self.d * self.logq # outer commitments
size += 2 * r2 * self.d * gaussian_entropy(sigh) # quadratic garbage polys
size += (2 * (r - 1) + 2 * r2) * self.d * self.logq # linear garbage polys
size += n * self.d * float(gaussian_entropy(sigz)) # masked opening

if get_verbose() >= 1:
    print("Tail:")
    print("Outer Commitments: kappa = %d" % kappa)
    print("Additional multiplicity: r2 = %d" % r2)
    print("Decomposition bases: b1 = %d b2 = %d" % (b1, b2))
    print("Expansion factors: t1 = %d t2 = %d" % (t1, t2))
    print("Final relation: n = %d  $\beta$  = %s" % (n, beta))
return size

def size(self, n, r, beta, nuvec):
    """
    Size in kilobytes
    """
    s = 0
    r, beta = [r], [RR(beta)]
    for i in range(len(nuvec)):
        size, n, r, beta = self.main(n, r, beta, nuvec[i], i < len(nuvec) - 1)
        s += size

    s += self.tail(n, r, beta)

    return round(s / 2**13, 2)

@cached_method
def __call__(self, n, r, beta, base, length, verbose=True):
    def i2v(i):
        return vector(ZZ(i).digits(base, padto=length)) + vector(
            ZZ, length, [1] * length
        )

    best = self.size(n, r, beta, i2v(0)), 0

    for i in range(base**length):
        current = self.size(n, r, beta, i2v(i)), i
        if current[0] < best[0]:
            best = current
            if verbose:
                print(f"{best[0]:.2f}kB, {i2v(best[1])}")

    return best[0], i2v(best[1])

```

F Parameter Estimates

The script is also [attached](#).

```

"""
Estimating sizes.

```

```

"""

from dataclasses import dataclass
from functools import partial
from sage.all import sqrt, ceil, log, cached_function
from labrador import LaBRADOR

# Parameters

class HashableDict(dict):
    def __hash__(self):
        return hash(frozenset(self.items()))

@dataclass
class Parameters:
    e_com: int
    sigma_com: float
    log_q: int
    n_p: int
    m: int
    tfhe: HashableDict

    def __hash__(self):
        return hash((self.e_com, self.sigma_com, self.log_q, self.n_p, self.m, self.tfhe))

# https://github.com/zama-ai/tfhe-rs/blob/0.2.4/tfhe/src/shortint/parameters/mod.rs#L200
PARAM_MESSAGE_2_CARRY_0 = HashableDict(
    {
        "lwe_dimension": 656,
        "glwe_dimension": 2,
        "polynomial_size": 512,
        "lwe_modular_std_dev": 0.000034119201269311964,
        "glwe_modular_std_dev": 0.00000004053919869756513,
        "pbs_base_log": 8,
        "pbs_level": 2,
        "ks_level": 4,
        "ks_base_log": 3,
        "pfks_level": 1,
        "pfks_base_log": 15,
        "pfks_modular_std_dev": 0.00000000037411618952047216,
        "cbs_level": 0,
        "cbs_base_log": 0,
        "message_modulus": 4,
        "carry_modulus": 1,
    }
)

OPRF = Parameters(
    e_com=2048, sigma_com=2, log_q=32, n_p=256, m=128, tfhe=PARAM_MESSAGE_2_CARRY_0
)

# https://github.com/zama-ai/tfhe-rs/blob/0.2.4/tfhe/src/shortint/parameters/mod.rs#L827
PARAM_MESSAGE_8_CARRY_0 = HashableDict(
    {
        "lwe_dimension": 1017,
        "glwe_dimension": 1,
        "polynomial_size": 32768,
        "lwe_modular_std_dev": 0.0000000460803851108693,
        "glwe_modular_std_dev": 0.0000000000000000002168404344971009,
        "pbs_base_log": 15,
        "pbs_level": 2,
        "ks_level": 5,
        "ks_base_log": 4,
        "pfks_level": 2,
    }
)

```



```

if not compress_ct1:
    zeta = tau
else:
    zeta = tau + log(d, 2)

return _mb(
    (params.e_com + e * params.log_q) * lwe_keep_bits
    + tau * 2 * d * glwe_keep_bits
    + tau * 2 * e * 2 * ell * d * glwe_keep_bits
    + zeta * d * ell_ * lwe_keep_bits
)

voprpf_pk_sizemb = partial(oprf_pk_sizemb, voprpf=True)

def oprf_pk_proof_old_sizemb(params, gamma1=10, gamma2=5, gammae=5):
    """
    Well-formedness proof of the OPRF public key using [LNP22].

    :param params:
    :param gamma1:
    :param gamma2:
    :param gammae:

    """
    ell = params.tfhe["pbs_level"]
    t1 = params.tfhe["ks_level"]
    eta = 1 # we absorbed the module dimension in the ring dimension
    d = params.tfhe["polynomial_size"] * params.tfhe["glwe_dimension"]
    e = params.tfhe["lwe_dimension"]
    Q = 2**params.log_q
    lamb = 2 * 128 // params.log_q
    kap = 2
    m1 = (params.e_com + e) // d + eta + 4 * e * ell
    delta = 1.0045
    ve = 4 + 2 * e * ell + t1
    bm1 = m1 + ve
    bcom = params.sigma_com * sqrt(params.e_com)
    s1 = gamma1 * 59 * sqrt(bcom**2 * d * ceil(params.e_com / d) + d * (bm1 - 1))

    bpk = params.tfhe["lwe_modular_std_dev"] * Q * sqrt(e * params.log_q)
    b = params.tfhe["glwe_modular_std_dev"] * Q * sqrt(2 * d)
    bp = b
    bpp = params.tfhe["lwe_modular_std_dev"] * Q * sqrt(eta * d)
    inside_se = (
        (m1 - 1) * d
        + 2 * bcom**2
        + bpk**2
        + b**2
        + e * ell * bp**2
        + t1 * bpp**2
        + 2 * log(bcom, 2)
        + log(bpk, 2)
        + log(b, 2)
        + e * ell * log(bp, 2)
        + t1 * log(bpp, 2)
    )
    se = gammae * sqrt(337) * sqrt(inside_se)

    for n in range(1, 10):
        m2 = 2 + n + lamb + 4
        s2 = gamma2 * 59 * sqrt(d * m2)

        # see if n large enough for the SIS problem
        beta = 8 * 59 * sqrt(2 * d * bm1 * s1**2 + 2 * d * m2 * s2**2)
        if beta > 2 ** (2 * sqrt(n * d * params.log_q * log(delta, 2))):

```



```

        continue

lwe_estimator_dim = (m2 - n - lamb - 4) * d
print(f"Check that ternary LWE in dimension {lwe_estimator_dim} is hard.")

tot = (
    (n + 2 + lamb + 2) * d * params.log_q
    + log(2 * kap + 1, 2) * d
    + bm1 * d * log(s1 * 2**2.57, 2)
    + m2 * d * log(s2 * 2**2.57, 2)
    + 256 * log(se * 2**2.57, 2)
)

return _mb(tot)

def voprk_pk_proof_old_sizeb(params, gamma1=10, gamma2=5, gammae=5):
    """
    Well-formedness proof of the OPRF public key using [LNP22].

    :param params:
    :param gamma1:
    :param gamma2:
    :param gammae:

    """
    ell = params.tfhe["pbs_level"]
    eta = 1
    t1 = params.tfhe["ks_level"]
    d = params.tfhe["polynomial_size"] * params.tfhe["glwe_dimension"]
    e = params.tfhe["lwe_dimension"]
    ecom = params.e_com
    Q = 2**params.log_q
    lamb = 128 // params.log_q
    kap = 2
    delta = 1.0045

    for n in range(1, 10):
        m1 = (ecom + 3 * e) // d + 2 * eta + 8 * e * ell
        m2 = 3 + n + t1 + lamb + 4
        ve = 5 + 4 * e * ell + 2 * t1
        bm1 = m1 + ve
        bcom = params.sigma_com * sqrt(ecom)
        s1 = gamma1 * 59 * sqrt(bcom**2 * d * ceil(ecom / d) + d * (bm1 - 1))
        s2 = gamma2 * 59 * sqrt(d * m2)

        # see if n large enough for the sis problem
        beta = 8 * 59 * sqrt(2 * d * bm1 * s1**2 + 2 * d * m2 * s2**2)
        if beta > 2 ** (2 * sqrt(n * d * log(Q, 2) * log(delta, 2))):
            continue

        bpk = params.tfhe["lwe_modular_std_dev"] * Q * sqrt(e * params.log_q)
        b = params.tfhe["glwe_modular_std_dev"] * Q * sqrt(2 * d)
        bp = b
        bpp = params.tfhe["lwe_modular_std_dev"] * Q * sqrt(eta * d)
        inside_se = (
            (m1 - 1) * d
            + 2 * bcom**2
            + bpk**2
            + 2 * b**2
            + 2 * e * ell * bp**2
            + 2 * t1 * bpp**2
            + 2 * log(bcom, 2)
            + log(bpk, 2)
            + 2 * log(b, 2)
            + 2 * e * ell * log(bp, 2)
            + 2 * t1 * log(bpp, 2)
        )

```

```

    se = gammae * sqrt(337) * sqrt(inside_se)

    tot = (
        (n + 2 + lamb + 2) * d * params.log_q
        + log(2 * kap + 1, 2) * d
        + bm1 * d * log(s1 * 2**2.57, 2)
        + m2 * d * log(s2 * 2**2.57, 2)
        + 256 * log(se * 2**2.57, 2)
    )

    return _mb(tot)

@cached_function
def oprf_pk_proof_sizekb(params, compress_ct1=True):
    """
    Well-formedness proof of the OPRF public key using LaBRADOR.
    """
    e = params.tfhe["lwe_dimension"]
    eta = 1 # we will absorb the module rank in the ring dimension
    d = params.tfhe["polynomial_size"] * params.tfhe["glwe_dimension"]
    sig = params.tfhe["lwe_modular_std_dev"]
    sigring = params.tfhe["glwe_modular_std_dev"]
    ell = params.tfhe["pbs_level"]
    t1 = params.tfhe["ks_level"]

    if not compress_ct1:
        zeta = eta
    else:
        zeta = eta + log(d, 2)

    witness_dim_original = (
        ceil((e + 2 * params.e_com + e * params.log_q) / d) + 3 + 8 * e * ell + zeta * t1
    )
    witness_dim = witness_dim_original * (d // 64)
    norm_bound = sqrt(
        e
        + eta * d
        + 4 * e * ell * d
        + params.sigma_com**2 * (2 * params.e_com)
        + sig**2 * (e * params.log_q + zeta * t1 * d)
        + sigring**2 * (2 * d + 4 * e * ell * d)
    )

    print(f"LaBRADOR witness dimension: {witness_dim}")

    r = 24
    n = ceil(witness_dim / 24)

    L = LaBRADOR(d=64, logq=32)
    size, params = L(n, r, norm_bound, 3, 5, verbose=False)
    return size, params

@cached_function
def voprf_pk_proof_sizekb(params, compress_ct1=True):
    """
    Well-formedness proof of the VOPRF public key using LaBRADOR.
    """
    e = params.tfhe["lwe_dimension"]
    eta = 1
    d = params.tfhe["polynomial_size"] * params.tfhe["glwe_dimension"]
    sig = params.tfhe["lwe_modular_std_dev"]
    sigring = params.tfhe["glwe_modular_std_dev"]
    ell = params.tfhe["pbs_level"]
    t1 = params.tfhe["ks_level"]

    if not compress_ct1:

```

```

        zeta = eta
    else:
        zeta = eta + log(d, 2)

    witness_dim_original = (
        ceil((3 * e + 2 * params.e_com + e * params.log_q) / d)
        + 4
        + 2 * eta
        + 16 * e * ell
        + 2 * zeta * tl
    )
    # need to expand this dimension to reach ring dimension 64
    witness_dim = witness_dim_original * ceil(d // 64)
    norm_bound = sqrt(
        3 * e
        + 2 * eta * d
        + 8 * e * ell * d
        + params.sigma_com**2 * (2 * params.e_com)
        + sig**2 * (e * params.log_q + 2 * tl * zeta * d)
        + sigring**2 * (4 * d + 8 * e * ell * d)
    )

    r = 256
    n = ceil(witness_dim / 256)
    L = LaBRADOR(d=64, logq=params.log_q)

    return L(n, r, norm_bound, 5, 5, verbose=False)

def oprf_ct0_sizekb(params):
    """
    OPRF request size in kilobytes.
    """

    keep_bits = ceil(-log(params.tfhe["lwe_modular_std_dev"], 2.0) + 1)

    return _kb(params.n_p * keep_bits + 256)

voprf_ct0_sizekb = oprf_ct0_sizekb

@cached_function
def oprf_ct0_proof_amortised_all_sizekb(
    params,
    gamma1=10,
    gamma2=5,
    compress=True,
    D=12,
    gamma=947 * 281 * 2,
    L=64,
    voprf=False,
):
    """
    Well-formedness proof of OPRF request using [LNP22].

    :param params:
    :param gamma1:
    :param gamma2:
    :param compress: Use Dilithium-style compression
    :param D:
    :param gamma:
    :param L: Amortise over this many proofs.
    :param voprf: Estimate VOPRF case.
    """
    d = 128
    lamb = ceil(128 / params.log_q)
    kappa = 2

```

```

e = params.tfhe["lwe_dimension"]
bcom = params.sigma_com * sqrt(params.e_com)
Bv = 128
Bct = 2*params.log_q * params.tfhe["lwe_modular_std_dev"] * sqrt(params.n_p)
if voprpf:
    m1 = ceil(params.e_com / d) + ceil(e / d) + 4 * L * ceil(params.n_p / d)
else:
    m1 = ceil(params.e_com / d) + ceil(e / d) + L * ceil(params.n_p / d)
ve = 2 + L
bm1 = m1 + ve - ceil(params.e_com / d)
if voprpf:
    mlwe_dim = 20
else:
    mlwe_dim = 11
delta = 1.0045

gammad = gamma2
gammae = gamma2
if voprpf:
    sd = gammad * sqrt(337) * sqrt(L * Bv**2 + L * params.n_p)
else:
    sd = gammad * sqrt(337) * sqrt(L * Bv**2)
if voprpf:
    voprpf_expand = 3
else:
    voprpf_expand = 1
se = (
    gammae
    * sqrt(337)
    * sqrt(
        e
        + voprpf_expand * L * params.n_p
        + 2 * bcom**2
        + L * Bct**2
        + 2 * log(bcom, 2)
        + L * log(Bct, 2)
    )
)
s1 = gamma1 * 59 * sqrt(bcom**2 + d * bm1)

print(
    f"Check that ternary MLWE in dimension {mlwe_dim} x {d} mod 2^{params.log_q} is hard."
)

for n in range(1, 32):
    m2 = mlwe_dim + n + lamb + 4
    s2 = gamma2 * 59 * sqrt(d * m2)

    beta = 8 * 59 * sqrt(2 * d * (m1 + ve) * s1**2 + 2 * d * m2 * s2**2)

    # Is n picked so that the sis problem is hard?
    if beta > 2 ** (2 * sqrt(n * d * params.log_q * log(delta, 2))):
        continue

    tot = (
        (n + 2 * ceil(256 / d) + lamb + 2) * d * params.log_q
        + log(2 * kappa + 1, 2) * d
        + (m1 + ve) * d * log(s1 * 2**2.57, 2)
        + m2 * d * log(s2 * 2**2.57, 2)
        + 256 * (log(se * 2.57, 2) + log(sd * 2.57, 2))
    )

    if compress:
        return round(
            _kb(tot)
            - dilithium_sistest(
                D, gamma, params, m2, n, gamma1, gamma2, L, voprpf=voprpf
            ),

```

```

        1,
    )
    else:
        return _kb(tot)

oprpf_ct0_proof_sizekb = partial(
    oprf_ct0_proof_amortised_all_sizekb,
    gamma1=10,
    gamma2=5,
    compress=False,
    D=11,
    gamma=957 * 281,
    L=1,
)

voprpf_ct0_proof_amortised_all_sizekb = partial(
    oprf_ct0_proof_amortised_all_sizekb, voprpf=True
)

voprpf_ct0_proof_sizekb = partial(
    voprpf_ct0_proof_amortised_all_sizekb,
    gamma1=10,
    gamma2=5,
    D=11,
    gamma=4051 * 41 * 2,
    L=1,
)

def dilithium_sistest(
    D, gamma, params, m2=32, n=13, gamma1=10, gamma2=5, L=1, voprpf=False
):
    d = 128
    e = params.tffe["lwe_dimension"]
    bcom = params.sigma_com * sqrt(params.e_com)

    expansion = 1 if not voprpf else 4

    m1 = ceil(params.e_com / d) + ceil(e / d) + L * expansion * ceil(params.n_p / d)
    ve = 2 + L
    bm1 = m1 + ve - ceil(params.e_com / d)
    s1 = gamma1 * 59 * sqrt(bcom**2 + d * bm1)
    s2 = gamma2 * 59 * sqrt(d * m2)
    kappa = 2

    delta = 1.0045
    b1 = 2 * s1 * sqrt(2 * (m1 + ve) * d)
    b2 = 2 * s2 * sqrt(2 * m2 * d) + 2**D * 59 * sqrt(n * d) + gamma * sqrt(n * d)
    betap = 4 * 59 * sqrt(b1**2 + b2**2)

    # make sure the sis problem is still hard and the compression params make sense
    if betap >= 2 ** (2 * sqrt(n * d * params.log_q * log(delta, 2))):
        raise ValueError(f"betap {betap} too large.")
    if 2 ** (D - 1) * kappa * d >= gamma:
        raise ValueError(f"gamma {gamma} too small.")

    return _kb((D - 2.25) * n * d)

def oprf_ct0_proof_amortised_sizekb(params, L=64):
    """
    Well-formedness proof of OPRF request per request using [LNP22].
    """
    return round(oprf_ct0_proof_amortised_all_sizekb(params, L=L) / L, 1)

def voprpf_ct0_proof_amortised_sizekb(params, L=64):

```

```

"""
Well-formedness proof of VOPRF request per request using [LNP22].
"""
return round(voprft_ct0_proof_amortised_all_sizekb(params, L=L) / L, 1)

def oprft_ct1_sizekb(params, compress_ct1=True):
"""
Response size in kilobytes.
"""
# TODO make number of bits we keep dependent on params
if compress_ct1 is False:
    e = params.tfhe["lwe_dimension"]
    return _kb(params.m * e * 24 + params.m * 16)
else:
    e11 = params.tfhe["glwe_dimension"]
    d = params.tfhe["polynomial_size"]
    return _kb(d * e11 * 24 + params.m * 16)

voprft_ct1_sizekb = oprft_ct1_sizekb

def oprft_online_sizekb(params, compress_ct1=True, amortise=True):
    r = oprft_ct0_sizekb(params)
    if amortise:
        r += oprft_ct0_proof_amortised_sizekb(params)
    else:
        r += oprft_ct0_proof_sizekb(params)
    r += oprft_ct1_sizekb(params, compress_ct1=compress_ct1)
    return round(r, 1)

def voprft_online_sizekb(params, amortise=True, compress_ct1=True):
    r = voprft_ct0_sizekb(params)
    if amortise:
        r += voprft_ct0_proof_amortised_sizekb(params)
    else:
        r += voprft_ct0_proof_sizekb(params)
    r += voprft_ct1_sizekb(params, compress_ct1=compress_ct1)
    return round(r, 1)

def oprft(params):
    consistency_check(params)

    print(
        f"""
% OPRF SIZES
/oprf/pk/sizeb/.initial={oprf_pk_sizeb(params)},
/oprf/pk/proof-old/sizeb/.initial={oprf_pk_proof_old_sizeb(params)},
/oprf/pk/proof/sizeb/.initial={oprf_pk_proof_sizekb(params)[0]},
/oprf/pk/proof/params/.initial={{oprf_pk_proof_sizekb(params)[1]}}},
/oprf/ct0/sizeb/.initial={oprf_ct0_sizekb(params)},
/oprf/ct0/proof/sizeb/.initial={oprf_ct0_proof_sizekb(params)},
/oprf/ct0/proof/amortised-all/sizeb/.initial={oprf_ct0_proof_amortised_all_sizekb(params)},
/oprf/ct0/proof/amortised/sizeb/.initial={oprf_ct0_proof_amortised_sizekb(params)},
/oprf/ct1/sizeb/.initial={oprf_ct1_sizekb(params, compress_ct1=False)},
/oprf/ct1-compressed/sizeb/.initial={oprf_ct1_sizekb(params)},
/oprf/online/amortised/sizeb/.initial={oprf_online_sizekb(params)},
/oprf/online/sizeb/.initial={oprf_online_sizekb(params, amortise=False)},
"""
    )

def voprft(params):
    consistency_check(params)

```

```

    print(
        f"""
% VOPRF SIZES
/voprf/pk/sizeb/.initial={voprf_pk_sizeb(params)},
/voprf/pk/proof-old/sizeb/.initial={voprf_pk_proof_old_sizeb(params)},
/voprf/pk/proof/sizeb/.initial={voprf_pk_proof_sizeb(params)[0]},
/voprf/pk/proof/params/.initial={{voprf_pk_proof_sizeb(params)[1]}},
/voprf/ct0/sizeb/.initial={voprf_ct0_sizeb(params)},
/voprf/ct0/proof/sizeb/.initial={voprf_ct0_proof_sizeb(params)},
/voprf/ct0/proof/amortised-all/sizeb/.initial={voprf_ct0_proof_amortised_all_sizeb(params)},
/voprf/ct0/proof/amortised/sizeb/.initial={voprf_ct0_proof_amortised_sizeb(params)},
/voprf/ct1/sizeb/.initial={voprf_ct1_sizeb(params, compress_ct1=False)},
/voprf/ct1-compressed/sizeb/.initial={voprf_ct1_sizeb(params)},
/voprf/online/amortised/sizeb/.initial={voprf_online_sizeb(params)},
/voprf/online/sizeb/.initial={voprf_online_sizeb(params, amortise=False)},
        """
    )

def both():
    oprf(OPRF)
    voprf(VOPRF)

```

$F_{\text{vpoprf}}.\text{Setup}(1^\lambda)$ $\mathbf{A}_{\text{pp}} \leftarrow \mathbb{Z}_Q^{N \times N}$ $\text{pp} \leftarrow \mathbf{A}_{\text{pp}}$ return pp	$F_{\text{vpoprf}}.\text{KeyGen}(1^\lambda)$ $\text{pk} \leftarrow \perp$ $\text{sk} \leftarrow \mathbb{Z}_p^{m_p \times n_p}$ $\mathbf{x}_i \leftarrow X$ for $i \in \mathbb{Z}_\beta$ $\mathbf{z} \leftarrow F_{\text{vpoprf}}.\text{Eval}(\text{sk}, t = \mathbf{t}, x = \mathbf{x})$ $\pi \leftarrow \text{NIZKAoK}_{R_t}(\mathbf{A}; \mathbf{t}, \mathbf{z}, \mathbf{x})$ return (pk, sk)	$F_{\text{vpoprf}}.\text{Eval}(\text{sk} = \mathbf{A}, t = \perp, x = \mathbf{x})$ $\mathbf{z}' := F_{\mathbf{A}}(\mathbf{t}, \mathbf{x})$ $\mathbf{z} := \text{RO}_{\text{fin}}(\mathbf{t}, \mathbf{x}, \mathbf{z}')$ return z
$F_{\text{vpoprf}}.\text{Request}(\text{pk}, t = \perp, x = \mathbf{x}; \text{pp})$ Parse pp as $\mathbf{A}_{\text{pp}}, \hat{\mathbf{x}}, \hat{\mathbf{z}}, \hat{\pi}$ if $\hat{\pi}$ does not verify, abort $\text{FHE.pk}^{(\text{pp})}, \text{FHE.sk} \leftarrow \text{FHE.KeyGen}^{(\text{pp})}()$ $\rho \leftarrow S_m$ $\mathbf{x}' \leftarrow (\mathbf{x}, \hat{\mathbf{x}})$ $\mathbf{x}' \leftarrow \mathbf{x}'_{\rho(i)}$ for $i \in \mathbb{Z}_r$ $\mathbf{y} \leftarrow \text{decomp}(\mathbf{G}_{\text{inp}} \cdot \mathbf{x}' \bmod p)$ $\text{ct} \leftarrow \text{FHE.Enc}(\text{FHE.sk}, \mathbf{y})$ $\pi \leftarrow \text{NIZKAoK}_C(\text{FHE.pk}^{(\text{pp})}, \text{ct}; \text{FHE.sk}, x)$ $\text{req} \leftarrow (\text{FHE.pk}^{(\text{pp})}, \text{ct}, \pi, t)$ return req, ρ	$F_{\text{vpoprf}}.\text{BlindEval}(\text{sk} = \mathbf{A}, t = \perp, \text{req}; \text{pp})$ $(\text{FHE.pk}^{(\text{pp})}, \text{ct}, \pi) \leftarrow \text{req}$ $\text{ct}' \leftarrow F_{\text{vpoprf}}.\text{HEEval}(\text{FHE.pk}^{(\text{pp})}, \mathbf{A}, \mathbf{t}, \text{ct})$ if π does not verify then $\text{ct}' = \perp$ $\text{rep} \leftarrow \text{ct}'$ return rep	$F_{\text{vpoprf}}.\text{Finalise}(\text{FHE.sk}, \text{rep} = \text{ct}, \rho)$ if ct' not a ctxt then return \perp $\mathbf{z}^* \leftarrow \text{FHE.Dec}(\text{FHE.sk}, \text{ct})$ if $F_{\text{vpoprf}}.\text{Verify}(\mathbf{z}, \mathbf{z}^*, \rho) = 0$, then return \perp $\mathbf{z}'_i \leftarrow \mathbf{z}^*_{\rho(i)}$ for $i \in \mathbb{Z}_\alpha$ $\mathbf{z} \leftarrow \text{RO}_{\text{fin}}(\mathbf{t}, \mathbf{x}, \mathbf{z}')$ return z
$F_{\text{vpoprf}}.\text{Verify}(\mathbf{z}, \mathbf{z}^*, \rho)$ if $\mathbf{z}_{j, \rho(k+\alpha)}^{(\beta)} \neq \mathbf{z}^*_{\rho(k+\alpha)}$ for any $k \in \mathbb{Z}_\beta$, return 0 if $\mathbf{z}_{\rho(i_0+r_\alpha \ell)}^{(\alpha)} \neq \mathbf{z}^*_{\rho(i_1+r_\alpha \ell)}$ for $i_0 \neq i_1 \in \mathbb{Z}_{r_\alpha}$, return 0 if $\mathbf{z}_{\rho(i_0+r_\alpha \ell_0)}^{(\alpha)} = \mathbf{z}^*_{\rho(i_1+r_\alpha \ell_1)}$ for $i_0, i_1 \in \mathbb{Z}_{r_\alpha}$ and $\ell_0 \neq \ell_1 \in \mathbb{Z}_\alpha$, return 0 return 1		




Fig. 5. Our construction for VPOPRF, F_{vpoprf} .