

# A Fast and Simple Partially Oblivious PRF, with Applications

Nirvan Tyagi<sup>1</sup>, Sofía Celi<sup>2</sup>, Thomas Ristenpart<sup>3</sup>,  
Nick Sullivan<sup>2</sup>, Stefano Tessaro<sup>4</sup>, and Christopher A. Wood<sup>2</sup>

<sup>1</sup> Cornell University

<sup>2</sup> Cloudflare

<sup>3</sup> Cornell Tech

<sup>4</sup> University of Washington

**Abstract.** We build the first construction of a partially oblivious pseudorandom function (POPRF) that does not rely on bilinear pairings. Our construction can be viewed as combining elements of the 2HashDH OPRF of Jarecki, Kiayias, and Krawczyk with the Dodis-Yampolskiy PRF. We analyze our POPRF’s security in the random oracle model via reduction to a new one-more gap strong Diffie-Hellman inversion assumption. The most significant technical challenge is establishing confidence in the new assumption, which requires new proof techniques that enable us to show that its hardness is implied by the  $q$ -DL assumption in the algebraic group model.

Our new construction is as fast as the current, standards-track OPRF 2HashDH protocol, yet provides a new degree of flexibility useful in a variety of applications. We show how POPRFs can be used to prevent token hoarding attacks against Privacy Pass, reduce key management complexity in the OPAQUE password authenticated key exchange protocol, and ensure stronger security for password breach alerting services.

## 1 Introduction

An oblivious pseudorandom function (OPRF) [FIPR05, JL09] allows a client holding a private input  $x$  and a server holding a key  $sk$  for a PRF  $f$  to engage in a protocol to *obliviously* evaluate  $f_{sk}$  on  $x$ . The client learns (and optionally verifies) the evaluation  $f_{sk}(x)$  while the server learns nothing. *Partially*-oblivious PRFs (POPRF), first introduced by Everspaugh et al. in the context of the Pythia password hardening system [ECS<sup>+</sup>15], extend this functionality to include a public input (or metadata tag)  $t$  for the PRF evaluation. A client learns (and, optionally, verifies)  $f_{sk}(t, x)$  where  $t$  is known by both server and client; the private input  $x$  remains hidden.

OPRFs are increasingly becoming a critical cryptographic tool for privacy-preserving protocols. Examples include one-time use anonymous credentials for spam prevention [DGS<sup>+</sup>18, HJK<sup>+</sup>21], private set intersection (PSI) for checking compromised credentials [LPA<sup>+</sup>19, TPY<sup>+</sup>19], de-identified authenticated logging [HIJ<sup>+</sup>21], and password-authenticated key exchange [JKX18, JKK14]. In all

these applications, we observe that there is a need to “partition” the PRF in a productive manner, i.e., allowing computation of  $f_{sk}(t, x)$  using domain separation on some public value  $t$ . OPRF blinding protocols do not support this in a secure manner, because the server cannot verify what  $t$  is used within a client’s oblivious request. Most OPRF applications therefore use a separate key instance for each  $t$ , with an associated increase in key management complexity. Huang et al. [HIJ<sup>+</sup>21] instead suggest using what they call an “attribute-based OPRF”, which can be instantiated via any POPRF. But the only known POPRF [ECS<sup>+</sup>15] relies on bilinear pairings, which slows performance relative to the best known OPRF and also complicates deployment given the lack of widespread implementation support for pairings.

In this work, we introduce a new POPRF that combines aspects of the 2HashDH OPRF of Jarecki et al. [JKK14], that is the de facto standard used in practice, with the Dodis-Yampolskiy (DY) verifiable random function [DY05]. Our POPRF is also closely related to a signature scheme suggested by Zhang, Safavi-Naini, and Susilo (ZSS) [ZSS04, ZSS03]. Our new POPRF, called 3HashS-DHI, is essentially as performant as 2HashDH and does not rely on pairings, thereby enabling support for a public input virtually *for free*. While 3HashS-DHI’s protocol is simple, its analysis is not, requiring a new interactive discrete log (DL) assumption whose security we reduce to  $q$ -DL in the algebraic group model [FKL18]. We also provide new formal security notions for POPRFs and (as a special case) OPRFs, which we believe will be of independent interest.

**Formal syntax and security notions for POPRFs.** We start with the latter contribution. We provide a new formalization for POPRFs, including syntax, semantics, and security definitions. Our formal syntax builds off of [ECS<sup>+</sup>15] and follows closely to that used in the related OPRF formalization [JL09, JKK14]. In terms of security, we propose new property-based security definitions that cover pseudorandomness (in the face of malicious clients) as well as request privacy and verifiability (in the face of malicious servers). Our property-based security games avoid the ideal function based formulations inherited from 2PC used in prior works on OPRF; they also avoid the non-standard “one-more” PRF security definition of [ECS<sup>+</sup>15].

Our *pseudorandomness* notion for POPRFs guarantees that the evaluation outputs look random to a malicious client, even when the malicious client has access to a blinded evaluation oracle. It is formalized with a simulation-based indistinguishability game that takes rough inspiration from the UC-style all-in-one OPRF security definition of [JKK14] and prior notions on partially blind signatures [AF96]. Here an adversary must distinguish between real evaluations of the PRF given access to a blind evaluation oracle, and evaluations of a random function given access to a simulated blind evaluation oracle. The simulator can receive random function evaluations on a limited number of points for any given public input  $t$ , where the limit is determined by the number of times the adversary has queried the blind evaluation oracle for that  $t$ . This restriction captures that only one random function evaluation is learned for each blind evaluation. Note that our accounting is more granular than the general “ticket-

ing” approaches of blind UC protocols [JKK14, KZ08, Fis06]), due to the need of tying invocations to particular  $t$  values.

Our other notion is *request privacy* which captures that nothing about a client message  $x$  should leak to a malicious server during an oblivious evaluation, and, moreover, the server should not be able to link an output  $f_{sk}(t, x)$  to particular oblivious request transcripts. The latter is often referred to as a linking attack, and is problematic in various applications of POPRFs. Our request privacy notion comes in two flavors, depending on whether the malicious server behaves passively or actively. The former allows us to analyze the privacy of schemes that do not allow verification that a server legitimately computed the blinded evaluation protocol; the latter requires schemes to allow client-side verification of the server’s response.

**The 3HashSDHI construction.** The main contribution of this work is a new construction of a POPRF, which we call 3HashSDHI. The name refers to its use of three hashes and its reliance on the strong Diffie-Hellman inversion assumption. Its starting point is the 2HashDH construction of Jarecki et al. [JKK14], whose full PRF evaluation we define as  $2\text{HashDH.Ev}(sk, x) = H_2(x, H_1(x)^{sk})$ . The blinded evaluation protocol has the client send  $B = H_1(x)^r$  for random  $r$ , and the server respond with  $B' = B^{sk}$ . The client can unblind to  $(B')^{1/r} = H_1(x)^{sk}$  in order to complete the evaluation of the function. Here operations are over a prime-order group (written multiplicatively) such as an elliptic curve. Proof of evaluation consists of a simple Chaum-Pedersen proof of discrete log equality [CP92] proving  $\log_g pk = \log_B B'$  where  $pk = g^{sk}$  is the server’s public key. As mentioned, 2HashDH is already in use in practice [DGS<sup>+</sup>18, TPY<sup>+</sup>19, HIJ<sup>+</sup>21] and is on track to become a standard [DFHSW20].

We want a way to extend 2HashDH to allow public tags. To do so, we take inspiration from the Dodis-Yampolskiy PRF, whose evaluation is defined as  $DY.Ev(sk, t) = g^{1/(sk+t)}$ . Put together, the 3HashSDHI scheme gives a PRF evaluated as:

$$3H.Ev(sk, t, x) = H_2\left(t, x, H_1(x)^{1/(sk+H_3(t))}\right).$$

It can therefore be interpreted as evaluating the Dodis-Yampolskiy PRF on the public input  $t$  over a random generator determined by the private input  $x$ , followed by a final hashing step. The basic structure of  $H_1(x)^{1/(sk+H_3(t))}$  was also described in an attempt to build secure partially blind signatures by ZSS [ZSS03]; we discuss ZSS further in Section 4.

To perform a blind evaluation, the client hashes and blinds their private input as  $B = H_1(x)^r$  using a random scalar  $r$  and sends  $B$  to the server holding  $sk$ . The server computes and sends back to the client the strong Diffie-Hellman inversion  $B' = B^{1/(sk+H_3(t))}$  of the blinded element using the secret key and public hash of the public input  $t$ . The client can unblind by computing  $(B')^{1/r} = H_1(x)^{1/(sk+H_3(t))}$  and then complete the evaluation by hashing appropriately. To provide verifiability, the server uses a Chaum-Pedersen zero-knowledge proof (ZKP) of discrete log equality to prove  $\log_g pk' = \log_{B'} B$  where  $pk' = pk \cdot g^{H_3(t)}$  which can be easily computed from public values by the client.

Our protocol incurs minimal overhead on top of the OPRF blind evaluation of 2HashDH, requiring only an extra hash computation, group operation, and scalar inversion. It makes use of the same Chaum-Pedersen proof for verifiability, which, as has been observed for 2HashDH, allows for evaluation of a batch of inputs whilst only constructing one Chaum-Pedersen proof [DGS<sup>+</sup>18,DFHSW20] (provided the batch is for the same public metadata tag  $t$ ).

We formally show request privacy against passive adversaries (without ZKP) holds based just on the randomness of the blinding, and that request privacy against malicious adversaries holds additionally assuming the ZKP is sound. The key technical challenge is proving the new POPRF is pseudorandom.

As is seemingly requisite for schemes with blinded evaluation protocols, we prove the pseudorandomness security of our scheme with respect to a one-more gap style assumption [BNPS03,Bol03]. In fact the algebraic structure exposed to adversarial clients by the 3HashSDHI blinded evaluation protocol — raising an arbitrary group element  $Y$  to  $1/(sk+H_3(t))$  for adversarial  $t$  — requires new proof techniques compared to prior approaches. We start by introducing a new one-more gap strong Diffie-Hellman inversion (OM-Gap-SDHI) assumption, based on the perceived hardness of computing  $Y^{1/(x+c)}$  for any base  $Y$  and (restricted) scalars  $c$ . We show via a relatively straightforward proof that this assumption is sufficient to prove POPRF pseudorandomness for 3HashSDHI, modeling the hash functions as random oracles. Additionally, the verifiable version requires that the ZKP is zero-knowledge.

The main difficulty is analyzing the security of our new computational assumption. In particular, for given distinct constants  $c_1, \dots, c_n$ , the assumption considers a setting with an oracle SDH returning  $B^{1/(x+c_i)}$  on input  $(B, i)$ . Given some additional random group elements  $Y_1, \dots, Y_m$ , it requires it to be hard to compute  $\ell$  elements  $Y_{i_1}^{1/(x+c_i)}, \dots, Y_{i_\ell}^{1/(x+c_i)}$ , for any  $i \in [n]$  and for distinct  $i_1, \dots, i_\ell \in [m]$ , by while fewer than  $\ell$  queries  $\text{SDH}(\cdot, i)$ . The challenge is that we do *not* restrict the number of queries  $\text{SDH}(X, j)$  for  $j \neq i$ , and this could be for group elements of  $X$  that depend on  $c_i$  (e.g.,  $X$  is a prior output of an  $\text{SDH}(\cdot, i)$  query). Ultimately, we show in the algebraic group model (AGM) [FKL18] that the assumption reduces to one of the uber assumptions from Bauer, Fuchsbauer, and Loss [BFL20], and therefore, in turn, is implied by the  $q$ -DL assumption, where  $q$  is a bound on the number of oracle queries. This AGM analysis implies hardness of the new assumption in the generic group model (GGM) [Sho97, Mau05].

In terms of concrete security, our analyses shows that, roughly speaking, 3HashSDHI is as hard as breaking the  $q$ -DL problem. Actually our main AGM proof is loose by a factor that is the maximum number of blind evaluation queries made by an adversary. Whether this AGM analysis can be tightened is an open question, but we observe in the body that a slight alternative to our AGM analysis gives a tight reduction in the GGM. We suggest using this tighter analysis to drive parameter selection: the best known attack against  $q$ -DL is due to Cheon [Che06] and indicates that a 256-bit group suffices for 80-bit security and a 384-bit group for 128-bit security. Importantly this matches the situation

for 2HashDH, and so moving to 3HashSDHI does not require changing group parameters to achieve the desired security levels.

**Applications of our POPRF.** Equipped with our new POPRF, we return to our motivating applications and show how swapping in a POPRF for the existing OPRF can lead to various benefits for deployments.

*One-time use anonymous credentials.* Privacy Pass [DGS<sup>+</sup>18, CDFH21] is a protocol in which clients may be issued one-time use tokens that can later be redeemed anonymously to authenticate themselves. It has been proposed for use in the context of content distribution networks and web advertising, requiring users to authenticate with a token, and thereby reducing malicious web requests, protecting against e.g., denial-of-service attacks and fraudulent advertisement conversions. Tokens are issued to users that prove trustworthiness, e.g., through a CAPTCHA challenge. The protocol is being considered for standardization by both the IETF and the World Wide Web Consortium (W3C), and a prototype deployment is already in production use by Cloudflare, hCaptcha and others.

An OPRF is the core component of the protocol. Tokens are issued via an OPRF in which users obtain evaluations at random points, storing the point  $x$  and evaluation  $y$ . Redeeming a token simply involves showing the pair  $(x, y)$ , which the server can check is valid, but cannot link  $x$  back to an issuance due to the oblivious evaluation. The server stores a strikelist of used tokens to prevent double spending. Additionally, all servers perform a global double-spend check to avoid clients from exploiting the possibility of spending tokens more than once against distributed token checking systems.

An abuse of the protocol that has been observed in its early use is individual users (or groups of users) gathering tokens over a long period of time and redeeming them all at once, e.g., in an attempt to overwhelm a website. We refer to such behavior as a hoarding attack. An easy way to mitigate the damage of a hoarding attack is by expiring old unspent tokens after an amount of time: the way to do this with an OPRF is by rotating the OPRF key (a mechanism that is limited in frequency). However, establishing trust in a frequently-rotating key is a challenging problem. More so, it is especially important clients have trust in the OPRF key, as a server that equivocates on their public key can link token issuances and redemptions, by, for example, using a unique public key for each issuance. As we show, POPRFs address the issue of expiring tokens without the need of rotating keys by using the public metadata input to encode an expiration epoch.

*Bucketized PSI for checking compromised credentials.* Password breach alerting protocols [TPY<sup>+</sup>19, LPA<sup>+</sup>19] allow a user to query to determine if their username, password pair  $(u, pw)$  has appeared in a dataset  $D$  of known breaches. If so, the user is vulnerable to credential stuffing attacks and should change their password. Current services for breach alerting rely on an ad hoc 2HashDH-based private-set membership protocol that achieves scalability via bucketization: the user sends a truncated hash  $H(u)$  of their username to identify a subset  $B \subseteq D$  that have matching truncated username hash. A 2HashDH-based protocol is

then performed over  $B$ : the client obviously evaluates  $2\text{HashDH.Ev}(sk, u \parallel pw)$  with  $sk$  held by server, and also obtains the OPRF outputs for all the values in the bucket  $B$ . Bucketization ensures scalability by limiting  $|B|$  despite  $|D|$  being on the order of billions of username, password pairs.

One issue is that currently deployed protocols provide no cryptographic binding between the bucket identifier  $H(u)$  and the blinded OPRF output: a malicious client can query for arbitrary usernames, not just ones that match  $H(u)$ . Whether this is a significant security problem in practice is not clear, but we note that POPRFs easily rectify it by setting  $t = H(u)$  and replacing  $2\text{HashDH}$  above with  $3\text{HashSDHI}$ .

*Asymmetric password-authenticated key exchange.* Password authenticated key exchange (PAKE) protocols [BM93] allow a client and server to establish a shared session key authenticated by a short password. Strong asymmetric PAKE (SaPAKE) protocols [JKX18] additionally ensure that the server can store (just) what amount to salted hashes of user passwords, thereby making it so that PAKEs can achieve the same resistance as standard password-based authentication in the case of a server breach. The OPAQUE [JKX18] SaPAKE protocol uses an OPRF as one of its core components; it is currently being considered for standardization by the IETF [KLW21]. The OPRF suggested for use is  $2\text{HashDH}$ .

In OPAQUE the server uses a separate OPRF key for each user. We show how we can instead use our  $3\text{HashSDHI}$  POPRF to allow OPAQUE to work with a single master key  $pk$ ; diversity across users can then be provided using usernames as the public input  $t$  to  $3\text{HashSDHI}$ . We believe that this will simplify deployments and potentially improve their security, as discussed in the body.

## 2 Preliminaries

### 2.1 Algebraic Group Model

In some of our security proofs, we consider security against *algebraic* adversaries which we model using the algebraic group model, following the treatment of [FKL18]. We call an algorithm  $\mathcal{A}$  *algebraic* if for all group elements  $Z$  that are output (either as final output or as input to oracles),  $\mathcal{A}$  additionally provides the representation of  $Z$  relative to all previously received group elements. The previous received group elements include both original inputs to the algorithm and outputs received from calls to oracles. More specifically, if  $[X]_i$  is the list of group elements  $[X_0, \dots, X_n] \in \mathbb{G}$  that  $\mathcal{A}$  has received so far, then, when producing group element  $Z$ ,  $\mathcal{A}$  must also provide a list  $[z]_i = [z_0, \dots, z_n]$  such that  $Z = \prod_i X_i^{z_i}$ .

### 2.2 Random oracle model

We will prove security using ideal primitives, modeling 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 of [JT20]. An ideal primitive  $P$  specifies algorithms  $P.\text{Init}$  and  $P.\text{Eval}$ . The initialization algorithm has syntax  $st_P \leftarrow P.\text{Init}(1^\lambda)$ . The stateful evaluation algorithm has syntax  $y \leftarrow P.\text{Eval}(x : st_P)$ . We sometimes use  $A^P$  as shorthand for giving algorithm  $A$  oracle access to  $P.\text{Eval}(\cdot : st_P)$ . The stateful formulation of the ideal primitive is used to allow for efficient instantiation in our security proofs, e.g., by “lazy sampling”. For example, a random oracle can be written to be stateless, but it would be inefficient to have to store a huge random table. We can combine access to multiple ideal primitives  $P = P_1 \times \dots \times P_m$  as follows:

$$\begin{array}{ll}
\frac{P.\text{Init}(1^\lambda)}{[st_{P,i}]_i^m \leftarrow [P_i.\text{Init}(1^\lambda)]_i^m} & \frac{P.\text{Eval}(x : [st_{P,i}]_i^m)}{(i, x) \leftarrow x} \\
\text{Return } [st_{P,i}]_i^m & y \leftarrow P_i.\text{Eval}(x : st_{P,i}) \\
& \text{Return } y
\end{array}$$

To concretize the above, we focus on random oracles. We define a random oracle that takes arbitrary input and produces random output from a sampling algorithm  $\text{Samp}$ . It is captured by the ideal primitive  $\text{RO}[\text{Samp}] = (\text{RO}.\text{Init}, \text{RO}.\text{H})$  defined as follows. When the range is clear from context,  $\text{Samp}$  may be omitted.

$$\begin{array}{ll}
\frac{\text{RO}.\text{Init}(1^\lambda)}{T \leftarrow [\cdot]} & \frac{\text{RO}.\text{Eval}(x : T)}{\text{If } x \notin T \text{ then } T[x] \leftarrow \text{Samp}()} \\
\text{Return } T & \text{Return } T[x]
\end{array}$$

When clear from context and in an abuse of notation (since we will use  $H_i$  to denote a hash function as well), we will write  $P = H_1 \times \dots \times H_m$  as the ideal primitive that gives access to  $m$  random oracles, accessible by querying directly an oracle labeled  $H_i$ .

**Algebraic algorithms in the random oracle model.** As in [FPS20], to support algebraic algorithms, we will require the structure of the domain and range to be specified for any random oracle  $\text{RO}$ . We assume an input can be efficiently checked to be a valid member of the domain and perform such checks implicitly returning  $\perp$  if they fail. We will require that algebraic algorithms provide representations for any group element input, specified as part of the domain of  $\text{RO}$ . And similarly, any group element output of  $\text{RO}$  is included in the list of received group elements for the algebraic adversary.

### 2.3 Non-interactive Zero Knowledge Proofs

We define a non-interactive proof system  $\text{NiZK}$  over an efficiently computable relation  $\mathcal{R}$  defined over pairs  $(x, w)$  where  $x$  is called the *statement* and  $w$  is called the *witness*. It is made up of the following algorithms. The setup algorithm produces the public parameters for execution,  $pp \leftarrow \text{NiZK}.\text{Setup}(\lambda)$ . The proving algorithm takes a witness and statement and produces a proof,  $\pi \leftarrow \text{NiZK}.\text{Prove}_{pp}^P(w, x)$ . The verification algorithm verifies the proof for a statement,  $b \leftarrow \text{NiZK}.\text{Ver}_{pp}^P(x, \pi)$ . We define the following security properties.

Game $\text{SOUND}_{\text{NiZK}, \mathcal{R}, \mathcal{P}}^{\mathcal{A}}(\lambda)$	Game $\text{ZK}_{\text{NiZK}, \mathcal{R}, \mathcal{S}, \mathcal{P}}^{\mathcal{A}, b}(\lambda)$	Oracle $\text{PROVE}(x, w)$	Oracle $\text{PRIM}(x)$
$pp \leftarrow \text{NiZK.Setup}(\lambda)$ $st_{\mathcal{P}} \leftarrow \text{P.Init}(\lambda)$ $(x, \pi) \leftarrow \mathcal{A}^{\mathcal{P}}(pp)$ Return $\bigwedge \left( \begin{array}{l} \text{NiZK.Ver}^{\mathcal{P}}(x, \pi) \\ \exists w : (x, w) \in \mathcal{R} \end{array} \right)$	$pp \leftarrow \text{NiZK.Setup}(\lambda)$ $st_{\mathcal{P}} \leftarrow \text{P.Init}(\lambda)$ $st_{\mathcal{S}} \leftarrow \text{S.Init}(pp)$ $b' \leftarrow \mathcal{A}^{\text{PRIM}, \text{PROVE}}(pp)$ Return $b'$	Require $(x, w) \in \mathcal{R}$ $\pi_1 \leftarrow \text{NiZK.Prove}^{\mathcal{P}}(x, w)$ $\pi_0 \leftarrow \text{S.Prove}(x : st_{\mathcal{S}})$ Return $\pi_b$	$y_1 \leftarrow \text{P.Eval}(x : st_{\mathcal{P}})$ $y_0 \leftarrow \text{S.Eval}(x : st_{\mathcal{S}})$ Return $y_b$

Fig. 1: Soundness (left) and zero knowledge (right) security games for non-interactive zero knowledge proof systems.

**Completeness.** A proof system is *complete* if given a true statement, a prover with a witness can convince the verifier. We will make use of a proof system with perfect completeness. A proof system has *perfect completeness* if for all  $(x, w) \in \mathcal{R}$ ,

$$\Pr \left[ \text{NiZK.Ver}_{pp}^{\mathcal{P}}(x, \text{NiZK.Prove}_{pp}^{\mathcal{P}}(w, x)) = 1 \right] = 1 .$$

**Knowledge soundness.** A proof system is computationally *knowledge sound* if whenever a prover is able to produce a valid proof for a statement  $x$ , it is a true statement, i.e., there exists some witness  $w$  such that  $(x, w) \in \mathcal{R}$ . Knowledge soundness is defined by the security game  $\text{SOUND}_{\text{NiZK}, \mathcal{R}, \mathcal{P}}^{\mathcal{A}}(\lambda)$  (Figure 1) in which an adversary is tasked with finding a verifying statement and proof where the statement is not in  $\mathcal{R}$ . The advantage of an adversary is defined as  $\text{Adv}_{\text{NiZK}, \mathcal{R}, \mathcal{P}, \mathcal{A}}^{\text{sound}}(\lambda) = \Pr[\text{SOUND}_{\text{NiZK}, \mathcal{R}, \mathcal{P}}^{\mathcal{A}}(\lambda) = 1]$  with respect to ideal primitive  $\mathcal{P}$ .

**Zero knowledge.** A proof system is computationally *zero-knowledge* if a proof does not leak any information besides the truth of a statement. Zero knowledge is defined by the security game  $\text{ZK}_{\text{NiZK}, \mathcal{R}, \mathcal{S}, \mathcal{P}}^{\mathcal{A}, b}(\lambda)$  (Figure 1) in which an adversary is tasked with distinguishing between proofs generated from a valid witness and simulated proofs generated without a witness. The advantage of an adversary is defined as

$$\text{Adv}_{\text{NiZK}, \mathcal{R}, \mathcal{S}, \mathcal{P}, \mathcal{A}}^{\text{zk}}(\lambda) = \left| \Pr[\text{ZK}_{\text{NiZK}, \mathcal{R}, \mathcal{S}, \mathcal{P}}^{\mathcal{A}, 1}(\lambda) = 1] - \Pr[\text{ZK}_{\text{NiZK}, \mathcal{R}, \mathcal{S}, \mathcal{P}}^{\mathcal{A}, 0}(\lambda) = 1] \right| ,$$

with respect to simulator algorithm  $\mathcal{S}$  and ideal primitive  $\mathcal{P}$ .

**Fiat-Shamir heuristic for Sigma protocols.** Our protocol requires a non-interactive zero knowledge proof for the relation including two pairs of group elements with equivalent discrete logs:

$$\mathcal{R} = \{(g, U, V, W), (\alpha) : U = g^{\alpha} \wedge W = V^{\alpha}\} .$$

This relation falls into a general family of relations of discrete log linear homomorphisms for which there exist so-called ‘‘Sigma protocols’’ [Cam98] to construct interactive proofs of knowledge. These can be made non-interactive using the Fiat-Shamir heuristic in the standard way. We denote  $\Sigma_{\mathcal{R}}[\text{GGen}]$  (shortened to  $\Sigma_{\mathcal{R}}$  for simplicity) as the resulting non-interactive proof system for  $\mathcal{R}$  known as the Chaum-Pedersen protocol [CP92] (shown in Figure 2); it is perfectly com-

$\Sigma_{\mathcal{R}}.\text{Prove}^{\text{H}}(\alpha, (g, U, V, W))$ $r \leftarrow \mathbb{S}_{\mathbb{Z}_p}$ $s_U \leftarrow g^r ; s_W \leftarrow V^r$ $c \leftarrow \text{H}(g \parallel U \parallel V \parallel W \parallel s_U \parallel s_W)$ $z \leftarrow r - c\alpha$ $\pi \leftarrow (z, c)$ Return $\pi$	$\Sigma_{\mathcal{R}}.\text{Ver}^{\text{H}}((g, U, V, W), \pi)$ $(z, c) \leftarrow \pi$ $s_U \leftarrow g^z U^c ; s_W \leftarrow V^z W^c$ Return $c = \text{H}(g \parallel U \parallel V \parallel W \parallel s_U \parallel s_W)$  $\mathcal{R} = \{(\alpha), (g, U, V, W) : U = g^\alpha \wedge W = V^\alpha\}$
--	--

Fig. 2: Description of Chaum-Pedersen discrete log equality Sigma protocol [CP92].

plete, computationally sound, and perfectly zero-knowledge in the random oracle model. We refer readers to [BS17, Figure 19.7] for construction of a simulator  $S_\Sigma$ , which leads to the following well-known result that we state for completeness:

**Theorem 1.** *The simulator  $S_\Sigma$  is such that for any RO-model adversary  $\mathcal{A}_{\text{zk}}$  against ZK of  $\Sigma_{\mathcal{R}}$ ,  $\text{Adv}_{\Sigma_{\mathcal{R}}, \mathcal{R}, S_\Sigma, \text{RO}, \mathcal{A}_{\text{zk}}}^{\text{zk}}(\lambda) \leq q_{\text{P}} \cdot (q_{\text{P}} + q_{\text{H}})/2^\lambda$ , where  $\mathcal{A}_{\text{zk}}$  makes at most  $q_{\text{P}}$  and  $q_{\text{H}}$  queries to PROVE and the random oracle  $\text{RO} : \mathbb{G}^6 \rightarrow \mathbb{Z}_p$ .*

### 3 Partially Oblivious Pseudorandom Functions

We provide a new formalization for POPRFs, including syntax, semantics, and security. Our formalization builds off that from [ECS<sup>+</sup>15], but we offer new security notions that cover simulation-based security as a PRF (in the presence of a blinded evaluation oracle), client input privacy, and verifiability.

**Syntax and semantics.** A partially-oblivious pseudorandom function (POPRF) scheme,  $\text{Fn}$ , is a tuple of algorithms

$$(\text{Fn.Setup}, \text{Fn.KeyGen}, \text{Fn.Req}, \text{Fn.BlindEv}, \text{Fn.Finalize}, \text{Fn.Ev}) .$$

The setup and key generation algorithm generate public parameters  $pp$  and a public key, secret key pair  $(pk, sk)$ , respectively. Oblivious evaluation is carried out as an interactive protocol run between client and server. The protocols we consider in this work make use of only a single round of interaction, so we simplify the syntax of the interactive oblivious evaluation protocol into algorithms ( $\text{Fn.Req}$ ,  $\text{Fn.BlindEv}$ ,  $\text{Fn.Finalize}$ ) that work as follows:

- (1) First, a client runs the algorithm  $\text{Fn.Req}_{pp}^{\text{P}}(pk, t, x)$ , which takes input a public key  $pk$ , tag (or public input)  $t$ , and message  $x$ , and outputs a local state  $st$  and a request message  $req$ . The message  $req$  is sent to a server.
- (2) A server runs algorithm  $\text{Fn.BlindEv}_{pp}^{\text{P}}(sk, t, req)$ , using as input a secret key, a tag  $t$ , and the request message. It produces a response message  $rep$  that should be sent back to the client.
- (3) Finally, the client runs the algorithm  $\text{Fn.Finalize}(rep : st)$  and outputs a PRF evaluation or  $\perp$  if the response message is rejected, for example, due to the verification check failing.

The unblinded evaluation algorithm  $\text{Fn.Ev}$  is deterministic, and takes as input a public key, secret key pair  $(pk, sk)$ , a message pair  $(t, x)$ , and outputs a PRF evaluation  $y$ . We also define sets  $\text{Fn.SK}$ ,  $\text{Fn.PK}$ ,  $\text{Fn.T}$ ,  $\text{Fn.M}$ , and  $\text{Fn.Out}$  representing the secret key, public key, tag, message, and output space, respectively. We define the input space  $\text{Fn.In} = \text{Fn.T} \times \text{Fn.M}$ . We assume efficient algorithms for sampling and membership queries on these sets. When it is clear from context, we drop the prefix  $\text{Fn}$  and subscript  $pp$  from algorithm names.

For correctness, we require that  $\text{Ev}$  is a function, and that the blinded and unblinded evaluations are consistent. To formalize the latter: we require that for any  $pp$  output from  $\text{Setup}$ , any  $pk, sk$  output by  $\text{KeyGen}$ , and any  $t, x$ , it holds that  $\Pr[\text{Ev}(sk, t, x) = y] = 1$  where the probability is taken over choice of  $y$  via the following process:

$$(st, req) \leftarrow \text{Req}^P(pk, t, x) ; rep \leftarrow \text{BlindEv}^P(sk, t, req) ; y \leftarrow \text{Finalize}^P(rep : st) .$$

**Security.** We introduce three new security definitions for POPRFs. We use code-based games mostly following the framework of Bellare and Rogaway [BR06].

**Pseudorandomness.** The first definition captures pseudorandomness, i.e., indistinguishability of the POPRF from a random function, even for malicious clients that have access to a blinded evaluation oracle. We borrow some elements from the UC definition for standard OPRFs from [JKK14], but opt for what we believe to be a simpler, standalone formulation. We also extend to handle partial obliviousness, which has some subtleties.

A pseudocode game appears in Figure 3. The game is parameterized by a security parameter  $\lambda$ , an adversary  $\mathcal{A}$ , a challenge bit  $b$ , a POPRF  $\text{Fn}$ , a simulator  $\text{S} = (\text{S.Init}, \text{S.BlindEv}, \text{S.Eval})$ , and an ideal primitive  $\text{P}$ . The last will be used for random oracles in our main result. A simulator is a triple of algorithms that share state (explicitly denoted by  $st_{\text{S}}$  in the game). Algorithm  $\text{S.Init}$  initializes the simulator state and outputs a public key for the game. Algorithm  $\text{S.BlindEv}$  simulates blinded evaluation response messages while  $\text{S.Eval}$  simulates random oracle queries. Importantly,  $\text{S.BlindEv}$  and  $\text{S.Eval}$  can obtain  $\text{EV}$  outputs, but they can only do so in a circumscribed way: the simulator has oracle access to  $\text{LIMEV}$  which limits the number of full evaluations it can obtain to be at most the number of queries so far made by the adversary to the  $\text{BLINDEV}$ . Importantly, this limit is per-metadata value  $t$  (indicated via the subscript): the  $\text{LIMEV}$  query on any particular  $t$  is bound by the total number of blinded evaluation queries on that particular  $t$ . This follows from similar granular restrictions in the partially blind signatures literature [AF96].

A weaker version of the game would simply cap the total number of queries to  $\text{LIMEV}$  by the total number of queries to  $\text{BLINDEV}$ . This notion is, however, too weak for applications because we would like to ensure that querying, say, three times on public input  $t_1$  cannot somehow help an adversary complete the evaluation for another public input  $t_2 \neq t_1$ . We note that a recent preprint [SS21] contained this weaker notion, couched in the context of Privacy Pass. (We discuss this paper further in Section 4.)

Game $\text{POPRF}_{\text{Fn,S,P}}^{\mathcal{A},b}(\lambda)$	Oracle $\text{EV}(t, x)$	Oracle $\text{BLINDEV}(t, \text{req})$
$\text{RandFn} \leftarrow \text{Fn.Gen}(\text{Fn.In}, \text{Fn.Out})$	$y_1 \leftarrow \text{Fn.Ev}^{\text{P}}(sk, x)$	$q_t \leftarrow q_t + 1$
$st_{\text{P}} \leftarrow \text{P.Init}(\lambda)$	$y_0 \leftarrow \text{RandFn}(t, x)$	$\text{rep}_1 \leftarrow \text{Fn.BlindEv}^{\text{P}}(sk, t, \text{req})$
$pp \leftarrow \text{Fn.Setup}(\lambda)$	Return $y_b$	$(\text{rep}_0, st_{\text{S}}) \leftarrow \text{S.BlindEv}^{\text{LIMEV}}(t, \text{req} : st_{\text{S}})$
$(sk, pk_1) \leftarrow \text{Fn.KeyGen}_{pp}^{\text{P}}()$	Oracle $\text{LIMEV}(t, x)$	Return $\text{rep}_b$
$(st_{\text{S}}, pk_0) \leftarrow \text{S.Init}(pp)$	$q_{t,s} \leftarrow q_{t,s} + 1$	Oracle $\text{PRIM}(x)$
$b' \leftarrow \mathcal{A}^{\text{EV, BLINDEV, PRIM}}(pp, pk_b)$	If $q_{t,s} \leq q_t$ then	$y_1 \leftarrow \text{P.Eval}(x : st_{\text{P}})$
Return $b'$	Return $\text{EV}(t, x)$	$(y_0, st_{\text{S}}) \leftarrow \text{S.Eval}^{\text{LIMEV}}(x : st_{\text{S}})$
	Return $\perp$	Return $y_b$

Fig. 3: Simulation-based security definition for pseudorandomness against malicious clients, with granular accounting for metadata in queries. The LIMEVAL oracle limits the number of evaluations the simulator can make on a per-metadata tag basis.

We let the advantage of a POPRF adversary  $\mathcal{A}$  be defined by

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

where the probability spaces are taken over the random choices made in the games and the events signify that the game outputs the value one.

One could relax our definition in various ways. For example, by setting a parameter  $q_{t,\max}$  that upper bounds the total number of BLINDEV queries on tag  $t$  over the course of the game and letting the simulator — at any point in the game — obtain  $q_{t,\max}$  full evaluations. This would seem to still provide qualitatively the same level of security, but our schemes meet the stronger notion that restricts the simulator over the course of the game. Another relaxation that does not preserve the same level of security would be to allow the simulator more queries than  $q_{t,\max}$ , for example,  $2 \cdot q_{t,\max}$ . But this degrades the security guarantee as it means that in  $q$  queries to BLINDEV on some  $t$  a malicious client can potentially compute up to  $2q$  POPRF outputs for that tag  $t$ .

**Request privacy and unlinkability.** Our second goal is to capture privacy for clients. This means not only that requests should hide the private message portion  $x$ , but also that request/response transcripts and output POPRF values should be unlinkable. We formalize two models for this goal, corresponding to the level of maliciousness by a misbehaving server.

Game POPRIV1 (Figure 4, top game) captures an indistinguishability experiment in which the adversary can query to obtain full transcripts (including output) resulting from honest blinded evaluation of a POPRF. The transcripts are either returned properly ( $b = 0$ ) or with the request, response pairs swapped relative to the outputs ( $b = 1$ ). Intuitively, if the adversary cannot distinguish between these two worlds, then there is no way to link a POPRF output value to a particular blinded evaluation, despite the adversary knowing the secret POPRF key. This captures also message privacy security: if a request reveals some information about the input  $x$  this can be used to win the POPRIV1 game. We sometimes refer to this as request privacy against passive adversaries, because the adversary cannot interfere with the server’s proper execution.

<p>Game <math>\text{POPRIV1}_{\text{Fn,S,P}}^{\mathcal{A},b}(\lambda)</math></p> <p><math>pp \leftarrow \\$ \text{Fn.Setup}(\lambda)</math>  <math>(pk, sk) \leftarrow \\$ \text{Fn.KeyGen}(pp)</math>  <math>st_P \leftarrow \\$ \text{P.Init}(\lambda)</math>  <math>b' \leftarrow \\$ \mathcal{A}^{\text{PRIM,TRANS}}(pp, pk, sk)</math>  Return <math>b'</math></p>	<p>Oracle <math>\text{TRANS}(t, x_0, x_1)</math></p> <p><math>(st_0, req_0) \leftarrow \text{Fn.Req}^P(pk, t, x_0)</math>  <math>(st_1, req_1) \leftarrow \text{Fn.Req}^P(pk, t, x_1)</math>  <math>rep_0 \leftarrow \text{Fn.BlindEv}^P(req_0)</math>  <math>rep_1 \leftarrow \text{Fn.BlindEv}^P(req_1)</math>  <math>y_0 \leftarrow \text{Fn.Finalize}^P(st_0, rep)</math>  <math>y_1 \leftarrow \text{Fn.Finalize}^P(st_1, rep')</math>  <math>\tau \leftarrow (req_b, rep_b, y_0)</math>  <math>\tau' \leftarrow (req_{1-b}, rep_{1-b}, y_1)</math>  Return <math>(\tau, \tau')</math></p>	<p>Oracle <math>\text{PRIM}(x)</math></p> <p><math>y \leftarrow \\$ \text{P.Eval}(x : st_P)</math>  Return <math>y</math></p>
<p>Game <math>\text{POPRIV2}_{\text{Fn,S,P}}^{\mathcal{A},b}(\lambda)</math></p> <p><math>pp \leftarrow \\$ \text{Fn.Setup}(\lambda)</math>  <math>st_P \leftarrow \\$ \text{P.Init}(\lambda)</math>  <math>i \leftarrow 0</math>  <math>b' \leftarrow \\$ \mathcal{A}^{\text{PRIM,REQ,FIN}}(pp)</math>  Return <math>b'</math></p>	<p>Oracle <math>\text{REQ}(pk, t, x_0, x_1)</math></p> <p><math>i \leftarrow i + 1</math>  <math>(st_{i,0}, req_0) \leftarrow \\$ \text{Fn.Req}^P(pk, t, x_0)</math>  <math>(st_{i,1}, req_1) \leftarrow \\$ \text{Fn.Req}^P(pk, t, x_1)</math>  Return <math>(req_b, req_{1-b})</math></p> <p>Oracle <math>\text{FIN}(j, rep, rep')</math></p> <p>If <math>j &gt; i</math> then Return <math>\perp</math>  <math>y_b \leftarrow \text{Fn.Finalize}^P(st_{j,b}, rep)</math>  <math>y_{1-b} \leftarrow \text{Fn.Finalize}^P(st_{j,1-b}, rep')</math>  If <math>y_0 = \perp</math> or <math>y_1 = \perp</math> then  Return <math>\perp</math>  Return <math>(y_0, y_1)</math></p>	<p>Oracle <math>\text{PRIM}(x)</math></p> <p><math>y \leftarrow \\$ \text{P.Eval}(x : st_P)</math>  Return <math>y</math></p>

Fig. 4: Security definitions for honest-but-curious server unlinkability (**top**) and malicious server unlinkability (**bottom**).

The advantage of a POPRIV1 adversary  $\mathcal{A}$  in the P-model is defined by

$$\text{Adv}_{\text{Fn,P},\mathcal{A}}^{\text{po-priv1}}(\lambda) = \left| \Pr \left[ \text{POPRIV1}_{\text{Fn,P}}^{\mathcal{A},1}(\lambda) \Rightarrow 1 \right] - \Pr \left[ \text{POPRIV1}_{\text{Fn,P}}^{\mathcal{A},0}(\lambda) \Rightarrow 1 \right] \right|$$

where the probability spaces are taken over the random choices made in the games and the events signify that the game outputs the value one. We say a Fn scheme is *perfectly private* if  $\text{Adv}_{\text{Fn,P},\mathcal{A}}^{\text{po-priv1}}(\lambda) = 0$  for all adversaries  $\mathcal{A}$ .

POPRIV1 security does not capture malicious servers that deviate from the protocol. So, for example, it doesn't rule out attacks in which the server replies with garbage to a blinded evaluation request.

Our next game POPRIV2 allows the adversary to choose the public keys used for request generation and leaves to the adversary how to reply to requests. The game therefore splits transcript generation across two oracles, a request oracle (REQ) and finalize oracle (FIN). The first oracle replies with a randomly ordered pair of request messages based on the challenge bit, and the second oracle can be queried with adversarially chosen response messages. The game requires that neither  $y_0$  nor  $y_1$  is equal to  $\perp$  — if either is then the finalize oracle returns  $\perp$ . This prevents the trivial attack of corrupting one reply but not the other.

The advantage of a POPRIV2 adversary  $\mathcal{A}$  in the P-model is defined by

$$\text{Adv}_{\text{Fn,P},\mathcal{A}}^{\text{po-priv2}}(\lambda) = \left| \Pr \left[ \text{POPRIV2}_{\text{Fn,P}}^{\mathcal{A},1}(\lambda) \Rightarrow 1 \right] - \Pr \left[ \text{POPRIV2}_{\text{Fn,P}}^{\mathcal{A},0}(\lambda) \Rightarrow 1 \right] \right|$$

where the probability spaces are taken over the random choices made in the games and the events signify that the game outputs the value one.

POPRIV2 is strictly stronger than POPRIV1. Looking ahead our new POPRF meets POPRIV1 when verification is omitted, and POPRIV2 when verification is required.

**Relation to partially blind signatures.** POPRFs are related to partially blind signatures, which were introduced by Abe and Fujisaki [AF96]. A partially blind signature is a tuple of algorithms

$$DS = (DS.Setup, DS.KeyGen, DS.Sign, DS.Ver, DS.Req, DS.BlindEv, DS.Finalize)$$

where the first four algorithms define a standard digital signature scheme for message space consisting of pairs  $(t, x)$ , called the public input (or tag) and private input, respectively. Signatures can also be generated via an interactive protocol which, like we did for POPRFs, we formalize simply as a single round trip protocol initiated by a client running  $DS.Req(pk, t, x)$  to generate a request message  $req$  and client state  $st$ , sending the former to the server which runs  $DS.BlindEv(sk, req)$  to generate and send a response  $rep$  back to the client, which then computes a signature via  $DS.Finalize(st, rep)$ . This protocol should achieve message privacy, which can be defined similarly to our definition above for POPRFs.

The security property targeted is one-more unforgeability, which, roughly speaking, states that an adversarial client can't generate  $q + 1$  unique triples  $(t_1, x_1, \sigma_1), \dots, (t_{q+1}, x_{q+1}, \sigma_{q+1})$  that all verify under a public key  $pk$  even when given the ability to query  $q$  times an oracle that returns the output of  $DS.BlindEv(sk, req)$  for the associated secret key  $sk$  and request message of the adversary's choosing.

A partially blind signature is unique if  $DS.Sign$  is deterministic and its output on  $(pk, t, x)$  matches that of the interactive protocol when initiated on the same triple. A blind signature is just a partially blind signature with  $t$  omitted. JKK observed that one can transform unique blind signatures into OPRFs by hashing the signature. (As far as we are aware there has been no formal treatment of this observation.) A similar transform should work to build a POPRF from a unique partially blind signature.

Most prior partially blind signature schemes are not unique, e.g., [AF96, AO00]. The only unique scheme we are aware of is due to Zhang, Safavi-Naini, and Susilo (ZSS) [ZSS03], but it relies on bilinear pairings and so this generic transformation will not achieve our goals for a POPRF. Moreover as mentioned in the introduction, the security analysis in ZSS is wrong (see Section 4). That said, our construction shares much of the underlying structure from the ZSS one.

## 4 The 3HashSDHI POPRF

We now turn to our main result: providing a new POPRF. Our construction combines elements of the 2HashDH construction with a technique used by Dodis and Yampolskiy for their verifiable PRF; it is also related to a partially blind

signature scheme suggested by Zhang, Safavi-Naini, and Susilo. We call our construction 3HashSDHI, which we often abbreviate to 3H. The name refers to its use of three hashes and reliance on the strong inverse Diffie-Hellman assumption.

**Algorithms.** Our protocol relies on a group  $\mathbb{G}$  of prime order  $p$  and with generator  $g$ . As mentioned in the introduction, the 3HashSDHI protocol computes a PRF output as

$$3\text{H.Ev}(sk, t, x) = \text{H}_2\left(t, x, \text{H}_1(x)^{1/(sk+\text{H}_3(t))}\right)$$

where  $\text{H}_1: \{0, 1\}^* \rightarrow \mathbb{G}$ ,  $\text{H}_2: \{0, 1\}^* \rightarrow \{0, 1\}^{\gamma_2}$ , and  $\text{H}_3: \{0, 1\}^* \rightarrow \{0, 1\}^{\gamma_3}$  are the titular hash functions. Note that  $\text{H}_1$  has range the group  $\mathbb{G}$ , whereas the second and third hashes output bit strings of length  $\gamma_2$  and  $\gamma_3$ . By default we set  $\gamma_2 = \lambda$  and  $\gamma_3 = 2\lambda$ . The third hash must be collision resistant for security to hold. Looking ahead to the security analysis, we will model the hash functions as random oracles. The setup, key generation, and full evaluation algorithms are shown in pseudocode below.

$3\text{H.Setup}(\lambda)$	$3\text{H.KeyGen}^{\text{H}_1 \times \text{H}_2 \times \text{H}_3}(pp)$	$3\text{H.Ev}^{\text{H}_1 \times \text{H}_2 \times \text{H}_3}(sk, t, x)$
$(p, g, \mathbb{G}) \leftarrow \text{GGen}(\lambda)$	$(p, g, \mathbb{G}) \leftarrow pp$	$Y \leftarrow \text{H}_1(x)^{1/(sk+\text{H}_3(t))}$
$pp \leftarrow (p, g, \mathbb{G})$	$sk \leftarrow \mathbb{Z}_p$ ; $pk \leftarrow g^{sk}$	$Z \leftarrow \text{H}_2(t, x, Y)$
Return $pp$	Return $(pk, sk)$	Return $Z$

Here,  $\text{GGen}$  denotes a group parameter generator outputting a triple  $(p, g, \mathbb{G})$  consisting of a prime  $p$ , (the description of) a group  $\mathbb{G}$  of order  $p$ , and a generator  $g$  of  $\mathbb{G}$ .

The blind evaluation protocol has a client compute  $\text{H}_1(x)$  and mask the resulting group element by raising it to a random scalar  $r$ . The client can send the resulting blinded value  $B$  to the server, who can then raise  $B$  to  $1/(sk+\text{H}_3(t))$  and return the result. The client then finalizes by raising the returned value to  $1/r$  in order to remove the blinding, followed by the final step of computing the final hash  $\text{H}_2$ . The blinding ensures request privacy.

We optionally can extend this blinded evaluation protocol to include a proof that the server properly exponentiated  $B$ . This is necessary to have the protocol enjoy POPRIV2 security, which is important in some (but not all) applications. At first, it may not be obvious how to prove to the client that the server is returning  $B' = B^{1/(sk+\text{H}_3(t))}$  relative to the public key  $g^{sk}$ , because the sum appears in the denominator. However, we can use the following trick: the server generates a standard DL proof that  $B = (B')^k$  for some  $k$ . The client runs verification by explicitly reconstructing  $g^k = pk \cdot g^{\text{H}_3(t)} = g^{sk+\text{H}_3(t)}$ . This means that the verification procedure checks the special structure of the exponent  $k$ .

The full protocol, including the NIZK (which uses its own hash  $\text{H}_4$ ), is shown in Figure 5. Here we show that  $t$  is sent from the client to server, though in some applications the server may receive  $t$  out-of-band. Execution requires just one round trip. It requires just two group exponentiations on the client side (and, when using the NIZK, those used for its verification). The server uses one exponentiation plus one for the NIZK proof.

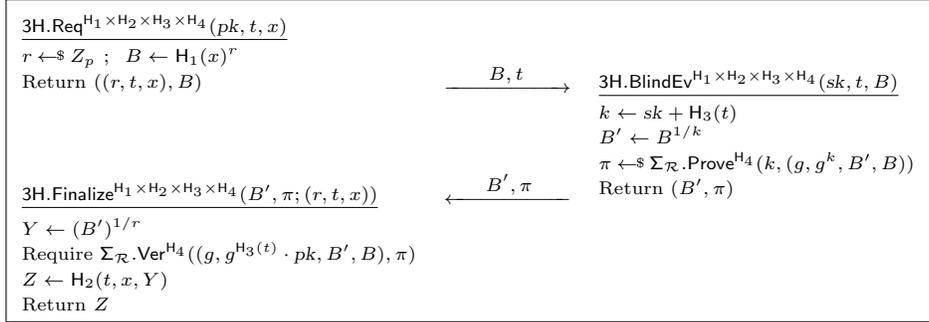


Fig. 5: Blind evaluation for our 3H POPRF construction. All three algorithms have implicit input the parameters  $pp = (p, g, \mathbb{G})$  that describe the group used. The NIZK uses relation  $\mathcal{R} = \{(g, U, V, W), (\alpha) : U = g^\alpha \wedge W = V^\alpha\}$ .

**Relation to partially blind signatures.** 3HashSDHI is closely related to a partially blind signature suggested by Zhang, Safavi-Naini, and Susilo [ZSS03]. It uses groups  $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$  each of order  $p$ , with generators  $g_1, g_2, g_T$ , and that come equipped with an efficient-to-compute pairing  $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$  such that for any  $\alpha, \beta \in \mathbb{Z}_p$  it holds that  $e(g_1^\alpha, g_2^\beta) = g_T^{\alpha\beta}$ . Then their signature is defined as

$$\text{ZSS.Ev}(sk, t, x) = \mathbb{H}_{\mathbb{G}_2}(t, x)^{1/(sk + \mathbb{H}_3(t))}$$

where  $\mathbb{H}_{\mathbb{G}_2}: \{0, 1\}^* \rightarrow \mathbb{G}_2$  hashes onto the group  $\mathbb{G}_2$  and  $\mathbb{H}_3$  is as defined above for 3HashSDHI. As can be seen, 3HashSDHI uses essentially the same structure, combined with a final hash but: (1) we dispense with the use of bilinear pairings using instead NIZKs to provide verifiability; and (2) our blinding protocol is simpler (ZSS's use of bilinear pairings means blinding via  $\mathbb{H}_{\mathbb{G}_2}(t, x)^r$  which allows linking attacks). We also comment that the security analysis of the partially blind signature scheme in [ZSS03] is incorrect. We believe a proof should be enabled rather easily using our techniques, in particular by using a variant of the new gap assumption discussed in the next section.

**Comparison to prior (O)PRFs.** Recall that the 2HashDH OPRF is defined by  $\text{2HashDH.Ev}(sk, x) = \mathbb{H}_2(x, \mathbb{H}_1(x)^{sk})$ . On the other hand, the DY PRF is evaluated on a message  $t$  via  $\text{DY.Ev}(sk, t) = h^{1/(sk+t)}$  for generator  $h$ . Thus, our 3HashSDHI can be seen as blending of the two approaches, basically defining, for each  $x$ , a separate instance of the DY PRF with generator  $h = \mathbb{H}_1(x)$  and input message  $t$ . The way we combine them retains the simple blinding mechanism of 2HashDH to allow hiding  $x$ . Despite the similarity to the prior constructions, analyzing security requires new techniques (see the next section).

Pythia [ECS<sup>+</sup>15] provided the first POPRF. It uses pairing-friendly groups  $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$  each of order  $p$ , with generators  $g_1, g_2, g_T$ . Then, the Pythia POPRF is defined by

$$\text{Pythia.Ev}(sk, t, x) = e(\mathbb{H}_{\mathbb{G}_1}(t), \mathbb{H}_{\mathbb{G}_2}(x))^{sk}$$

where  $H_{\mathbb{G}_2}$  and  $H_{\mathbb{G}_1}$  are hash functions that map to the groups  $\mathbb{G}_1, \mathbb{G}_2$ . This construction enables blinded evaluation by sending  $B \leftarrow H_{\mathbb{G}_2}(x)^r$ , and having the server respond with  $e(H_{\mathbb{G}_1}(t)^{sk}, B)$ . It also has some other features that were desirable in the password hardening context for which Pythia was designed, specifically, that one can have compact key rotation tokens of the form  $\Delta = sk'/sk$ . The token  $\Delta$  can be shared with a client to help it update previously computed POPRF values for any  $t, x$ .

The 3HashSDHI construction does not support compact key rotations even if one omits the final  $H_2$  evaluation — one would need a token for each value  $t$ . This may make key rotations less efficient in password hardening settings, but in other applications like those we explore in Section 7 we do not need compact key rotation tokens. At the same time, 3HashSDHI avoids use of pairings. This makes it significantly faster to compute and saves bandwidth.

We also compare to the recent attribute-based verifiable OPRF (AB-VOPRF) suggested by Huang et al. of Facebook [HJK<sup>+</sup>21] for use with Privacy Pass. An attribute-based VOPRF is a POPRF that separates out an explicit algorithm for converting a secret key and attribute  $t$  (what we call a tag) into a tag-specific public key, secret key pair. As with other VOPRFs, there is a verifiable, blinded evaluation protocol by which a client can obtain an output on some  $(t, x)$  pair without revealing  $x$ . Any AB-VOPRF gives a POPRF, and vice versa.

The Facebook construction of an AB-VOPRF combines the 2HashDH approach with the Naor-Reingold PRF [NR97]. Evaluation is defined by

$$\text{FB.Ev}(sk, t, x) = H_1(x)^{a_0 \cdot \prod_i a_i^{t[i]}}$$

where  $sk = a_0, a_1, \dots, a_{|t|}$  and  $t[i]$  indicates the  $i^{\text{th}}$  bit of  $t$ . To make this verifiable, the scheme must provide a more complex NIZK involving  $|t|$  group elements, making it expensive to transmit and verify, particularly in applications where a wide variety of tags  $t$  will be used. In comparison 3HashSDHI is as efficient as 2HashDH.

Finally, a concurrent, independent work by Silde and Strand [SS21] describe what we call the 3HashSDHI protocol and how it could be useful for Privacy Pass and the Facebook de-identified logging application. They formalize a notion of anonymous token security that is more tailored to Privacy Pass style applications (compared to our general POPRF definitions), but this definition contains the aforementioned problem (see Section 3) of not performing query accounting on a per public input basis, making it too weak of a security notion for their applications. In addition, the security analysis relative to this notion is incomplete, and so the paper does not yet provide a proof even of this weaker security notion. Nevertheless, their work underscores the benefits of the 3HashSDHI protocol in the applications they explore and our proof techniques (in particular, the new one-more gap SDHI assumption discussed in the next section) should enable improvements to their analysis.

## 5 Security Analysis

We show formally that the 3HashSDHI PO-PRF enjoys pseudorandomness and request privacy. The former is the more complex analysis; we start with it.

### 5.1 Pseudorandomness

The main technical challenge is showing that 3HashSDHI meets our pseudorandomness definition, captured by game POPRF (Figure 3 in Section 3). We start with an overview of our proof strategy, and then state our main result.

**Proof strategy.** Our proof of pseudorandomness proceeds in several steps. First, we introduce a new discrete log (DL) type cryptographic hardness assumption: the one-more gap strong Diffie-Hellman inversion problem, denoted  $(m, n)$ -OM-GAP-SDHI for parameters  $m, n$  that we will explain. The new assumption is a generalization of prior one-more DL assumptions, but extended with two oracles and a more involved one-more winning condition which depends on the number of queries with a specific form to one of the oracles. We show that we can build a POPRF simulator such that, in the ROM, distinguishing between the real ( $b = 1$ ) and ideal worlds ( $b = 0$ ) reduces to breaking an instance of  $(m, n)$ -OM-GAP-SDHI where  $m$  is the number of  $H_3$  queries made by  $\mathcal{A}$  and  $n$  is the number of  $H_1$  queries.

We also analyze the security of our new assumption, showing that, in the Algebraic Group Model (AGM) [FKL18] it reduces to one of the uber assumptions from Bauer, Fuchsbauer, and Loss (BFL) [BFL20]. In turn, we can use a result from BFL to finally show that our new assumption is implied (again, in the AGM) by the  $q$ -DL assumption. This provides good evidence of the difficulty of the problem, and allows us to derive precise concrete security bounds.

**The one-more gap SDHI assumption.** Game  $(m, n)$ -OM-GAP-SDHI is shown in Figure 7. The game generates a group instance and a challenge secret  $sk$ . The adversary  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  runs in two stages. In the first stage it receives the group description  $p, \mathbb{G}$  and outputs a sequence of  $n$  scalar values  $c_1, \dots, c_n$ . Importantly  $\mathcal{A}_1$  does not receive  $g$ , forcing it to commit to the  $c_i$  values in a way independently of the generator  $g$ . We assume that  $g$  is randomly chosen; this will be important in our analysis. Then, the second stage  $\mathcal{A}_2$  is run on input the generator  $g$ ,  $g^{sk}$ , and a vector of  $m$  group elements  $g^{y_1}, \dots, g^{y_n}$ . The adversary is given access to two oracles. The SDH oracle returns  $B^{1/(sk+c_i)}$  for arbitrary  $B$  and one of the previously specified  $c_i$  values. The SDDH oracle is a decision oracle that helps the adversary determine whether  $Z = Y^{1/(sk+c_i)}$  for arbitrary  $Y, Z$  and one of the previously specified  $c_i$  values.

The adversary outputs a distinguished index  $\gamma$  indicating a  $c_\gamma$  value, as well as a set of  $\ell$  pairs  $(Z_i, \alpha_i) \in \mathbb{G} \times [0..m]$ . The adversary wins if  $\ell > q_\gamma$  and  $Z_i = Y_{\alpha_i}^{1/(sk+c_\gamma)}$  for all  $1 \leq i \leq \ell$ . Here  $q_\gamma$  is the number of queries to the SDH with second input set to  $\gamma$ . Without the “one more” restriction of  $\ell > q_\gamma$ , it is trivial

Game $(m, n)$ -OM-GAP-SDHI $_{\text{GGen}}^{\mathcal{A}}(\lambda)$	Oracle SDH $(B, i)$
$(p, g, \mathbb{G}) \leftarrow_{\mathcal{S}} \text{GGen}(\lambda)$	If $i \notin [1, n]$ then
$sk \leftarrow_{\mathcal{S}} \mathbb{Z}_p$ ; $[y_i]_i^m \leftarrow_{\mathcal{S}} [\mathbb{Z}_p]_i^m$	Return $\perp$
$(st_{\mathcal{A}}, [c_i]_i^n) \leftarrow_{\mathcal{S}} \mathcal{A}_1(p, \mathbb{G})$	$q_i \leftarrow q_i + 1$
Require $\forall_{i \neq j} c_i \neq c_j$	$Z \leftarrow B^{1/(sk+c_i)}$
$(\gamma, [Z_i, \alpha_i]_i^\ell) \leftarrow_{\mathcal{S}} \mathcal{A}_2^{\text{SDH, SDDH}}(g, g^{sk}, [g^{y_i}]_i^m : st_{\mathcal{A}})$	Return $Z$
Require $q_\gamma < \ell \wedge \forall_{i \neq j} \alpha_i \neq \alpha_j$	Oracle SDDH $(Y, Z, i)$
Return $[Z_i]_i^\ell = [g^{y\alpha_i/(sk+c_\gamma)}]_i^\ell$	Return $Z = Y^{1/(sk+c_i)}$

Fig. 7: The one-more gap strong Diffie-Hellman inversion security game.

to win. We define the  $(m, n)$ -OM-GAP-SDHI-advantage of an adversary  $\mathcal{A}$  by

$$\text{Adv}_{\text{GGen}, \mathcal{A}}^{(m, n)\text{-om-gap-sdhi}}(\lambda) = \Pr \left[ (m, n)\text{-OM-GAP-SDHI}_{\text{GGen}}^{\mathcal{A}}(\lambda) \Rightarrow \text{true} \right].$$

An adversary  $\mathcal{A}$  has query budget  $(\vec{q}, q_{\text{SDDH}})$  for  $\vec{q} = [\vec{q}_1, \dots, \vec{q}_n]$  if at the end of the game  $\mathcal{A}$  has made at most  $\vec{q}_i$  queries to SDH with index  $i$  and has made at most  $q_{\text{SDDH}}$  queries to SDDH.

We note that a weakening of the assumption dispenses with the more granular per- $c$ -value accounting, instead just asking that the adversary can't come up with  $\ell > q$  solutions for any mixture of  $Y_i$  and  $c_j$  values. This variant is much easier to analyze in the AGM, but is not sufficient for our analysis.

**Reducing  $(m, n)$ -OM-GAP-SDHI to  $q$ -DL.** In two steps, we show how to reduce this assumption, in the AGM, to the difficulty of  $q$ -DL. The latter involves a game  $q$ -DL (Figure 6) that generates a group instance  $p, g, \mathbb{G}$  for security parameter  $\lambda$ , and gives an adversary  $g, g^x, g^{x^2}, \dots, g^{x^q}$  for a random scalar  $x$ . The adversary must output  $x$ . We define the advantage of a  $q$ -DL-adversary  $\mathcal{A}$  to be  $\text{Adv}_{\text{GGen}, \mathcal{A}}^{q\text{-dl}}(\lambda) = \Pr [q\text{-DL}_{\text{GGen}}^{\mathcal{A}}(\lambda) \Rightarrow \text{true}]$ .

As a convenient middle layer, we rely on BFL's "Uber-assumption" [BFL20], formalized via the game  $m$ -UBER in Figure 8. It involves a game where the adversary can obtain  $g^{\rho(\vec{x})}$  by querying an arbitrarily chosen  $m$ -variate polynomial  $\rho(\vec{X})$  to an oracle EV, for a secret vector  $\vec{x} \leftarrow_{\mathcal{S}} \mathbb{Z}_p^m$ . The adversary wins if it outputs successfully  $g^{\mu(\vec{x})}$  for some polynomial  $\mu(\vec{X})$  which is *independent* of the polynomials  $\rho_1(\vec{X}), \dots, \rho_q(\vec{X})$  queried to EV, i.e.,  $\mu(\vec{X})$  cannot be expressed as an affine combination  $\mu(\vec{X}) = \alpha_1 \rho_1(\vec{X}) + \dots + \alpha_q \rho_q(\vec{X}) + \beta$ . The adversary can also query an additional DECIDE oracle with a polynomial  $\rho(\vec{X})$ , as well as group

Game $q$ -DL $_{\text{GGen}}^{\mathcal{A}}(\lambda)$
$(p, g, \mathbb{G}) \leftarrow_{\mathcal{S}} \text{GGen}(\lambda)$
$x \leftarrow_{\mathcal{S}} \mathbb{Z}_p$
$x' \leftarrow_{\mathcal{S}} \mathcal{A}(p, \mathbb{G}, g, [g^{x^i}]_{i=1}^q)$
Return $x = x'$

Fig. 6: The  $q$ -type discrete log security game.

Game $m\text{-UBER}_{\mathbb{G}\text{Gen}}^{\mathcal{A}}(\lambda)$	Oracle $\text{Ev}(\rho(\vec{X}))$
$(p, g, \mathbb{G}) \leftarrow \$ \mathbb{G}\text{Gen}(\lambda)$	$Q \leftarrow Q \cup \{\rho(\vec{X})\}$
$Q \leftarrow \{\}$	Return $g^{\rho(\vec{x})}$
$\vec{x} = [x_i]_i^n \leftarrow \$ [\mathbb{Z}_p]_i^n$	
$(U, \mu(\vec{X})) \leftarrow \$ \mathcal{A}^{\text{Ev, DECIDE}}(p, \mathbb{G}, g)$	Oracle $\text{DECIDE}(\rho(\vec{X}), [Y_i]_i^n)$
Return $(U = g^{\mu(\vec{x})} \wedge Q \perp\!\!\!\perp \{\mu(\vec{X})\})$	$\vec{y} = [y_i]_i^n \leftarrow [\log_g Y_i]_i^n$
	Return $\rho(\vec{y}) \equiv_p 0$

Fig. 8: The interactive, flexible-output, polynomial uber assumption with decision oracle. Here,  $\perp\!\!\!\perp$  denotes algebraic independence.

elements  $g^{y_1}, \dots, g^{y_m}$ , and learn whether  $g^{\rho(y_1, \dots, y_m)} = 0$  or not. We denote the corresponding advantage as  $\text{Adv}_{\mathbb{G}\text{Gen}, \mathcal{A}}^{m\text{-uber}}(\lambda) = \Pr [m\text{-UBER}_{\mathbb{G}\text{Gen}}^{\mathcal{A}}(\lambda) \Rightarrow \text{true}]$ .

We are going to prove the following theorem in Appendix A. Here and subsequently we use ‘ $\approx$ ’ to denote that runtimes are equal up to small constant factors.

**Theorem 1.** *For any algebraic adversary  $\mathcal{A}_{\text{sdhi}}$  of  $(m, n)$ -OM-GAP-SDHI with query budget  $(\vec{q} = [q_1, \dots, q_n], q_{\text{SDDH}})$ , and any  $\mathbb{G}\text{Gen}$  outputting  $(p, g, \mathbb{G})$ , where  $g$  is a uniformly chosen element of  $\mathbb{G}$ , we give adversary  $\mathcal{A}_{\text{uber}}$  such that*

$$\text{Adv}_{\mathbb{G}\text{Gen}, \mathcal{A}_{\text{sdhi}}}^{(m, n)\text{-om-gap-sdhi}}(\lambda) \leq (q_{\max} + 1) \cdot \text{Adv}_{\mathbb{G}\text{Gen}, \mathcal{A}_{\text{uber}}}^{(m+1)\text{-uber}}(\lambda) + \frac{q}{p},$$

where  $q = \sum_i^n q_i$  and  $q_{\max} = \max\{q_i\}_i^n$ . Also,  $\mathcal{A}_{\text{uber}}$  makes at most  $q$  queries to its polynomial evaluation oracle with maximum degree  $q + 1$ , and outputs a polynomial of degree at most  $q$ . Further,  $T(\mathcal{A}_{\text{sdhi}}) \approx T(\mathcal{A}_{\text{uber}})$ .

It is important here to note that the theorem assumes that the query budgets  $q_i$  corresponding to different  $i$ ’s are *fixed* a priori, rather than being chosen adaptively.

Combined with a basic reduction from [BFL20], this gives us the following immediate corollary.

**Corollary 1.** *For any algebraic adversary  $\mathcal{A}_{\text{sdhi}}$  of  $(m, n)$ -OM-GAP-SDHI, with query budget  $(\vec{q} = [q_1, \dots, q_n], q_{\text{SDDH}})$ , and any  $\mathbb{G}\text{Gen}$  outputting  $(p, g, \mathbb{G})$ , where  $g$  is a uniformly chosen element of  $\mathbb{G}$ , we give adversary  $\mathcal{A}_{\text{dl}}$  such that*

$$\text{Adv}_{\mathbb{G}\text{Gen}, \mathcal{A}_{\text{sdhi}}}^{(m, n)\text{-om-gap-sdhi}}(\lambda) \leq (q_{\max} + 1) \cdot \text{Adv}_{\mathbb{G}\text{Gen}, \mathcal{A}_{\text{dl}}}^{(q+1)\text{-dl}}(\lambda) + \frac{q}{p},$$

where  $q = \sum_i^n q_i$  and  $q_{\max} = \max\{q_i\}_i^n$ . Further,  $T(\mathcal{A}_{\text{sdhi}}) \approx T(\mathcal{A}_{\text{dl}})$ .

The main difficulty of the proof of Theorem 1 in Appendix A stems from the one-more requirement  $\ell > q_\gamma$  in the winning condition, which is defined in a way that depends on the specific number of queries  $q_\gamma$  to  $\text{SDH}(\cdot, \gamma)$ . To gain some intuition, it is convenient to think of the game in algebraic terms (and this point of view is also accurate when casting our proof in the AGM).<sup>5</sup> Specifically, let

<sup>5</sup> We ignore the SDDH oracle in this discussion, and it will be easy to handle in the actual proof via the DECIDE oracle.

we describe exponents of the elements provided to  $\mathcal{A}_2$  as formal polynomials  $X_0$  (standing for the secret key) and  $X_1, \dots, X_m$  (for the values  $y_1, \dots, y_m$ ). Initially, the adversary has these polynomials available, and now a call to  $\text{SDH}(P, i)$  can also be thought of as dividing some polynomial  $P$  (or more generally, a rational function) by  $(X_0 + c_i)$ . The rational function  $P$  can be any affine combination of the functions obtained so far, and  $\text{SDH}(P, i)$  adds a new rational function to this set of available rational functions. In other words, consecutive queries induce a *transcript*  $\tau$  consisting of the initial functions  $X_0, X_1, \dots, X_m$ , and the functions returned by  $\text{SDH}$ . The goal of the adversary is to ensure that, for some  $\gamma$ , the span<sup>6</sup> of  $\tau$  contains  $\ell > q_\gamma$  functions of the form

$$\frac{X_{\alpha_1}}{X_0 + c_\gamma}, \dots, \frac{X_{\alpha_\ell}}{X_0 + c_\gamma}.$$

An adversary cannot achieve this goal naively by querying  $\text{SDH}(X_{\alpha_j}, \gamma)$  for  $j \in [\ell]$  without violating the query budget. Still, the key difficulty here is that the adversary *could*, after learning (say)  $X_1/(X_0 + c_\gamma)$  make a further query that would give  $X_1/(X_0 + c_\gamma)(X_0 + c_{\gamma'})$  for some  $\gamma' \neq \gamma$ . This second query would *not* count towards  $q_\gamma$ , and could potentially be helpful, as it *does* involve  $c_\gamma$ .

The bulk of our proof shows that arbitrary queries to  $\text{SDH}$  cannot, in fact, help the adversary. We do so via a careful inductive analysis which shows that the transcript  $\tau$  can be rewritten in an equivalent way, call it  $\tau'$ , without affecting its span. In particular,  $\tau'$  only involves rational functions whose denominators have form  $(X_0 + c_i)^k$  for some  $i$  and  $k$ , but no products involving multiple  $c_i$ 's appear in the denominators. We leverage this structure to show that the span of such  $\tau'$  can include at most  $q_\gamma$  rational functions of the form  $\frac{X_i}{X_0 + c_\gamma}$ .

Now, given the above algebraic game cannot be won, an adversary winning the game must necessarily produce an output  $(\gamma, [Z_i, \alpha_i]_i^\ell)$  where for at least one  $i \in [\ell]$ , we have that the polynomial  $X_{\alpha_i}/(X_0 + c_i)$  is not in the span of the queries to  $\text{SDH}$ . This lends itself naturally to a reduction to the Uber-assumption, which we describe in full in the proof.

**Reducing to  $(m, n)$ -OM-GAP-SDHI.** We now turn to showing that we can reduce the pseudorandomness security of  $3\text{HashSDHI}$  to our new assumption. We focus on the verifiable version of  $3\text{HashSDHI}$ ; an analysis for the non-verifiable version is easily derived from our analysis here. Our analysis is in the RO model; we model all four hash functions as ROs.

We start by describing the simulator used in the proof. The simulator's goal is to respond to blind evaluation and RO queries so that the resulting transcript of values is indistinguishable from real responses. Importantly, the simulator must do this without making too many calls to the full evaluate oracle for each  $\text{BLINDEVAL}$ -queried public input  $t$ . Intuitively, achieving this security enforces that a malicious client can not exploit the blinded evaluation oracle to do more than help it compute a single  $\text{POPRF}$  output for the particular requested  $t$ .

<sup>6</sup> By "span" we mean the set of rational functions that can be obtain by taking *affine* combinations of the functions in  $\tau$ .

The simulator works as follows. It chooses its own secret key  $sk$  and answers the  $H_1$  and  $H_3$  queries with random group elements and scalars, respectively. To answer a blinded evaluation query, it runs the scheme’s blind evaluation algorithm  $\text{Fn.BlindEv}(sk, B)$ , except that it uses the NIZK’s simulator to generate the proof  $\pi$  (and to simulate any ideal primitive underlying the NIZK, i.e.,  $H_4$ ). The key challenge is in simulating  $H_2$  queries, that which enables the adversary to “complete” a blinded evaluation. The simulator must arrange that the value it returns in response to  $H_2$  queries is consistent with the random value returned by  $\text{EV}$ . To do so, the simulator checks whether a queried point  $(t, x, Y)$  is such that  $Y = H_1(x)^{1/(sk+H_3(t))}$  and, if so, it queries  $\text{LIMEV}(t, x)$  and returns the output. Otherwise, it chooses a random point to return. The simulator can perform this check because it chose  $sk$ . See Figure 10 in Appendix B for the full details of the simulator.

The simulation can fail should the adversary be able to query it on a point  $Y = H_1(x)^{1/(sk+H_3(t))}$  when the simulator cannot make another call to  $\text{LIMEV}$  for that value  $t$ . This can only arise should the adversary query  $H_2$  on more such values  $t, Y$  than queries it so far made to  $\text{BLINDEV}$  on that  $t$ . We show that an adversary, that can do so, can also win the  $(m, n)$ -OM-GAP-SDHI game where  $m, n$  are the total number of queries involving a distinct  $x$  value and distinct  $t$  value, respectively. (We define this more precisely below.) This step also relies on the collision resistance of  $H_3$ , which holds in the ROM.

To formalize this, we state below a theorem using the ideal primitive model in which  $\mathsf{P} = H_1 \times H_2 \times H_3 \times H_4$  for random oracles over  $H_1 : * \rightarrow \mathbb{G}$ ,  $H_2 : * \times * \times \mathbb{G} \rightarrow \{0, 1\}^\lambda$ ,  $H_3 : * \rightarrow \mathbb{Z}_p$ ,  $H_4 : \mathbb{G}^6 \rightarrow \mathbb{Z}_p$  for  $(p, \mathbb{G})$  determined by  $\text{GGen}(\lambda)$ . Here ‘ $*$ ’ denotes the set of arbitrary inputs. We define the query budget for an adversary  $\mathcal{A}_{\text{prf}}$  in the  $\mathsf{P}$  model to be a tuple  $(m, n, q_E, \vec{q}, q_{H_1}, q_{H_2}, q_{H_3}, q_{H_4})$  where:

- $m$  is the maximum number of distinct  $x$  values queried by  $\mathcal{A}_{\text{prf}}$  to  $\text{EV}$ ,  $H_1$ , or  $H_2$ ;
- $n$  is the maximum number of distinct  $t$  values queried by  $\mathcal{A}_{\text{prf}}$  to  $\text{EV}$ ,  $\text{BLINDEV}$ ,  $H_2$ , or  $H_3$ ;
- $\vec{q} = [\vec{q}_1, \dots, \vec{q}_n]$  is a vector where each  $\vec{q}_i$  is the maximum number of queries by  $\mathcal{A}_{\text{prf}}$  to  $\text{BLINDEV}(t_i, \text{req})$  for any  $\text{req}$  and where we  $t_1, \dots, t_n$  are the (at most)  $n$  values  $t_i$  queried in the course of the game in the order of when they are queried. (That is,  $t_1$  is the first  $t$  value queried,  $t_2$  is the second, etc.) In words, the adversary is limited to some number  $n$  of public inputs  $t$  that it can target, and makes a limited number of blinded evaluation queries for each of those inputs  $t$ .
- $q_E, q_{H_1}, q_{H_2}, q_{H_3}$ , and  $q_{H_4}$  are the maximum number of queries made by  $\mathcal{A}_{\text{prf}}$  to the  $\text{EV}$ ,  $H_1$ ,  $H_2$ ,  $H_3$ , and  $H_4$  oracles, respectively.

Note that our query budget requirement  $\vec{q}$  does not restrict *which* values  $t$  the adversary can use; these can be picked adaptively. But the number of times each  $t$  value is queried is restricted by the order in which they are queried. The granular accounting of blinded evaluation queries via  $\vec{q}$  will be important when combining the following theorem with Theorem 1.

**Theorem 2.** *Let  $\mathcal{A}_{\text{prf}}$  be a P-model POPRF adversary against 3H with query budget  $(m, n, q_E, \vec{q}, q_{H_1}, q_{H_2}, q_{H_3}, q_{H_4})$ . Then we give a  $H_4$ -model adversary  $\mathcal{A}_{zk}$  and adversary  $\mathcal{A}_{\text{sdhi}}$  such that*

$$\text{Adv}_{3H, S[S_\Sigma], P, \mathcal{A}_{\text{prf}}}^{\text{po-prf}}(\lambda) \leq \text{Adv}_{\Sigma_{\mathcal{R}}, \mathcal{R}, H_4, S_\Sigma, \mathcal{A}_{zk}}^{\text{zk}}(\lambda) + \text{Adv}_{\text{GGen}, \mathcal{A}_{\text{sdhi}}}^{(m,n)\text{-om-gap-sdhi}}(\lambda) + \frac{n^2}{p},$$

where  $S$  is the simulator defined in Figure 10 that makes use of NIZK simulator  $S_\Sigma$ . Adversary  $\mathcal{A}_{sk}$  makes  $q_{H_4}$  queries to its random oracle and  $\mathcal{A}_{\text{sdhi}}$  has query budget  $(\vec{q}, q_{H_2})$ . Further,  $T(\mathcal{A}_{\text{prf}}) \approx T(\mathcal{A}_{zk}) \approx T(\mathcal{A}_{\text{sdhi}})$ .

A detailed proof is given in Appendix B. It proceeds via a sequence of games, starting with the real world  $\text{POPFRF}_{3H, P, S}^{\mathcal{A}_{\text{po-prf}, 1}}(\lambda)$  and first transitioning to a game that replaces the NIZK  $\pi$  with one generated by the NIZK simulator  $S_\Sigma$ . Then we change how EV queries are handled. Instead of computing the POPRF using  $sk$ , we pick a random value and add it to a table  $R$ . We also modify the handling of  $H_2$  queries to check if  $R$  has been set on a relevant value and, if so, patch up  $H_2$ 's response so that it maintains consistency. This does not change the distribution of responses to the adversary. Finally, we are in position to perform a reduction to  $(q_{H_1}, q_{H_3})$ -OM-GAP-SDHI: the only difference between this game and the ideal world  $\text{POPFRF}_{3H, P, S[S_\Sigma]}^{\mathcal{A}_{\text{prf}, 0}}(\lambda)$  is when  $H_2$  needs to repair a  $R$  value more often than queries to BLINDEV. This reduction step is made relatively simple by our new assumption, which provides the values and oracles necessary to simulate  $\mathcal{A}_{\text{prf}}$ 's view in a straightforward way.

We can combine the two main theorems with a standard result about the NIZK that we use (restated in Section 2) to give the following corollary.

**Corollary 2.** *Let  $\mathcal{A}_{\text{prf}}$  be a P-model POPRF adversary against 3H with query budget  $(m, n, q_E, \vec{q}, q_{H_1}, q_{H_2}, q_{H_3}, q_{H_4})$  and GGen any group parameter generator outputting  $(p, g, \mathbb{G})$ , where  $p$  is a prime  $g$  is a uniformly chosen element of  $\mathbb{G}$ . Then, we give adversary  $\mathcal{A}_{\text{dl}}$  such that*

$$\begin{aligned} \text{Adv}_{3H, S[S_\Sigma], P, \mathcal{A}_{\text{prf}}}^{\text{po-prf}}(\lambda) &\leq (q_{\max} + 1) \cdot \text{Adv}_{\text{GGen}, \mathcal{A}_{\text{dl}}}^{(q+1)\text{-dl}}(\lambda) \\ &\quad + \frac{q + n^2}{p} + \frac{3q^2 + q(q_{H_4} + 4) + 2}{2^\lambda}, \end{aligned}$$

where  $q = \sum_i^n q_i$ ,  $q_{\max} = \max\{q_i\}_i^n$ . Further,  $S$  is the simulator defined in Figure 10 that makes use of NIZK simulator  $S_\Sigma$ , and  $T(\mathcal{A}_{\text{prf}}) \approx T(\mathcal{A}_{\text{dl}})$ .

**Concrete security and parameter selection.** Corollary 2 is interpreted best in the generic-group model (GGM) [Sho97, Mau05], as this yields an absolute bound in terms of  $\mathcal{A}_{\text{prf}}$ 's resources. The advantage of a generic algorithm  $\mathcal{A}_{\text{dl}}$  running in time  $T$  (or more precisely, making  $T$  queries to the generic-group oracle) against  $(q + 1)$ -DL in a group of order  $p$  is  $\text{Adv}_{\text{GGen}, \mathcal{A}_{\text{dl}}}^{(q+1)\text{-dl}}(\lambda) \leq (T + q + 2)^2(q + 1)/(p - 1)$  (see, e.g., [BFL20] for a proof). This advantage is multiplied by  $q_{\max}$  to obtain the dominating term in our final bound. We conjecture however that the bound is somewhat pessimistic, and that the factor  $q_{\max}$

is an artifact of the proof. In fact, as we discuss below, a different interpretation of our proof flow, which is particularly meaningful in the GGM, avoids this factor altogether. This improved bound omitting  $q_{\max}$  is also essentially tight, since Cheon’s attack [Che06] extracts<sup>7</sup> the secret key from  $q$  BLINDEVAL queries in time  $\sqrt{p/q}$ , as long as  $q$  divides  $p - 1$  or  $p + 1$ .

Cheon’s attack can therefore guide parameter selection, as is also done for 2HashDH deployments. For example, a 256-bit group may be sufficient to achieve security for up to  $T = 2^{80}$ , as this would still accommodate up to  $q \approx 2^{96}$  blind evaluations without violating our bounds. In contrast, to ensure security up to  $T = 2^{128}$ , moving to a 384-bit curve appears necessary. The conclusion being that our choice of parameters is consistent with that for 2HashDH, meaning we achieve the same group operation performance while adding public inputs.

We also note that our reduction to the uber assumption requires the generator to be uniformly chosen. We cannot envision any security issues when the generator is instead fixed, and the need for uniformly chosen generators is likely just an artifact of our proof technique.

**Tighter GGM bound.** We only sketch the main idea behind the tighter GGM proof, as it is the result of a minor modification of our AGM proof flow. First, note that the  $(q_{\max} + 1)$  factor in Corollary 2 is inherited from Theorem 1 and is due to our inability to *efficiently* find, within the adversary  $\mathcal{A}_{\text{uber}}$ , a good index  $j \in [\ell]$  that leads to a break of the uber-assumption. Therefore, we are left with guessing. However, an alternative is to find such  $j$  by computing all  $\ell$  possible polynomials  $\mu(\vec{X})$ , and outputting the one which is independent from those input to Ev. Unfortunately, this is computationally expensive, and requires time at least  $\Omega(q_{\max}^2)$ . In other words, we could make the proof tight with respect to advantage while losing tightness with respect to time complexity. While in our proof flow this needs to be taken into account, in the GGM only the number of group operations matters (i.e., the number of oracle calls), whereas “additional” running time is for free. Thus, if  $\mathcal{A}_{\text{prf}}$  makes  $T$  queries to its GGM oracles, our proof flow yields (with the proposed modification) an adversary  $\mathcal{A}_{\text{dl}}$  with roughly the same number of GGM queries and advantage against  $(q + 1)$ -DL.

## 5.2 Request Privacy

We now turn to request privacy, which is simpler to analyze. Intuitively, 3HashSDHI client requests leak no information because the blinding makes them independent of other requests and finalized outputs. The following theorem formalizes this for the case of POPRIV1 for the non-verifiable version of 3HashSDHI.

**Theorem 2.** *For any POPRIV1 adversary  $\mathcal{A}_{\text{po-priv1}}$  against 3H (without client verification) we have that  $\text{Adv}_{3\text{H}, \mathcal{A}_{\text{po-priv1}}}^{\text{po-priv1}}(\lambda) = 0$ .*

<sup>7</sup> Define  $x = (sk + H_3(t))^{-1}$  for some fixed  $t$ . Then, the attacker can just obtain, via consecutive iterative queries, the values  $g^x, g^{x^2}, \dots, g^{x^q}$ , and then recover  $x$  via Cheon’s attack. Finally,  $sk = x^{-1} - H_3(t)$ .

Note that the theorem makes no assumptions about the hash functions or group, instead privacy derives directly from the information-theoretic blinding.

*Proof.* Let  $G$  be the same as game  $\text{POPRIV1}_{3\text{H},\text{P}}^{\mathcal{A},b}(\lambda)$  except that we replace  $(\text{req}_d, \text{rep}_d) = (\text{H}_1(x_d)^{r_d}, \text{H}_1(x_d)^{r_d \cdot sk})$  with  $(\text{req}_d, \text{rep}_d) = (g^{r_d}, g^{r_d \cdot sk})$  for  $d \in \{0, 1\}$  and where  $r_0, r_1$  are the random exponents chosen in the two invocations of  $3\text{H.Req}$ . Observe that in game  $G$  the values returned by  $\text{REQ}$  are independent of the challenge bit  $b$ . Then we have that

$$\Pr \left[ \text{POPRIV1}_{3\text{H},\text{P}}^{\mathcal{A},1}(\lambda) \Rightarrow 1 \right] = \Pr [G \Rightarrow 1] = \Pr \left[ \text{POPRIV1}_{3\text{H},\text{P}}^{\mathcal{A},0}(\lambda) \Rightarrow 1 \right].$$

■

Non-verifiable PO-PRFs, including the non-verifiable version of  $3\text{HashSDHI}$ , cannot achieve our stronger notion of malicious request privacy. The attack is straightforward since the adversary can simply replace one of the two responses with garbage, and determine the challenge bit. In detail for the case of  $3\text{H}$ , adversary  $\mathcal{A}_{\text{po-priv2}}$  can pick  $sk \in \mathbb{G}$  arbitrarily, let  $pk = g^{sk}$ , and then query  $\text{REQ}(pk, t, x_0, x_1)$  for some arbitrary  $t, x_0, x_1$ . It obtains back from the oracle  $\text{req}, \text{req}'$ , and then parses  $\text{req}$  as a pair  $(B, t)$ . It then queries  $\text{FIN}(B^{1/(sk+\text{H}_3(t))}, g)$  to get back reply  $(y_0, y_1)$ . It checks if  $y_0 = \text{H}_2(\text{H}_3(x_0)^{1/(sk+\text{H}_3(t))})$  and returns 0 if so. Otherwise it returns one. This adversary wins with probability 1.

The verifiable version of  $3\text{HashSDHI}$  achieves our stronger notion of malicious request privacy, due to the ZKP forcing the malicious server to respond honestly to blinded requests (relative to the public key being used). The following theorem formalizes this, where we model the hash used by the ZKP as a random oracle, and all other hashes as standard model.

**Theorem 3.** *Let  $\mathcal{A}_{\text{po-priv2}}$  be a  $\text{POPRIV2}$  adversary in the  $\text{H}_4$ -model against  $3\text{H}$  that makes at most  $q$  queries to  $\text{FIN}$ . We give in the proof below a  $\text{SOUND}_{\text{NiZK},\mathcal{R},\text{P}}^{\mathcal{A}}$  adversary  $\mathcal{B}_{\text{sound}}$  such that*

$$\text{Adv}_{3\text{H},\text{H}_4,\mathcal{A}_{\text{po-priv2}}}^{\text{po-priv2}}(\lambda) \leq 4q \cdot \text{Adv}_{\text{NiZK},\mathcal{R},\text{H}_4,\mathcal{B}_{\text{sound}}}^{\text{sound}}(\lambda).$$

Further,  $T(\mathcal{B}_{\text{sound}}) \approx T(\mathcal{A}_{\text{po-priv2}})$ .

*Proof.* Consider game  $\text{POPRIV2}_{3\text{H},\text{P}}^{\mathcal{A}_{\text{po-priv2}},1}(\lambda)$ . We consider the event that, in the course of the game, a  $\text{FIN}(j, (B'_1, \pi_1), (B'_2, \pi_2))$  query is made such that either

1.  $B'_1 \neq B_{j,b}^{1/(sk+\text{H}_3(t_j))}$  but  $\Sigma_{\mathcal{R}}.\text{Ver}^{\text{H}_4}((g, g^{\text{H}_3(t_j)} \cdot pk_j, B'_1, B_{j,b}), \pi_1) = 1$ ; or
2.  $B'_2 \neq B_{j,1-b}^{1/(sk+\text{H}_3(t_j))}$  but  $\Sigma_{\mathcal{R}}.\text{Ver}^{\text{H}_4}((g, g^{\text{H}_3(t_j)} \cdot pk_j, B'_2, B_{j,1-b}), \pi_2) = 1$ .

Here  $pk_j, t_j$  are the values queried to the  $j^{\text{th}}$  call to  $\text{REQ}$  and we let  $sk_j = \text{dlog}_g pk_j$ . Recall that here  $\mathcal{R} = \{(g, U, V, W), (\alpha) : U = g^\alpha \wedge W = V^\alpha\}$ , and verification is therefore checking, in case (1), that

$$B'_2 = B_{j,b}^{\text{H}_3(t_j)+sk_j} \Leftrightarrow (B'_2)^{1/(sk_j+\text{H}_3(t_j))} = B_{j,b}$$

and a similar equality for case (2). So if this event occurs, this means the adversary has violated the soundness of the ZKP: only a single value  $\alpha = sk_j + \text{H}_3(t_j)$  can be the witness for  $\mathcal{R}$ .

To formally reduce to ZKP soundness, first let game  $G_0$  be the same as  $\text{POPRIV2}_{3\text{H},\mathcal{P},b}^{\mathcal{A}_{\text{po-priv2}}}(\lambda)$  but the bit  $b$  is chosen at random from  $\{0, 1\}$ . Let “ $G_0 \Rightarrow b$ ” be the event in game  $G_0$  that the game returns the value  $b$ . (We use this event notation for subsequent games analogously.) Further we let  $G_{0_{\text{bad}}}$  be the same as  $G_0$  except that within each  $\text{FIN}$  it first computes  $sk_j = \text{dlog}_g pk_j$  and checks if conditions (1) and (2) hold. If either does not, then it sets a flag  $\text{bad}$ . Clearly  $G_{0_{\text{bad}}}$  is not computationally efficient; our reduction will avoid this computationally inefficient step. Finally we let  $G_1$  be the same as game  $G_0$  except that all  $\text{FIN}(j, \text{rep}, \text{rep}')$  queries are handled by first replacing  $\text{rep}$  and  $\text{rep}'$  with the correct values, i.e.,  $\text{rep} \leftarrow B_{j,b}^{sk_j}$  and  $\text{rep}' \leftarrow B_{j,1-b}^{sk_j}$  where  $sk_j \leftarrow \text{dlog}_g(pk_j)$ . Notice that  $G_{0_{\text{bad}}}$  and  $G_1$  are identical until the first query, if any, that sets the flag  $\text{bad}$ . We have that

$$\text{Adv}_{3\text{H},\mathcal{P},\mathcal{A}_{\text{po-priv2}}}^{\text{po-priv2}}(\lambda) = 2 \cdot \Pr[G_0 \Rightarrow b] - 1.$$

and that

$$\Pr[G_0 \Rightarrow b] = \Pr[G_{0_{\text{bad}}} \Rightarrow b] \leq \Pr[G_1 \Rightarrow b] + \Pr[G_{0_{\text{bad}}} \text{ sets bad }],$$

where the inequality comes from the fact that  $G_{0_{\text{bad}}}$  and  $G_1$  are identical-until- $\text{bad}$  and application of the fundamental lemma of game playing [BR06]. We now bound the probability that  $\Pr[G_{0_{\text{bad}}} \text{ sets bad}]$  via reduction to the soundness of the ZKP.

Adversary  $\mathcal{B}_{\text{sound}}$  works as follows. First, it randomly chooses a number  $q^* \in [1, 2q]$  to serve as its guess for which ZKP  $\pi$  will be forged by the adversary. Here  $q$  is the maximum number of  $\text{FIN}$  queries made by  $\mathcal{A}_{\text{po-priv2}}$ ; each such query includes two proofs. Then  $\mathcal{B}_{\text{sound}}$  runs  $G_0$ , stopping when  $\mathcal{A}_{\text{po-priv2}}$  has made  $j = \lceil q^*/2 \rceil$  queries to  $\text{FIN}$ . At this point,  $\mathcal{B}_{\text{sound}}$  stops outputting  $((g, g^{\text{H}_3(t_j)} pk_j, B'_1, B_{j,b}), \pi_1)$  if  $q^*$  is odd and  $((g, g^{\text{H}_3(t_j)} pk_j, B'_2, B_{j,1-b}), \pi_2)$  otherwise. Adversary  $\mathcal{B}_{\text{sound}}$  avoids computing  $sk_j$ ; it simply guesses which of the proofs would have caused  $\text{bad}$  to be set to true, had  $sk_j$  been computed and the conditions (1) and (2) been checked. A standard argument yields that

$$\Pr[G_{0_{\text{bad}}} \text{ sets bad}] \leq 2q \cdot \text{Adv}_{\text{NiZK},\mathcal{R},\mathcal{P},\mathcal{B}_{\text{sound}}}^{\text{sound}}(\lambda).$$

To finish the proof, we can observe that  $G_1$  always correctly computes responses, and a similar argument as we used for  $\text{POPRIV1}$  gives that the transcript observed by  $\mathcal{A}_{\text{po-priv2}}$  is independent of the challenge bit  $b$ , and so  $\Pr[G_1 \Rightarrow b] = 1/2$ . Combining all the above yields the advantage statement in the theorem. ■

## 6 Performance Evaluation

We implemented  $3\text{HashSDHI}$  to measure the computational cost of the protocol in comparison to related protocols, including the baseline  $2\text{HashDH}$  VOPRF from [DFHSW20], Pythia [ECS<sup>+</sup>15], and the recent attribute-based VOPRF (ABVOPRF) from Facebook [HIJ<sup>+</sup>21]. Each protocol was implemented in a minimal fashion, e.g., by omitting domain separating hash function invocations, in order to emphasize the cost of core public key operations. Our implementations

Scheme	L (B)	T (B)	KeyGen	KeyVerify	Req	BlindEv	Finalize	Ev
2HashDH	16	1	0	0	73	222	392	77
2HashDH	16	8	0	0	75	229	404	79
2HashDH	16	64	0	0	84	256	447	89
3HashSDHI	16	1	0	0	85	369	527	125
3HashSDHI	16	8	0	0	76	334	477	112
3HashSDHI	16	64	0	0	75	328	471	110
Pythia	16	1	168	0	849	4068	6070	2871
Pythia	16	8	180	0	831	4099	6092	2881
Pythia	16	64	171	0	809	3922	5849	2768
ABVOPRF	16	1	294	292	74	517	684	370
ABVOPRF	16	8	1910	2386	76	2135	2789	1981
ABVOPRF	16	64	15053	19196	80	15305	19702	15163

Table 1: Average operation time for various POPRF protocols. All times are measured in  $\mu\text{s}$ , and a zero time represents some value much less than a single microsecond.

use the ristretto255 group [dVGT+20] where prime-order groups are required, and the bn256 curve for Pythia, where pairing-friendly curves are required. We implemented each protocol in Go using the CIRCL experimental cryptographic library [FHK19] and `bn256` package. These benchmarks were evaluated on a machine with a 2.6 GHz 6-Core Intel Core i7 CPU and 32 GB RAM running macOS 10.15.7. Below we report on the average time over 1000 measurements of each operation.

Our benchmarks profile the functions needed to configure a client for the protocol, including the `KeyGen` and `KeyVerify` (for ABVOPRF), as well as the functions used to carry out the protocol, including `Req`, `BlindEv`, and `Finalize`. We also profiled the `FullEvaluation` routine to compare the cost of the blinding and proof verification operations in the protocol. For each protocol scheme, we keep the input size ( $L$ ) fixed but vary the metadata size ( $T$ ) in bytes. The results of our analysis are given in Table 1.

We comment on two key results in this data. First, the difference between the baseline VOPRF protocol and POPRF protocol are minuscule (as a function of metadata size). In particular, the POPRF introduces approximately a 25% overhead (in terms of  $\mu\text{s}$  to compute), though this is likely negligible at scale. Second, there is nearly an order of magnitude difference between the POPRF and the ABVOPRF as a function of metadata size. This is primarily due to the online `KeyGen` process and its resulting linear cost in proof evaluation. This suggests that the POPRF construction scales better in the presence of arbitrary-size tag values.

Some of the protocols included allow certain operations to be computed offline, thereby improving online protocol performance. For example, if the tag value is known in advance, 3HashSDHI can pre-compute the private key used in the `BlindEv` call. Likewise, in the ABVOPRF protocol, the key generation and

Scheme	T (B)	Request	Evaluate	Finalize	Ev
2HashDH	1	73	223	392	77
2HashDH	8	75	231	403	79
2HashDH	64	83	255	447	89
3HashDH	1	84	367	526	124
3HashDH	8	76	332	478	112
3HashDH	64	76	328	472	110
Pythia	1	847	3913	6074	2832
Pythia	8	824	3935	6085	2840
Pythia	64	806	3764	5843	2715
ABVOPRF	1	74	224	685	77
ABVOPRF	8	76	228	2790	79
ABVOPRF	64	75	230	19721	80

Table 2: Average times for performance, excluding precomputation cost. All times are measured in  $\mu s$ .

verification operations can be computed offline and clients can cache the results. Table 2 summarizes the cost of operations, discounting precomputation costs. Note that although the performance of the ABVOPRF construction improves, it still introduces substantially more overhead than the 3HashSDHI construction.

## 7 Applications

POPRFs provide a new degree of flexibility that we observe to be useful in a variety of applications. One example is Facebook’s de-identified logging system [HIJ<sup>+</sup>21], where 3HashSDHI provides a more efficient solution compared to the current Naor-Reingold style approach. But essentially anywhere an OPRF is used we see opportunity for POPRFs to provide potential benefits in terms of increasing deployment flexibility, reducing key management challenges, and/or improving security. Here we briefly discuss three previously mentioned motivating applications: anonymous one-time-use tokens, password breach alerting, and password-based authenticated key exchange.

### 7.1 Privacy Pass

Privacy Pass [DGS<sup>+</sup>18, CDFH21] is a protocol in which users are allowed to receive one-time-use anonymous tokens, using an issuance protocol, that can later be used to anonymously authenticate themselves using a redemption protocol. It can be used as a solution that works with any system that employs challenges to assert whether a user is honest or not, without falling back to the usage of block-listing or a “catch-all” solution (e.g., prompting all users with authentication or CAPTCHA challenges).

The primary component underlying Privacy Pass is a verifiable OPRF (VOPRF). Current implementations use 2HashDH [JKK14]. The issuance protocol has a client request VOPRF output for a random input  $x$  under a Privacy Pass server held VOPRF secret key  $sk$ . The client must verify that the received token  $y = 2\text{HashDH.Ev}(sk, x)$  is correct relative to the server’s public key  $pk$ . A token  $(x, y)$  can then be redeemed by sending to the server  $(x, \text{MAC}_y(\text{data}))$  where MAC is a message authentication code such as HMAC and  $\text{data}$  is some application-specific bit string (called the binding data in [DGS<sup>+</sup>18]). The server can recompute  $y$  and then use it to check the MAC value.

The core security properties achieved by Privacy Pass are unlinkability and unforgeability. Unlinkability derives from the request privacy properties of the VOPRF, which ensures that a malicious server learns nothing about a client’s input  $x$  nor can they link  $(x, y)$  to a particular issuance query. Unforgeability derives from the pseudorandomness security of the VOPRF: given the ability to obtain  $q$  tokens, the adversary can at most compute  $q$  outputs of the VOPRF.

However, one abuse of Privacy Pass not prevented by the current design is what we refer to as a hoarding attack, also called a farming attack [DGS<sup>+</sup>18]. A malicious user can gather a large number of tokens by running the token issuance protocol as many times as possible over some period of time. Later, the malicious user can redeem all the gathered tokens at once in an effort to render the provided service unavailable in a (D)DoS attack (by, for example, overwhelming a website with expensive requests).

One potential defense against hoarding attacks is to force periodic rotation of the VOPRF secret key, such as once per week. But this is clumsy because it requires clients to verify that key rotations are made legitimately: a server that rotates too often can violate unlinkability. In the limit, a malicious server could pick a separate  $pk$  for each client issuance, thereby completely violating unlinkability. Forcing clients to use gossip protocols (to verify that the same public keys are used) or using public ledgers to record public keys for monitoring purposes require further complicating infrastructure.

Use of a verifiable POPRF provides a simpler, more elegant solution. Tokens can be bound to a public input  $t$  that lets the server and client agree upon a scope for token issuance. In this case, one replaces the Privacy Pass’  $2\text{HashDH.Ev}(sk, x)$  with  $\text{DY.Ev}(sk, t, x)$  in both the issuance and redemption phases. One could use as  $t$  a coarse timestamp, such as the current day or week at which an issuance occurred, and then redemption could enforce a policy about the staleness of tokens, forcing them to be redeemed within some time period. Alternatively, one could imagine using  $t$  to bind issuance and redemption to be performed from the same server-visible client IP address (though, verifiability would require that client’s can infer their server-visible IP address).

We note that all these hoarding mitigations reduce the anonymity set of a redemption to only the clients issued tokens under  $t$ . This is true also for the key rotation approach, reducing the anonymity set to clients that were issued tokens under the particular public key. Some degradation of privacy seems fundamental to restrict token use, and choosing how to make use of  $t$  requires care and further

work to identify best practices. But the point is that POPRFs provide a degree of flexibility that easily allows a variety of choices.

## 7.2 Private Set Membership and Breach Alerting

One widely deployed use of OPRFs is for password breach alerting (also called compromised credential checking) [TPY<sup>+</sup>19, Hun]. These use an OPRF-based bucketized private set membership protocol [TPY<sup>+</sup>19, LPA<sup>+</sup>19] that works as follows. The server generates a long-lived 2HashDH secret key  $sk$  and computes  $y_{u,pw} = 2\text{HashDH.Ev}(sk, u, pw)$  for each username, password pair  $(u, pw) \in D$ . Here  $D$  is a breach database of known-compromised pairs. Then, to perform a lookup for  $(u^*, pw^*)$ , the client sends a truncated hash of the username that forms a bucket identifier  $\beta = H(u^*)$ , as well as a blind evaluation request for the client’s username, password pair  $(u^*, pw^*)$  that they want to check. The server computes the OPRF response and sends it back along with the bucket  $B = \{y_{u,pw} \mid H(u) = \beta\}$  of values that have matching truncated username hash. The client finishes computing  $y^* = 2\text{HashDH.Ev}(sk, u^*, pw^*)$  and checks if  $y^* \in B$ . If so, their credential is known to be compromised; otherwise, it is not.

In the currently deployed protocols, there is no way to enforce that the client’s query indicates the correct bucket identifier  $B$ . Instead, they could query for  $b' = H(u')$  for  $u' \neq u^*$  and complete the protocol execution checking for values in some other bucket. Whether this opens up password breach alerting services to abuse is not clear.

Nevertheless, we observe that replacing the OPRF with a POPRF provides a simple way to cryptographically bind the buckets to particular bucket identifiers, by setting  $t = \beta$ . One could similarly do so using per-bucket secret keys  $sk_\beta$  for a standard OPRF, but this would complicate key management.

## 7.3 OPAQUE

As mentioned in the introduction, OPAQUE [JKX18] is a strong aPAKE protocol (SaPAKE) that provides password-based mutual authentication in a client-server setting without having to rely on Public Key Infrastructure (PKI). It provides security against pre-computation attacks upon server compromise, and increases the protection against offline dictionary attacks, as an attacker will have to perform an exhaustive per-user attack upon server’s data compromise. OPAQUE is a protocol also amenable to a multi-server distributed implementation where an offline dictionary attack is only possible if a threshold of servers is compromised.

OPAQUE can be thought of as a protocol that works as a “compiler” by transforming a suitable AKE protocol (resistant to key compromise impersonation attacks and forward secrecy) into a secure aPAKE protocol using an OPRF. Current implementations use 2HashDH. It consists of two phases: an offline registration phase and an online authenticated key exchange phase.

The purpose of the offline phase is to register a user’s account using their unique user identity and their password. The user identity could be a username or email address. To do so, the user opaquely registers its password without

the server ever knowing it. The client and server obliviously compute  $\text{rwd} = 2\text{HashDH.Ev}(sk_u, x)$ , where  $x$  is the user’s password and  $sk_u$  is a server’s random, per-user generated OPRF secret key. The output of this functionality,  $\text{rwd}$ , is used to encrypt its AKE private key. The resulting ciphertext is stored on the server-side alongside with other user’s credentials, such as its corresponding user id. The server will store, as well, the per-user OPRF key  $sk_u$  used to generate  $\text{rwd}$ .

To perform online authentication, the client and server obliviously recompute  $\text{rwd}$ , enabling the client to recover their AKE private key by decrypting the ciphertext (sent back to the client during the OPRF flows), and complete a (now password-authenticated) AKE exchange.

A complexity for deployment of OPAQUE is that the server has to maintain and keep a consistent view of the per-user OPRF keys and associated AKE private key ciphertexts. Implementation vulnerabilities may arise should servers incorrectly use the same OPRF key across multiple users, allowing potentially for cross-user attacks that allow logging in as the wrong user. Implementations are also likely to store the per-user OPAQUE secret keys in the same user database alongside other per-user data, meaning that compromises that allow exfiltrating the database (e.g., SQL injection) will reveal all information needed to brute-force recover passwords.

One potential approach instead would be to, again, replace the per-user OPRF key  $sk_u$  with a POPRF with global  $sk$  and use the user identity as  $t$ . This would allow protecting  $sk$  by storing it in a separate hardened crypto service (similar to deployment models used in password hardening, see [ECS<sup>+</sup>15]) only once. One potential complication is that performing periodic key rotations for  $sk$  would be more challenging, since it would require somehow either resetting all users passwords (not reasonable in most contexts by asking users to reset) or rolling clients to new keys as they login by maintaining old  $sk$  for some period of time. In contrast, the per-user  $sk_u$  approach can selectively rotate OPRF keys for each user as needed.

## Acknowledgments

The authors would like to thank Tjerand Silde and Martin Strand for discussions about their preprint.

## References

- AF96. Masayuki Abe and Eiichiro Fujisaki. How to date blind signatures. In *ASIACRYPT*, volume 1163 of *Lecture Notes in Computer Science*, pages 244–251. Springer, 1996.
- AO00. Masayuki Abe and Tatsuaki Okamoto. Provably secure partially blind signatures. In *CRYPTO*, volume 1880 of *Lecture Notes in Computer Science*, pages 271–286. Springer, 2000.
- BFL20. Balthazar Bauer, Georg Fuchsbauer, and Julian Loss. A classification of computational assumptions in the algebraic group model. In *CRYPTO*

- (2), volume 12171 of *Lecture Notes in Computer Science*, pages 121–151. Springer, 2020.
- BM93. Steven Bellovin and Michael Merritt. Augmented encrypted key exchange: a password based protocol secure against dictionary attacks and password file compromise. In *CCS*, pages 244–250. ACM, 1993.
- BNPS03. Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The one-more-rsa-inversion problems and the security of chaum’s blind signature scheme. *J. Cryptol.*, 16(3):185–215, 2003.
- Bol03. Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-diffie-hellman-group signature scheme. In *Public Key Cryptography*, volume 2567 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2003.
- BR06. Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426. Springer, 2006.
- BS17. Dan Boneh and Victor Shoup. *A Graduate Course in Applied Cryptography*. 2017. Version 0.4.
- Cam98. Jan Camenisch. *Group signature schemes and payment systems based on the discrete logarithm problem*. PhD thesis, ETH Zurich, Zürich, Switzerland, 1998.
- CDFH21. Sofia Celi, Alex Davidson, and Armando Faz-Hernandez. Privacy Pass Protocol Specification. Internet-Draft draft-ietf-privacypass-protocol-00, Internet Engineering Task Force, January 2021. Work in Progress.
- Che06. Jung Hee Cheon. Security analysis of the strong diffie-hellman problem. In Serge Vaudenay, editor, *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, volume 4004 of *Lecture Notes in Computer Science*, pages 1–11. Springer, 2006.
- CP92. David Chaum and Torben P. Pedersen. Wallet databases with observers. In *CRYPTO*, volume 740 of *Lecture Notes in Computer Science*, pages 89–105. Springer, 1992.
- DFHSW20. Alex Davidson, Armando Faz-Hernandez, Nick Sullivan, and Christopher A. Wood. Oblivious Pseudorandom Functions (OPRFs) using Prime-Order Groups. Internet-Draft draft-irtf-cfrg-voprf-05, Internet Engineering Task Force, November 2020. Work in Progress.
- DGS<sup>+</sup>18. Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. Privacy pass: Bypassing internet challenges anonymously. *Proc. Priv. Enhancing Technol.*, 2018(3):164–180, 2018.
- dVGT<sup>+</sup>20. Henry de Valence, Jack Grigg, George Tankersley, Filippo Valsorda, Isis Lovecruft, and Mike Hamburg. The ristretto255 and decaf448 Groups. Internet-Draft draft-irtf-cfrg-ristretto255-decaf448-00, Internet Engineering Task Force, October 2020. Work in Progress.
- DY05. Yevgeniy Dodis and Aleksandr Yampolskiy. A verifiable random function with short proofs and keys. In *Public Key Cryptography*, volume 3386 of *Lecture Notes in Computer Science*, pages 416–431. Springer, 2005.
- ECS<sup>+</sup>15. Adam Everspaugh, Rahul Chatterjee, Samuel Scott, Ari Juels, and Thomas Ristenpart. The pythia PRF service. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 547–562. USENIX Association, 2015.

- FHK19. Armando Faz-Hernandez and Kris Kwiatkowski. *Introducing CIRCL: An Advanced Cryptographic Library*. Cloudflare, June 2019. <https://github.com/cloudflare/circl>.
- FIPR05. Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In *TCC*, volume 3378 of *Lecture Notes in Computer Science*, pages 303–324. Springer, 2005.
- Fis06. Marc Fischlin. Round-optimal composable blind signatures in the common reference string model. In *CRYPTO*, volume 4117 of *Lecture Notes in Computer Science*, pages 60–77. Springer, 2006.
- FKL18. Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In *CRYPTO (2)*, volume 10992 of *Lecture Notes in Computer Science*, pages 33–62. Springer, 2018.
- FPS20. Georg Fuchsbauer, Antoine Plouviez, and Yannick Seurin. Blind schnorr signatures and signed elgamal encryption in the algebraic group model. In *EUROCRYPT (2)*, volume 12106 of *Lecture Notes in Computer Science*, pages 63–95. Springer, 2020.
- HIJ<sup>+</sup>21. Sharon Huang, Subodh Iyengar, Sundar Jeyaraman, Shiv Kushwah, Chen-Kuei Lee, Zutian Luo, Payman Mohassel, Ananth Raghunathan, Shaahid Shaikh, Yen-Chieh Sung, and Albert Zhang. PrivateStats: De-Identified Authenticated Logging at Scale, January 2021.
- HJK<sup>+</sup>21. Sharon Huang, Subodh Iyengar, Sundar Jeyaraman, Shiv Kushwah, Chen-Kuei Lee, Zutian Luo, Payman Mohassel, Ananth Raghunathan, Shaahid Shaikh, Yen-Chieh Sung, and Albert Zhang. Privatestats: De-identified authenticated logging at scale. [Online; last accessed Feb 16 2021], 2021. [https://research.fb.com/wp-content/uploads/2021/01/PrivateStats-De-Identified-Authenticated-Logging-at-Scale\\_final.pdf](https://research.fb.com/wp-content/uploads/2021/01/PrivateStats-De-Identified-Authenticated-Logging-at-Scale_final.pdf).
- Hun. Troy Hunt. Have i been pwned. <https://haveibeenpwned.com/>.
- JKK14. Stanislaw Jarecki, Aggelos Kiayias, and Hugo Krawczyk. Round-optimal password-protected secret sharing and T-PAKE in the password-only model. In *ASIACRYPT (2)*, volume 8874 of *Lecture Notes in Computer Science*, pages 233–253. Springer, 2014.
- JKX18. Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. OPAQUE: an asymmetric PAKE protocol secure against pre-computation attacks. In *EUROCRYPT (3)*, volume 10822 of *Lecture Notes in Computer Science*, pages 456–486. Springer, 2018.
- JL09. Stanislaw Jarecki and Xiaomin Liu. Efficient oblivious pseudorandom function with applications to adaptive OT and secure computation of set intersection. In *TCC*, volume 5444 of *Lecture Notes in Computer Science*, pages 577–594. Springer, 2009.
- JT20. Joseph Jaeger and Nirvan Tyagi. Handling adaptive compromise for practical encryption schemes. In *CRYPTO (1)*, volume 12170 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2020.
- KLW21. Hugo Krawczyk, Kevin Lewi, and Christopher A. Wood. The OPAQUE Asymmetric PAKE Protocol. Internet-Draft draft-irtf-cfrg-opaque-02, Internet Engineering Task Force, February 2021. Work in Progress.
- KZ08. Aggelos Kiayias and Hong-Sheng Zhou. Equivocal blind signatures and adaptive uc-security. In *TCC*, volume 4948 of *Lecture Notes in Computer Science*, pages 340–355. Springer, 2008.

- LPA<sup>+</sup>19. Lucy Li, Bijeeta Pal, Junade Ali, Nick Sullivan, Rahul Chatterjee, and Thomas Ristenpart. Protocols for checking compromised credentials. In *CCS*, pages 1387–1403. ACM, 2019.
- Mau05. Ueli M. Maurer. Abstract models of computation in cryptography. In Nigel P. Smart, editor, *Cryptography and Coding, 10th IMA International Conference, Cirencester, UK, December 19-21, 2005, Proceedings*, volume 3796 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2005.
- NR97. Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *FOCS*, pages 458–467. IEEE Computer Society, 1997.
- Sho97. Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT*, volume 1233 of *Lecture Notes in Computer Science*, pages 256–266. Springer, 1997.
- SS21. Tjerand Silde and Martin Strand. Anonymous tokens with public metadata and applications to private contact tracing. Cryptology ePrint Archive, Report 2021/203, 2021. <https://eprint.iacr.org/2021/203>.
- TPY<sup>+</sup>19. Kurt Thomas, Jennifer Pullman, Kevin Yeo, Ananth Raghunathan, Patrick Gage Kelley, Luca Invernizzi, Borbala Benko, Tadek Pietraszek, Sarvar Patel, Dan Boneh, and Elie Bursztein. Protecting accounts from credential stuffing with password breach alerting. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 1556–1571. USENIX Association, 2019.
- ZSS03. Fangguo Zhang, Reihaneh Safavi-Naini, and Willy Susilo. Efficient verifiably encrypted signature and partially blind signature from bilinear pairings. In *INDOCRYPT*, volume 2904 of *Lecture Notes in Computer Science*, pages 191–204. Springer, 2003.
- ZSS04. Fangguo Zhang, Reihaneh Safavi-Naini, and Willy Susilo. An efficient signature scheme from bilinear pairings and its applications. In *Public Key Cryptography*, volume 2947 of *Lecture Notes in Computer Science*, pages 277–290. Springer, 2004.

## A Security of $(m, n)$ -OM-GAP-SDHI

*Proof (Of Theorem 1).* This proof proceeds in two parts. (1) First, we present  $\mathcal{A}_{\text{uber}}$  and argue that it almost perfectly simulates the  $(m, n)$ -OM-GAP-SDHI game for  $\mathcal{A}_{\text{sdhi}}$ , up to a small statistical difference. (2) Second, we argue that if  $\mathcal{A}_{\text{sdhi}}$  outputs a set of values that wins the  $(m, n)$ -OM-GAP-SDHI game, it must be that at least one of those values has a non-trivial representation in the exponent that wins the  $m$ -UBER game.

### A.1 Simulation of $(m, n)$ -OM-GAP-SDHI environment

We construct adversary  $\mathcal{A}_{\text{uber}}$  as shown in Figure 9. The adversary first constructs a polynomial  $G(\vec{X}) \leftarrow \prod_i^n (X_0 + c_i)^{q_i}$  from the chosen  $c$  values in the first stage of  $\mathcal{A}_{\text{sdhi}}$ . The adversary receives an evaluation  $\hat{g} = g^{G(\vec{x})}$  from its EV oracle that it will pass as the generator in its simulation to the second stage of  $\mathcal{A}_{\text{sdhi}}$ . The public key  $\hat{g}^{x_0}$  is simulated by evaluating  $X_0 \cdot G(\vec{X})$ ; and the  $m$

challenge points will each be simulated with additional variables  $X_i$  for  $i \in [1, m]$  by evaluating polynomial  $X_i \cdot G(\vec{X})$ .

To simulate the SDH oracle on input  $(Y, i)$ , first assume that  $\mathcal{A}_{\text{uber}}$  knows the polynomial  $P(\vec{X})$  such that  $Y = g^{P(\vec{x})}$ . We will show shortly how  $\mathcal{A}_{\text{uber}}$  can compute  $P(\vec{X})$  from the algebraic representation of  $Y$ . Given  $P(\vec{X})$ ,  $\mathcal{A}_{\text{uber}}$  simulates by computing the polynomial  $P(\vec{X})/(X_0 + c_i)$  and returning its evaluation. Now to compute  $P(\vec{X})$ ,  $\mathcal{A}_{\text{uber}}$  takes the algebraic representation of  $Y$  in terms of elements given to  $\mathcal{A}_{\text{sdhi}}$ . The elements given to  $\mathcal{A}_{\text{sdhi}}$  are the initial elements plus the elements output from previous queries to SDH. The polynomial exponents of all of these elements are known to  $\mathcal{A}_{\text{uber}}$ , so they can be combined via the linear combination indicated by the algebraic representation to compute  $P(\vec{X})$ . Furthermore,  $P(\vec{X})/(X_0 + c_i)$  will always be computable due to the construction of  $G(\vec{X})$  including  $q_i$  factors of  $(X_0 + c_i)$  where  $q_i$  is the maximum number of queries for  $c_i$  to SDH. All initial elements have the polynomial  $G(\vec{X})$  as a factor of the exponent, and while subsequent elements returned from SDH divide out different factors of  $(X_0 + c_j)$ , they never exhaust a factor  $(X_0 + c_j)$  unless the maximum number of queries for that index has been reached. Lastly,  $\mathcal{A}_{\text{uber}}$  simulates SDDH by passing the query on to its own DECIDE oracle.

As long as the maximum query counts that  $\mathcal{A}_{\text{uber}}$  uses to create  $G(\vec{X})$  are abided by, the environment for  $\mathcal{A}_{\text{sdhi}}$  is almost perfectly simulated, as we argue below. Moreover,  $\mathcal{A}_{\text{uber}}$  selects an output of  $\mathcal{A}_{\text{sdhi}}$  at random and returns it along with its polynomial representation.

Concretely, let  $\pi_0$  be the probability that at least one of the values in the output of  $\mathcal{A}_{\text{sdhi}}$  has a non-trivial representation in the exponent that wins the  $m$ -UBER game in the original game  $(m, n)$ -OM-GAP-SDHI. Let  $\pi_1$  be the probability that the same happens within the simulation of  $\mathcal{A}_{\text{uber}}$ . Then, we argue that

$$\pi_0 - \pi_1 \leq \frac{q}{p}.$$

This can be seen as follows: The only difference between the simulation and the original game is that the former uses  $(\hat{g}, \hat{X})$ , where  $\hat{g} = g^{G(x)}$  and  $\hat{X} = g^{x \cdot G(x)}$  for a  $x \leftarrow_s \mathbb{Z}_p$  and a polynomial  $G(X)$  of degree  $q$ , instead of  $(g, g^x)$  in the latter, where  $g$  is a random generator. In particular,  $\pi_0 - \pi_1$  is upper bounded by the statistical distance between  $(y \cdot G(x), y \cdot x \cdot G(x))$  and  $(y, y \cdot x)$ , where  $(x, y) \leftarrow_s \mathbb{Z}_p^2$ . Now, let  $S \subseteq \mathbb{Z}_p^2$  be the set pairs  $(x, y)$  such that  $G(x) \neq 0$ . For any  $(z_1, z_2) \in \mathbb{Z}_p^2$ , and  $(x, y) \leftarrow_s S$ , we now have

$$\begin{aligned} \Pr[(y, y \cdot x) = (z_1, z_2)] &= \Pr[(y, z_1 \cdot x) = (z_1, z_2)] \\ &= \Pr[(y \cdot G(x), z_1 \cdot x) = (z_1, z_2)] = \Pr[(y \cdot G(x), y \cdot x \cdot G(x)) = (z_1, z_2)]. \end{aligned}$$

Therefore, the statistical distance is upper bounded by the probability that  $(x, y) \leftarrow_s \mathbb{Z}_p^2$  is in  $S$ , which in turn is the probability that  $G(x) = 0$ . The latter is at most  $q/p$  by the Schwartz-Zippel Lemma.

In conclusion, the advantage of  $\mathcal{A}_{\text{uber}}$  is

$$\text{Adv}_{\text{Gen}, \mathcal{A}_{\text{uber}}}^{(m+1)\text{-uber}}(\lambda) \geq \frac{\pi_1}{\ell} \geq \frac{\pi_1}{q_{\max} + 1} \geq \frac{\pi_0}{q_{\max} + 1} - \frac{q}{p(q_{\max} + 1)},$$

<p>Adversary <math>\mathcal{A}_{\text{uber}}^{\text{Ev,Decide}}(p, \mathbb{G}, g)</math></p> <p><math>(st_{\mathcal{A}}, [c_i]_i^n) \leftarrow \mathcal{A}_1(p, \mathbb{G})</math></p> <p><math>G(\vec{X}) \leftarrow \prod_i^n (X_0 + c_i)^{q_i}</math></p> <p><math>\hat{g} \leftarrow \text{Ev}(G(\vec{X})); \vec{X} \leftarrow \text{Ev}(X_0 \cdot G(\vec{X}))</math></p> <p><math>[\hat{A}]_{i=1}^m \leftarrow [\text{Ev}(X_i \cdot G(\vec{X}))]_{i=1}^m</math></p> <p><math>(\gamma, [Z_i, \alpha_i]_i^\ell) \leftarrow \mathcal{A}_2^{\text{SDH, SDDH}}(\hat{g}, \hat{X}, [\hat{A}]_i^m; st_{\mathcal{A}})</math></p> <p><math>j \leftarrow \mathbb{N}_\ell; \mu(\vec{X}) = X_{\alpha_j} (X_0 + c_\gamma)^{q_\gamma - 1} \prod_{i \neq \gamma}^n (X_0 + c_i)^{q_i}</math></p> <p>Return <math>(Z_j, \mu(\vec{X}))</math></p>	<p>Oracle <math>\text{SDH}(Y, i)</math></p> <p>Compute <math>P(\vec{X}) : Y = g^{P(\vec{x})}</math> from representation of <math>Y</math></p> <p><math>Z \leftarrow \text{Ev}(P(\vec{X}) / (X_0 + c_i))</math></p> <p>Return <math>Z</math></p> <p>Oracle <math>\text{SDDH}(Y, Z, i)</math></p> <p><math>P(A_1, A_2, A_3) \leftarrow (A_1 + c_i) \cdot A_3 - A_2</math></p> <p>Return <math>\text{Decide}(P(A_1, A_2, A_3), [\vec{X}, Y, Z])</math></p>
---	---

Fig. 9: Adversary  $\mathcal{A}_{\text{uber}}$  in the security proof of  $(m, n)$ -OM-GAP-SDHI.

because, without loss of generality, we can assume that  $\ell \leq q_{\max} + 1$ . Below, we are going to prove that  $\text{Adv}_{\text{GGen}, \mathcal{A}_{\text{sDHI}}}^{(m, n)\text{-om-gap-sdhi}}(\lambda) = \pi_0$ , which concludes the proof.

## A.2 Linear independence of winning elements

Next, we argue that one of the strong Diffie-Hellman values output by a winning  $\mathcal{A}_{\text{sDHI}}$  corresponds to a group element with a non-trivial polynomial in the exponent, i.e., a polynomial that is linearly independent from all queried polynomials. Our ultimate goal will be to claim that if the adversary produces  $\ell$  winning group elements for a value  $c_\gamma$  having only queried SDH  $q_\gamma < \ell$  times, then one of the winning elements must have a non-trivial (linearly-independent) polynomial exponent from all group elements given to the adversary during initialization and as output from SDH.

We will find that it is easy to argue the linear independence of the initial group elements and winning elements. The main challenge we will face is reasoning about elements output from SDH. In previous formulations of one-more assumptions [Bol03], the number of winning elements was required to be greater than the number of queries to the “one-more” oracle. In these cases, the typical proof strategy is show linear independence of the set of winning elements and initial elements – since this set is strictly larger than the set of query elements and initial elements, it must be that at least one winning element is linearly independent of the set of query elements and initial elements. However, this strategy will not work for OM-Gap-SDHI, which simply requires the number of winning elements to be greater than the number of queries *for the chosen* winning  $c_\gamma$ .

Intuitively, we would like to say that elements returned from SDH for other  $c_{j \neq \gamma}$  will not be helpful in constructing a winning element for  $c_\gamma$ . If this is the case, then we can conclude by arguing that a set of  $q_\gamma < \ell$  elements cannot construct a set of  $\ell$  linearly-independent winning elements.

Unfortunately, it is not immediately clear this is the case, as elements that were previously output for queries to  $c_\gamma$  may be passed back in to SDH under a different  $c_j$ . This leads to the possibility that many elements returned from SDH ( $> q_\gamma$ ) may have a “dependence” on  $c_\gamma$ . The main step in our proof shows that whenever a query to SDH on  $c_i$  is made, the output element can be refactored as an element that only depends on the queried  $c_i$ , regardless of whether the input element includes previous outputs from SDH dependent on other  $c_{j \neq i}$ .



Then, we observe that the terms included in  $\tau_i$  correspond to the terms for binary strings  $2^{i-1}$  to  $2^i - 1$ . We get the following expression for  $\tau_i$  based on the above interpretation:

$$\begin{aligned}\tau_i &= \frac{\mathcal{Y}_i + \sum_{j=1}^{i-1} a_{i,j} \tau_j}{X_0 + c_1} \\ &= \sum_{s=2^{i-1}}^{2^i-1} \Omega(s)\end{aligned}\tag{2}$$

The above form follows from an induction argument on the form of all  $\tau_{j < i}$ . Multiplying  $\tau_j$  by  $a_{i,j}/(X_0 + c_1)$  corresponds exactly to transforming all  $\Omega(s')$  for  $s' \in [2^{j-1}, 2^j - 1]$  to  $\Omega(s)$  for  $s \in [2^{i-1} + 2^{j-1}, 2^{i-1} + 2^j - 1]$ . Summing up the terms performing this transformation on all  $\tau_{j < i}$  corresponds to all  $s \in [2^{i-1} + 1, 2^i - 1]$ , and the final term corresponding to  $s = 2^{i-1}$  is added separately.

Lastly, we rearrange the terms of  $\tau_i$  grouping them by  $\mathcal{Y}_j$  which will be useful later on:

$$\begin{aligned}\tau_i &= \sum_{s=2^{i-1}}^{2^i-1} \Omega(s) \\ &= \Omega(2^{i-1}) + \sum_{j=1}^{i-1} \sum_{s=2^{i-j-1}}^{2^{i-j}-1} \Omega(2^j * s + 2^{j-1}) \\ &= \frac{\mathcal{Y}_i}{X_0 + c_1} + \sum_{j=1}^{i-1} \sum_{s=2^{i-j-1}}^{2^{i-j}-1} \frac{\mathcal{Y}_j \cdot \prod_{k=1}^{|\omega(2^j * s + 2^{j-1})|-1} a_{\omega(2^j * s + 2^{j-1})_k, \omega(2^j * s + 2^{j-1})_{k+1}}}{(X_0 + c_1)^{|\omega(2^j * s + 2^{j-1})|}}\end{aligned}\tag{3}$$

Next, we consider a query made to a different  $c_{j \neq 1}$ , following a transcript of  $m$  queries  $\tau = [\tau_i]_{i=1}^m$  made to  $c_1$ . Without loss of generality, we will denote  $c_j = c_2$ . This  $(m + 1)^{th}$  query output takes the following form:

$$\hat{\tau}_{m+1} = \frac{\mathcal{Y}_{m+1} + \sum_{i=1}^m a_{m+1,i} \tau_i}{X_0 + c_2}$$

The above is the same as  $\tau_{m+1}$  (Equation 2) except the numerator is divided by  $(X_0 + c_2)$  instead of  $(X_0 + c_1)$ . Thus, we can rewrite using the same binary string notation (again, grouped by  $\mathcal{Y}_j$ ):

$$\begin{aligned}
\hat{\tau}_{m+1} &= \frac{\mathcal{Y}_{m+1} + \sum_{i=1}^m a_{m+1,i} \tau_i}{X_0 + c_2} = \frac{X_0 + c_1}{X_0 + c_2} \cdot \tau_{m+1} \\
&= \frac{X_0 + c_1}{X_0 + c_2} \left( \sum_{s=2^{i-1}}^{2^i-1} \Omega(s) \right) \\
&= \frac{X_0 + c_1}{X_0 + c_2} \left( \Omega(2^{i-1}) + \sum_{j=1}^{i-1} \sum_{s=2^{i-j-1}}^{2^{i-j}-1} \Omega(2^j * s + 2^{j-1}) \right) \\
&= \frac{\mathcal{Y}_{m+1}}{X_0 + c_2} + \sum_{j=1}^m \sum_{s=2^{m-j}}^{2^{m-j+1}-1} \frac{\mathcal{Y}_j \cdot \prod_{k=1}^{|\omega(2^j * s + 2^{j-1})|-1} a_{\omega(2^j * s + 2^{j-1})_k, \omega(2^j * s + 2^{j-1})_{k+1}}}{(X_0 + c_1)^{|\omega(2^j * s + 2^{j-1})|-1} (X_0 + c_2)}
\end{aligned}$$

The query output  $\hat{\tau}_{m+1}$  is a sum of terms with mixed  $(X_0 + c_1)$  and  $(X_0 + c_2)$  factors in the denominator. We will show that  $\hat{\tau}_{m+1}$  to  $\tau'_{m+1}$  of the form:

$$\tau'_{m+1} = \frac{\mathcal{Y}_{m+1} + \sum_{i=1}^m b_i \mathcal{Y}_i}{X_0 + c_2}, \quad (4)$$

for some set of coefficients  $b_1, \dots, b_m$ , such that the span of the query outputs is preserved:

*Claim.* We provide  $[b_i]_{i=1}^m$  to construct  $\tau'_{m+1}$  (Equation 4) such that

$$\text{Span}([\tau_1, \dots, \tau_m, \hat{\tau}_{m+1}]) = \text{Span}([\tau_1, \dots, \tau_m, \tau'_{m+1}]).$$

*Proof of Claim:* We choose  $[b_i]_{i=1}^m$  such that  $\hat{\tau}_{m+1} \in \text{Span}([\tau_1, \dots, \tau_m, \tau'_{m+1}])$ . This is sufficient to complete the claim that the two spans are equivalent. We solve the following system of equations:

$$\begin{aligned}
\hat{\tau}_{m+1} &= \alpha_{m+1} \tau'_{m+1} + \sum_{i=1}^m \alpha_i \tau_i \\
&= \sum_{i=1}^{m+1} \beta_i \frac{\mathcal{Y}_i}{X_0 + c_2} + \sum_{i=1}^m \alpha_i \tau_i
\end{aligned} \quad (5)$$

for unknowns  $\alpha_1, \dots, \alpha_{m+1}, b_1, \dots, b_m$ , or equivalently, when reformulated, for unknowns  $\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_{m+1}$ , where  $\beta_{m+1} = \alpha_{m+1}$  and  $\beta_{i \neq m+1} = \alpha_{m+1} b_i$ . Now consider the expanded Equation 5:

$$\begin{aligned}
&\frac{\mathcal{Y}_{m+1}}{X_0 + c_2} + \sum_{j=1}^m \sum_{s=2^{m-j}}^{2^{m-j+1}-1} \frac{\mathcal{Y}_j \cdot \prod_{k=1}^{|\omega(2^j * s + 2^{j-1})|-1} a_{\omega(2^j * s + 2^{j-1})_k, \omega(2^j * s + 2^{j-1})_{k+1}}}{(X_0 + c_1)^{|\omega(2^j * s + 2^{j-1})|-1} (X_0 + c_2)} \\
&= \sum_{i=1}^{m+1} \beta_i \frac{\mathcal{Y}_i}{X_0 + c_2} + \sum_{i=1}^m \alpha_i \left( \frac{\mathcal{Y}_i}{X_0 + c_1} + \sum_{j=1}^{i-1} \sum_{s=2^{i-j-1}}^{2^{i-j}-1} \frac{\mathcal{Y}_j \cdot \prod_{k=1}^{|\omega(2^j * s + 2^{j-1})|-1} a_{\omega(2^j * s + 2^{j-1})_k, \omega(2^j * s + 2^{j-1})_{k+1}}}{(X_0 + c_1)^{|\omega(2^j * s + 2^{j-1})|}} \right)
\end{aligned}$$

One approach for solving this equation is by solving each of the partial equations, equating the coefficients for a particular  $\mathcal{Y}_i$ . The partial equation for  $\mathcal{Y}_{m+1}$

is trivial and is easily solveable by setting  $\beta_{m+1} = 1$ :

$$\frac{\mathcal{Y}_{m+1}}{X_0 + c_2} = \beta_{m+1} \frac{\mathcal{Y}_{m+1}}{X_0 + c_2}$$

The other partial equations for  $1 \leq i \leq m$  are more complex. We refer to the partial equation for  $\mathcal{Y}_i$  as  $PE_i$ , and it is defined as follows:

$$\begin{aligned} & \sum_{s=2^{m-i}}^{2^{m-i+1}-1} \frac{\mathcal{Y}_i \cdot \prod_{k=1}^{|\omega(2^i * s + 2^{i-1})|-1} a_{\omega(2^i * s + 2^{i-1})_k, \omega(2^i * s + 2^{i-1})_{k+1}}}{(X_0 + c_1)^{|\omega(2^i * s + 2^{i-1})|-1} (X_0 + c_2)} \\ &= \frac{\beta_i \mathcal{Y}_i}{X_0 + c_2} + \frac{\alpha_i \mathcal{Y}_i}{X_0 + c_1} + \sum_{j=i+1}^m \sum_{s=2^{j-i-1}}^{2^j-1} \frac{\mathcal{Y}_i \cdot \alpha_j \cdot \prod_{k=1}^{|\omega(2^i * s + 2^{i-1})|-1} a_{\omega(2^i * s + 2^{i-1})_k, \omega(2^i * s + 2^{i-1})_{k+1}}}{(X_0 + c_1)^{|\omega(2^i * s + 2^{i-1})|}} \end{aligned}$$

Note that  $\mathcal{Y}_i$  can be canceled out in  $PE_i$ . We next show that there exists (and that we can solve for) a satisfying assignment of variables  $[\alpha_i]_{i=1}^m$  and  $[\beta_i]_{i=1}^{m+1}$  for Equation 5 through a strong induction argument on the satisfiability of the partial equations.

*Induction hypothesis.*  $H(i)$ : There exists a satisfying assignment of variables  $[\alpha_j]_{j=i}^m$  and  $[\beta_j]_{j=i}^{m+1}$  for the system of equations  $PE_m, PE_{m-1}, \dots, PE_i$ .

Note that  $H(1)$  implies that Equation 5 is satisfiable. We proceed by proving the base case,  $H(m)$ , and then proving the induction step showing that  $H(i+1) \Rightarrow H(i)$ .

*Base case.* We prove  $H(m)$ . Consider  $PE_m$ :

$$\frac{\beta_m}{X_0 + c_2} + \frac{\alpha_m}{X_0 + c_1} = \frac{a_{m+1,m}}{(X_0 + c_1)(X_0 + c_2)}$$

We solve  $PE_m$  as follows:

$$\begin{aligned} \beta_m(X_0 + c_1) + \alpha_m(X_0 + c_2) &= a_{m+1,m} \\ \beta_m + \alpha_m &= 0 \\ \beta_m c_1 + \alpha_m c_2 &= a_{m+1,m} \\ \beta_m &= \frac{a_{m+1,m}}{c_1 - c_2} \\ \alpha_m &= \frac{a_{m+1,m}}{c_2 - c_1} \end{aligned}$$

*Induction step.* We prove  $H(i+1) \Rightarrow H(i)$ . Consider the left hand side of  $PE_i$ . It contains terms for  $(m+1)$ -length binary strings, where  $*$  is a wildcard for either 0/1:

$$1 \parallel *^{m-i} \parallel 1 \parallel 0^{i-1}$$

Now consider the left hand side of  $PE_j$  for  $j > i$ . It contains terms for the following binary strings:

$$1 \parallel *^{m-j} \parallel 1 \parallel 0^{j-i} \parallel 0^{i-1}$$

We observe that if we flip the  $i^{\text{th}}$  least significant bit of the binary strings from  $PE_j$ , we obtain a subset of binary strings from  $PE_i$ . In fact, if we do this for all  $PE_j$  for  $m \geq j > i$ , we obtain all binary strings from  $PE_i$  except for  $1 \parallel 0^{m-i} \parallel 1 \parallel 0^{i-1}$ .

A similar cancellation occurs on the right hand side of the equations. The right hand side of  $PE_i$  contains terms for the  $m$ -length binary strings:

$$\star^{m-i} \parallel 1 \parallel 0^{i-1}$$

And the right hand side of  $PE_j$  for  $m \geq j > i$  contains terms for the  $m$ -length binary strings:

$$\star^{m-j} \parallel 1 \parallel 0^{j-i} \parallel 0^{i-1}$$

Again, we observe that if we flip the  $i^{\text{th}}$  least significant bit of binary strings for all  $PE_j$ , we obtain all binary strings from  $PE_i$  except for  $0^{m-i} \parallel 1 \parallel 0^{i-1}$ .

This leads us to the following approach. We sum up all  $PE_j$  for  $m \geq j > i$ , transforming each one appropriately to “flip the  $i^{\text{th}}$  bit” of all represented binary strings. This sum can then be subtracted from  $PE_i$  to cancel out all terms on the left hand side of the equation except for the term corresponding to  $1 \parallel 0^{m-i} \parallel 1 \parallel 0^{i-1}$ , and most terms on the right hand side of the equation except for the term corresponding to  $0^{m-i} \parallel 1 \parallel 0^{i-1}$  and a set of terms  $\beta_u (X_0 + c_1)^{m-i+1}$  in each  $PE_u$  that does not correspond to a binary string.

We create the summed equation:

$$\sum_{j=i+1}^m \frac{a_{j,i}}{X_0 + c_1} \cdot PE_j \quad (6)$$

The following equation is what remains after taking the difference of  $PE_i$  and summed Equation 6:

$$\frac{a_{m+1,i}}{(X_0 + c_1)(X_0 + c_2)} = \frac{\alpha_i}{X_0 + c_1} + \frac{\beta_i}{X_0 + c_2} + \sum_{j=i+1}^m \frac{\beta_j a_{j,i}}{(X_0 + c_1)(X_0 + c_2)}$$

By induction hypothesis  $H(i+1)$ , we have that there exists a satisfying assignment for  $\beta_j$  for  $j > i$ . We fix those variables, then solve the above equation for  $\alpha_i$  and  $\beta_i$ :

$$\beta_i (X_0 + c_1) + \alpha_i (X_0 + c_2) = a_{m+1,i} - \sum_{j=i+1}^m \beta_j a_{j,i}$$

$$\beta_i + \alpha_i = 0$$

$$\beta_i c_1 + \alpha_i c_2 = a_{m+1,i} - \sum_{j=i+1}^m \beta_j a_{j,i}$$

$$\beta_i = \frac{a_{m+1,i} - \sum_{j=i+1}^m \beta_j a_{j,i}}{c_1 - c_2}$$

$$\alpha_i = \frac{a_{m+1,i} - \sum_{j=i+1}^m \beta_j a_{j,i}}{c_2 - c_1}$$

This concludes proof of the induction hypothesis and of the claim.

**Generalizing transcript rewrite to all  $c_i$ .** Next, we argue that the above claim, which holds for a transcript of queries to  $c_1$  followed by one query to  $c_2$ , generalizes to a transcript that makes queries to arbitrary different  $c_i$  values. We assume the transcript up until this point is made up of elements that are “separated by  $c_j$ ”, i.e., each only have powers of a single  $(X_0 + c_j)$  in the denominator. Given this, we show that a new query to any  $c_v$  (taking in a linear combination of all previous elements of the transcript) can be rewritten so that it too only depends on powers of  $(X_0 + c_v)$ , such that the span of the full transcript is preserved.

As such, we assume the transcript has elements of the following form, where a transcript element  $\tau_{i,j}$  is the  $i^{\text{th}}$  query to  $c_j$  (following from Equation 2 and 3). Element  $\tau_{i,j}$  depends only on previous  $\tau_{1,j}, \dots, \tau_{i-1,j}$  and all only have powers of  $(X_0 + c_j)$  in the denominator:

$$\begin{aligned} \tau_{i,j} &= \frac{\mathcal{Y}_{i,j} + \sum_{u=1}^{i-1} a_{u,j} \tau_{u,j}}{X_0 + c_j} \\ &= \frac{\mathcal{Y}_{i,j}}{X_0 + c_j} + \sum_{u=1}^{i-1} \sum_{s=2^{i-u}-1}^{2^i-u-1} \frac{\mathcal{Y}_{u,j} \cdot \prod_{k=1}^{|\omega(2^u * s + 2^{u-1})|-1} a_{\omega(2^u * s + 2^{u-1})_k, \omega(2^u * s + 2^{u-1})_{k+1}}}{(X_0 + c_j)^{|\omega(2^u * s + 2^{u-1})|}} \end{aligned} \quad (7)$$

A new query to  $c_v$  will have the following output form, where  $q_j$  denotes the number of queries that have been made to  $c_j$ :

$$\tau'_{q_v+1,v} = \frac{\mathcal{Y}'_{q_v+1,v} + \sum_{j=1}^n \sum_{i=1}^{q_j} a'_{q_v+1,v,i,j} \tau_{i,j}}{X_0 + c_v}$$

We show that we can replace this output with  $\tau_{q_v+1,v}$  of the following form which matches the form from Equation 7:

$$\begin{aligned} \tau_{q_v+1,v} &= \frac{\mathcal{Y}'_{q_v+1,v} + \sum_{\substack{j=1 \\ j \neq v}}^n \sum_{i=1}^{q_j} b_{i,j} \mathcal{Y}_{i,j} + \sum_{i=1}^{q_v} a'_{q_v+1,v,i,v} \tau_{i,v}}{X_0 + c_v} \\ &= \frac{\mathcal{Y}_{q_v+1,v} + \sum_{i=1}^{q_v} a_{i,v} \tau_{i,v}}{X_0 + c_v} \\ \text{where } \mathcal{Y}_{q_v+1,v} &= \mathcal{Y}'_{q_v+1,v} + \sum_{\substack{j=1 \\ j \neq v}}^n \sum_{i=1}^{q_j} b_{i,j} \mathcal{Y}_{i,j} \\ a_{i,v} &= a'_{q_v+1,v,i,v} \end{aligned}$$

We want that the span of the new transcript is preserved by replacing the new output with the rewritten output, and thus we prove the following:

*Claim.* We provide  $[[b_i]_{i=1}^{q_j}]_{j=1}^n$  to construct  $\tau_{q_v+1,v}$  such that

$$\text{Span}\left(\left[[[\tau_{i,j}]_{i=1}^{q_j}]_{j=1}^n, \tau'_{q_v+1,v}\right]\right) = \text{Span}\left(\left[[[\tau_{i,j}]_{i=1}^{q_j}]_{j=1}^n, \tau_{q_v+1,v}\right]\right)$$

*Proof of Claim:* We prove the claim by providing  $\alpha_{i,j}$  and  $b_{i,j}$  values such that:

$$\tau'_{q_v+1,v} = \alpha_{q_v+1,v} \tau_{q_v+1,v} + \sum_{j=1}^n \sum_{i=1}^{q_j} \alpha_{i,j} \tau_{i,j}$$

This is implied by our previous claim which shows that for each  $j \neq v$ , we can find  $\alpha_{i,j}$  and  $b_{i,j}$  such that:

$$\sum_{i=1}^{q_j} \frac{a_{q_v+1,v,i,j} \tau_{i,j}}{X_0 + c_v} = \alpha_{q_v+1,v} \left( \sum_{i=1}^{q_j} \frac{b_{i,j} \mathcal{Y}_{i,j}}{X_0 + c_v} \right) + \sum_{i=1}^{q_j} \alpha_{i,j} \tau_{i,j}$$

This concludes the argument that any new query output can be rewritten to be in the form of Equation 7 in which it only has powers of the queried  $c_v$  in the denominator, while preserving the span of the query output transcript.

**Using rewritten transcript to show non-trivial winning element.** We will continue to use the quotient notation, however, recall that in the simulation of the adversary's environment (Section A.1), the rational fractions are multiplied by a least common multiple,  $\prod_{i=1}^n (X_0 + c_i)^{q_i}$ . The notions of independence we show for polynomial rational fractions hold true for the the polynomials once multiplied by the LCM as well.

The adversary is given initial group elements which we represent as polynomial rational fractions as follows:  $\hat{g} \mapsto 1$ ,  $\hat{X} \mapsto X_0$ ,  $[A_i]_{i=1}^m \mapsto [X_i]_{i=1}^m$ . It is evident that these polynomials are linearly-independent as they each include a different formal variable (with exception of  $\hat{g}$  which is the only element with a constant).

The adversary will also receive group elements from the SDH oracle. We will denote the polynomial rational fraction outputs of the  $q = \sum_{i=1}^n q_i$  queries to SDH as  $[\tau'_i]_{i=1}^q$ .

Lastly, the adversary will output a set of  $\ell$  winning group elements for a selected  $\gamma$ ,  $[Z_i, \alpha_i]_i^\ell$ , where  $\ell > q_\gamma$ . These elements are represented by the following polynomial rational fractions:  $\left[ \frac{X_{\alpha_i}}{X_0 + c_\gamma} \right]_{i=1}^\ell$ . These  $\ell$  elements are linearly-independent as they each include a different formal variable.

We argue that at least one element from  $[X_{\alpha_i}/(X_0 + c_\gamma)]_{i=1}^\ell$  is linearly independent from the initial elements and SDH elements given to the adversary. If so, then  $\mathcal{A}_{\text{uber}}$  wins if it guesses correctly and the proof is complete. In other words, we want that:

$$\exists i \in [1, \ell] : \frac{X_{\alpha_i}}{X_0 + c_\gamma} \notin \text{Span}([1, [X_i]_{i=0}^m, [\tau'_i]_{i=1}^q])$$

We use Claim A.2 to rewrite the transcript  $\tau'$  into a new transcript  $\tau$  that contains elements  $[[\tau_{i,j}]_{i=1}^{q_j}]_{j=1}^n$  of form defined in Equation 7. The transcript  $\tau$  is built element by element by repeatedly applying Claim A.2 to the next query

in  $\tau'$ . Ultimately we have that:

$$\begin{aligned} \text{Span}([\tau'_i]_{i=1}^q) &= \text{Span}([\tau_{i,j}]_{i=1}^{q_j}]_{j=1}^n) \\ \text{Span}([1, [X_i]_{i=0}^m, [\tau'_i]_{i=1}^q]) &= \text{Span}([1, [X_i]_{i=0}^m, [\tau_{i,j}]_{i=1}^{q_j}]_{j=1}^n]) \end{aligned}$$

So instead we will show

$$\exists i \in [1, \ell] : \frac{X_\alpha}{X_0 + c_\gamma} \notin \text{Span}([1, [X_i]_{i=0}^m, [\tau_{i,j}]_{i=1}^{q_j}]_{j=1}^n])$$

Next, consider a linear combination of the above that results in a winning element  $X_\alpha/(X_0 + c_\gamma)$  with linear combination coefficients  $r, s_i, t_{i,j}$ :

$$\begin{aligned} r \cdot 1 + \sum_{i=0}^m s_i X_i + \sum_{j=1}^n \sum_{i=1}^{q_j} t_{i,j} \tau_{i,j} &= \frac{X_\alpha}{X_0 + c_\gamma} \\ r \cdot 1 + \sum_{i=0}^m s_i X_i + \sum_{\substack{j=1 \\ j \neq \gamma}}^n \sum_{i=1}^{q_j} t_{i,j} \tau_{i,j} &= \frac{X_\alpha}{X_0 + c_\gamma} - \sum_{i=1}^{q_\gamma} t_{i,\gamma} \tau_{i,\gamma} \end{aligned} \quad (8)$$

Now if we multiply both side of the above Equation 8 by the LCM expression  $\prod_{i=1}^n (X_0 + c_i)^{q_i}$ , the quotients are removed and we have an equation of two polynomials. By the structure of  $\tau_{i,j}$  from Equation 7, we have that none of the  $\tau_{i,j}$  for  $j \neq \gamma$  have a  $(X_0 + c_\gamma)$  term in the denominator. Thus, the left hand side polynomial has a factor of  $(X_0 + c_\gamma)^{q_\gamma}$ .

On the right hand side, because every  $\tau_{i,\gamma}$  term has  $(X_0 + c_\gamma)$  in the denominator,  $(X_0 + c_\gamma)^{q_\gamma}$  does not divide the right hand side polynomial. This implies that the two polynomials cannot be equal unless they are the zero polynomial, which is only possible if  $r, [s_i]_{i=1}^m$  coefficients are equal to 0.

Thus, if  $X_\alpha/(X_0 + c_\gamma) \in \text{Span}([1, [X_i]_{i=0}^m, [\tau_{i,j}]_{i=1}^{q_j}]_{j=1}^n])$ , then it must be:

$$\sum_{i=1}^{q_\gamma} t_{i,\gamma} \tau_{i,\gamma} = \frac{X_\alpha}{X_0 + c_\gamma}$$

However, since there are only  $q_\gamma < \ell$   $[\tau_{i,\gamma}]_{i=1}^{q_\gamma}$  terms, they can at most generate a  $q_\gamma$ -dimension space. Since the  $\ell$  winning elements are linearly-independent and generate a  $\ell$ -dimension space, it is not possible that they can all be generated from a linear combination of  $[\tau_{i,\gamma}]_{i=1}^{q_\gamma}$ . This concludes the proof.

## B Security Proofs for 3H

For ease of reference, we restate the theorem here:

**Theorem 2.** *Let  $\mathcal{A}_{\text{prf}}$  be a P-model POPRF adversary against 3H with query budget  $(m, n, q_E, \vec{q}, q_{H_1}, q_{H_2}, q_{H_3}, q_{H_4})$ . Then we give a  $H_4$ -model adversary  $\mathcal{A}_{\text{zk}}$  and adversary  $\mathcal{A}_{\text{sdhi}}$  such that*

$$\text{Adv}_{3\text{H}, \mathcal{S}[\mathcal{S}_\Sigma], \text{P}, \mathcal{A}_{\text{prf}}}^{\text{po-prf}}(\lambda) \leq \text{Adv}_{\Sigma, \mathcal{R}, \mathcal{H}_4, \mathcal{S}_\Sigma, \mathcal{A}_{\text{zk}}}^{\text{zk}}(\lambda) + \text{Adv}_{\text{GGen}, \mathcal{A}_{\text{sdhi}}}^{(m,n)\text{-om-gap-sdhi}}(\lambda) + \frac{n^2}{p},$$

$\underline{\text{S.Init}(\lambda, pp)}$ $R_1 \leftarrow [\cdot]; R_2 \leftarrow [\cdot]; R_3 \leftarrow [\cdot]$ $sk \leftarrow \mathbb{Z}_p$ $st_\Sigma \leftarrow \text{S}_\Sigma.\text{Init}(pp)$ $st_5 \leftarrow (pp, sk, R_1, R_2, R_3, st_\Sigma)$ $\text{Return}(st_5, g^{sk})$ $\underline{\text{S.BlindEv}^{\text{EV}}(t, req : (pp, sk, R_1, R_2, R_3, st_\Sigma))}$ $h \leftarrow \text{S.H}_3^{\text{EV}}(t : st_5)$ $B \leftarrow req; B' \leftarrow B^{1/(sk+h)}$ $\pi \leftarrow \text{S}_\Sigma.\text{Prove}((g, g^{sk+h}, B', B) : st_\Sigma)$ $\text{Return}(B', \pi)$	$\underline{\text{S.H}_1^{\text{LIMEV}}(x : (pp, sk, R_1, R_2, R_3, st_\Sigma))}$ $\text{If } x \notin R_1 \text{ then } R_1[x] \leftarrow \mathbb{G}$ $\text{Return } R_1[x]$ $\underline{\text{S.H}_2^{\text{LIMEV}}(t, x, Y : (pp, sk, R_1, R_2, R_3, st_\Sigma))}$ $\text{If } t \notin R_3 \text{ then } R_3[x] \leftarrow \mathbb{Z}_p$ $\text{If } x \notin R_1 \text{ then } R_1[x] \leftarrow \mathbb{G}$ $\text{If } (t, x, Y) \notin R_2 \text{ then}$ $\quad \text{If } Y = R_1[x]^{1/(sk+R_3[t])} \text{ then } R_2[t, x, Y] \leftarrow \text{LIMEV}(t, x)$ $\quad \text{Else } R_2[t, x, Y] \leftarrow \{0, 1\}^\lambda$ $\text{Return } R_2[t, x, Y]$ $\underline{\text{S.H}_3^{\text{LIMEV}}(x : (pp, sk, R_1, R_2, R_3, st_\Sigma))}$ $\text{If } x \notin R_3 \text{ then } R_3[x] \leftarrow \mathbb{Z}_p$ $\text{Return } R_3[x]$ $\underline{\text{S.H}_4^{\text{LIMEV}}(x : (pp, sk, R_1, R_2, R_3, st_\Sigma))}$ $\text{Return } \text{S}_\Sigma.\text{H}(x : st_\Sigma)$
--	--

Fig. 10: Simulator  $\text{S}[\text{S}_\Sigma]$  for PRF security of  $\text{3H}$  where  $\text{S}_\Sigma$  is the zero knowledge simulator for  $\Sigma_{\mathcal{R}}$ .

where  $\text{S}$  is the simulator defined in Figure 10 that makes use of NIZK simulator  $\text{S}_\Sigma$ . Adversary  $\mathcal{A}_{sk}$  makes  $q_{H_4}$  queries to its random oracle and  $\mathcal{A}_{\text{sdhi}}$  has query budget  $(\vec{q}, q_{H_2})$ . Further,  $T(\mathcal{A}_{\text{prf}}) \approx T(\mathcal{A}_{zk}) \approx T(\mathcal{A}_{\text{sdhi}})$ .

*Proof.* We bound the advantage of  $\mathcal{A}_{\text{prf}}$  by bounding the advantage of each of a series of game hops. We define  $G_0 = \text{POPRF}_{\text{3H,P,S}}^{\text{A}_{\text{po-prf},1}}(\lambda)$  and intermediate games  $G_1, G_2, G_3, G_4, G_5$  to gradually transform the view of the adversary until  $G_5 = \text{POPRF}_{\text{3H,P,S}}^{\text{A}_{\text{po-prf},0}}(\lambda)$ . The most important games,  $G_2 - G_4$  are shown in Figure 11). There we use tables  $R_1, R_2, R_3, R_4$  for RO simulation, and we restrict collisions in  $R_3$  (as described below) by sampling from  $\mathbb{Z}_p$  without replacement, which we do by defining  $\mathbb{Z}_p \setminus R_3$  to be the set of points in  $\mathbb{Z}_p$  that do not appear in any entries of  $R_3$ .

The advantage bound follows from the following claims which we will justify:

- (1)  $|\Pr[G_0 = 1] - \Pr[G_1 = 1]| \leq \frac{n(n-1)}{2p}$
- (2)  $|\Pr[G_1 = 1] - \Pr[G_2 = 1]| = \text{Adv}_{\Sigma_{\mathcal{R}}, \mathcal{R}, \text{RO}_4, \mathcal{A}_{zk}}^{\text{zk}}(\lambda)$
- (3)  $|\Pr[G_2 = 1] - \Pr[G_3 = 1]| = 0$
- (4)  $|\Pr[G_3 = 1] - \Pr[G_4 = 1]| \leq \text{Adv}_{\text{GGen}, \mathcal{A}_{\text{sdhi}}}^{(q_{H_1}, q_{H_2})\text{-om-gap-sdhi}}(\lambda)$
- (5)  $|\Pr[G_4 = 1] - \Pr[G_5 = 1]| \leq \frac{n(n-1)}{2p}$

*Claim 1:* We first transition to a game  $G_1$  in which we disallow collisions in the output of  $\text{RO}_3$ , i.e., we replace sampling from  $\mathbb{Z}_p$  for new range values with sampling without replacement from  $\mathbb{Z}_p$ . Recall that at most  $n$  unique  $t$  values are queried by  $\mathcal{A}_{\text{prf}}$  in the course of the game, and so a standard birthday analysis establishes the claim's upper bound of  $n(n-1)/2p$ .

*Claim 2:* In  $G_2$ , the proof generated in  $\text{BLINDEV}$  is simulated along with the corresponding random oracle  $\text{RO}_4$ . Since this is the only change between  $G_1$  and

<p>Games <math>G_2, G_3, G_4</math></p> <p><math>R_1 \leftarrow [ ]; R_2 \leftarrow [ ]; R_3 \leftarrow [ ]</math>  <math>R \leftarrow [ ]</math>  <math>i_t \leftarrow 0 ; j_t \leftarrow 0 ; \mathbf{bad} \leftarrow 0</math>  <math>pp \leftarrow \text{3H.Setup}(\lambda)</math>  <math>sk \leftarrow \mathbb{Z}_p ; pk \leftarrow g^{sk}</math>  <math>st_\Sigma \leftarrow \text{S}_\Sigma.\text{Init}(pp)</math>  <math>b' \leftarrow \mathcal{A}_{\text{prf}}^{\text{P, EV, BLINDEV}}(pp, pk)</math>  Return <math>b'</math></p>	<p>Oracle <math>\text{EV}(t, x)</math></p> <p>If <math>t \notin R_3</math> then <math>R_3[t] \leftarrow \mathbb{Z}_p \setminus R_3</math>  If <math>x \notin R_1</math> then <math>R_1[x] \leftarrow \mathbb{G}</math>  <math>Y \leftarrow R_1[x]^{1/(sk+R_3[t])}</math>  If <math>(t, x, Y) \notin R_2</math> then <math>R_2[t, x, Y] \leftarrow \{0, 1\}^\lambda</math>  <math>Z \leftarrow R_2[t, x, Y]</math>  <b>If <math>(t, x) \notin R</math> then</b>  <math>R[t, x] \leftarrow \{0, 1\}^\lambda</math>  <math>Z \leftarrow R[t, x]</math>  Return <math>Z</math></p> <p>Oracle <math>\text{BLINDEV}(t, \text{req})</math></p> <p>If <math>t \notin R_3</math> then <math>R_3[t] \leftarrow \mathbb{Z}_p \setminus R_3</math>  <math>B \leftarrow \text{req} ; B' \leftarrow B^{1/(sk+R_3[t])}</math>  <math>j_t \leftarrow j_t + 1</math>  <math>\pi \leftarrow \text{S}_\Sigma.\text{Prove}(g, g^{sk+h}, B', B) : st_\Sigma</math>  Return <math>(B', \pi)</math></p>	<p>Oracle <math>\text{H}_1(x)</math></p> <p>If <math>x \notin R_1</math> then <math>R_1[x] \leftarrow \mathbb{G}</math>  Return <math>R_1[x]</math></p> <p>Oracle <math>\text{H}_2(t, x, Y)</math></p> <p><b>If <math>t \notin R_3</math> then <math>R_3[t] \leftarrow \mathbb{Z}_p \setminus R_3</math></b>  <b>If <math>x \notin R_1</math> then <math>R_1[x] \leftarrow \mathbb{G}</math></b>  If <math>(t, x, Y) \notin R_2</math> then  <math>R_2[t, x, Y] \leftarrow \{0, 1\}^\lambda</math>  <b>If <math>Y = R_1[x]^{1/(sk+R_3[t])}</math> then</b>  <math>i_t \leftarrow i_t + 1</math>  <b>If <math>i_t &gt; j_t</math> then <math>\mathbf{bad} \leftarrow 1 ; \text{Return } \perp</math></b>  <b>If <math>(t, x) \notin R</math> then <math>R[t, x] \leftarrow \{0, 1\}^\lambda</math></b>  <math>R_2[t, x, Y] \leftarrow R[t, x]</math>  <b>Else <math>R_2[t, x, Y] \leftarrow \{0, 1\}^\lambda</math></b>  Return <math>R_2[t, x, Y]</math></p> <p>Oracle <math>\text{H}_3(x)</math></p> <p>If <math>x \notin R_3</math> then <math>R_3[x] \leftarrow \mathbb{Z}_p \setminus R_3</math>  Return <math>R_3[x]</math></p> <p>Oracle <math>\text{H}_4(x)</math></p> Return $\text{S}_\Sigma.H(x : st_\Sigma)$
--	--	---

Fig. 11: The key game transitions used in pseudorandomness security for 3H. Greyed highlighted statements are only included in  $G_2$ , blue highlighted statements in  $G_3$  and  $G_4$ , and boxed statements only in  $G_4$ .

$G_2$ , we can equate the distinguishing advantage between the two games to that of the zero-knowledge security game of  $\Sigma_{\mathcal{R}}$ . We construct  $\mathcal{A}_{zk}$  that runs  $G_1$  and generates proofs using its  $\text{PROVE}$  oracle and responds to queries to  $\text{RO}_4$  using its own  $\text{PRIM}$  oracle.

*Claim 3:* In  $G_3$ , the  $\text{EV}$  oracle generates outputs independently from the random oracles and stores its choices in table  $R$ . For consistency, it must be that for  $t, x$  and  $Y = R_1[x]^{1/(sk+R_3[t])}$ , the random output stored in  $R$  is the same as the one stored in  $R_2$  for responding to  $\text{RO}_2$  queries of  $(t, x, Y)$ .  $G_3$  checks if this is the case in  $\text{RO}_2$ , and if  $(t, x, Y)$  are of the above form, it repairs  $R$  and  $R_2$  to be consistent. Thus, from the adversary's perspective, there is no change between  $G_2$  and  $G_3$ .

*Claim 4:* In  $G_4$ , the repair between  $R$  and  $R_2$  in  $\text{RO}_2$  only occurs if there have been more calls to  $\text{BLINDEV}$  on  $t$  than calls of valid  $(t, x, Y)$  tuples to  $\text{RO}_2$ . Otherwise, a bad flag is set and the oracle returns  $\perp$  to the adversary. This matches the functionality of  $\text{S}$  in  $\text{POPFRF}_{3\text{H}, \text{P}, \text{S}}^{\text{A}_{\text{po-prf}}, 0}(\lambda)$  since the simulator is restricted to not run  $\text{LIMEV}$  on  $t$  more than calls made to  $\text{BLINDEV}$  on  $t$ . By an identical-until-bad argument via the fundamental lemma of game playing [BR06],

$$|\Pr[G_3 = 1] - \Pr[G_4 = 1]| \leq \Pr[\mathbf{bad} = 1]$$

where  $\mathbf{bad} = 1$  is the event that  $\mathbf{bad}$  is sent in game  $G_4$ . We bound the probability of this event by the advantage of an adversary  $\mathcal{A}_{\text{sdhi}}$ , i.e., if  $\mathbf{bad}$  is set,  $\mathcal{A}_{\text{sdhi}}$  wins the  $(m, n)$ -OM-GAP-SDHI game.

Adversary  $\mathcal{A}_{\text{sdhi}} = (\mathcal{A}_1, \mathcal{A}_2)$  (shown in Figure 12) runs  $G_4$  with the help of the  $(m, n)$ -OM-GAP-SDHI game. The  $[Y_i]_i^m$  group elements are used for the values returned by  $\text{RO}_1$ . (This is often called “programming”  $\text{RO}_1$ .) The  $n$  values

<p><u><math>\mathcal{A}_1(p, \mathbb{G})</math></u>  <math>C \leftarrow \emptyset</math>  For <math>i = 1</math> to <math>n</math>:  <math>c_i \leftarrow \mathbb{Z}_p \setminus C</math>; <math>C \leftarrow C \cup \{c_i\}</math>  <math>st_{\mathcal{A}} \leftarrow (p, \mathbb{G}, [c_i]_i^n)</math>  Return <math>(st_{\mathcal{A}}, [c_i]_i^n)</math>  <u><math>\mathcal{A}_2^{\text{SDH, SDDH}}(g, pk, [Y_i]_i^m : (p, \mathbb{G}, [c_i]_i^n)</math></u>  <math>i \leftarrow 1</math>; <math>j_* \leftarrow 0</math>; <math>\ell \leftarrow 1</math>  <math>K \leftarrow []</math>; <math>\hat{Z} \leftarrow []</math>; <math>\gamma \leftarrow \perp</math>  <math>R_1 \leftarrow []</math>; <math>R_2 \leftarrow []</math>; <math>R_3 \leftarrow []</math>; <math>R \leftarrow []</math>  <math>pp \leftarrow (p, g, \mathbb{G})</math>  <math>st_{\Sigma} \leftarrow \mathcal{S}_{\Sigma}.\text{Init}(pp)</math>  <math>b' \leftarrow \mathcal{A}_{\text{prf}}^{\text{P, EV, BLINDEV}}(pp, pk)</math>  Return <math>(\gamma, \hat{Z}[\gamma])</math>  <u>Oracle EV(<math>t, x</math>)</u>  If <math>(t, x) \notin R</math> then <math>R[t, x] \leftarrow \{0, 1\}^\lambda</math>  <math>Z \leftarrow R[t, x]</math>  Return <math>Z</math>  <u>Oracle BLINDEV(<math>t, req</math>)</u>  If <math>t \notin R_3</math> then <math>R_3[t] \leftarrow \ell</math>; <math>\ell \leftarrow \ell + 1</math>  <math>j_t \leftarrow j_t + 1</math>  <math>B \leftarrow req</math>  <math>B' \leftarrow \text{SDH}(B, R_3[t])</math>  <math>\pi \leftarrow \mathcal{S}_{\Sigma}.\text{Prove}((g, X \cdot g^{c_{R_3[t]}}), B', B) : st_{\Sigma})</math>  Return <math>(B', \pi)</math></p>	<p><u>Oracle <math>\mathbf{H}_1(x)</math></u>  If <math>x \notin R_1</math> then  <math>R_1[x] \leftarrow Y_i</math>  <math>K[x] \leftarrow i</math>; <math>i \leftarrow i + 1</math>  Return <math>R_1[x]</math>  <u>Oracle <math>\mathbf{H}_2(t, x, Y)</math></u>  If <math>t \notin R_3</math> then  <math>R_3[t] \leftarrow \ell</math>; <math>\ell \leftarrow \ell + 1</math>  If <math>x \notin R_1</math> then  <math>R_1[x] \leftarrow Y_i</math>  <math>K[x] \leftarrow i</math>; <math>i \leftarrow i + 1</math>  If <math>(t, x, Y) \notin R_2</math> then  If SDDH(<math>Y, R_1[x], R_3[t]</math>) then  <math>\hat{Z}[R_3[t]] \leftarrow \hat{Z}[R_3[t]] \parallel (Y, K[x])</math>  If <math> \hat{Z}[R_3[t]]  &gt; j_t</math> then  <math>\gamma \leftarrow R_3[t]</math>; abort <math>\mathcal{A}_{\text{prf}}</math>  If <math>x \notin R</math> then <math>R[x] \leftarrow \{0, 1\}^\lambda</math>  <math>R_2[t, x, Y] \leftarrow R[x]</math>  Else <math>R_2[t, x, Y] \leftarrow \{0, 1\}^\lambda</math>  Return <math>R_2[t, x, Y]</math>  <u>Oracle <math>\mathbf{H}_3(x)</math></u>  If <math>x \notin R_3</math> then  <math>R_3[x] \leftarrow \ell</math>; <math>\ell \leftarrow \ell + 1</math>  Return <math>c_{R_3[x]}</math>  <u>Oracle <math>\mathbf{H}_4(x)</math></u>  Return <math>\mathcal{S}_{\Sigma}.\mathbf{H}(x : st_{\Sigma})</math></p>
---	--

Fig. 12: Adversary  $\mathcal{A}_{\text{s d h i}} = (\mathcal{A}_1, \mathcal{A}_2)$  used in POPRF security proof of 3H.

$[c_i]_i^n$  from  $\mathbb{Z}_p$  output by  $\mathcal{A}_1$  for use in the strong Diffie-Hellman queries are used for the return values from  $\text{RO}_3$ . By assumption on the query budget for  $\mathcal{A}_{\text{prf}}$ , the number of  $Y_i$  values and the number of  $c_i$  values are sufficient for simulating the queries made by  $\mathcal{A}_{\text{prf}}$ . And since both of these sets of values are chosen at random, they have the same distribution as the random oracle responses in  $\mathbb{G}_4$ . Queries to  $\text{BLINDEV}$  are answered by computing the strong Diffie-Hellman evaluation with the appropriate  $c_i$  value using  $\text{SDH}$ . Note that  $\mathcal{A}_{\text{s d h i}}$  has query budget  $\bar{q}$  because by assumption  $\mathcal{A}_{\text{prf}}$  queries at most  $\bar{q}_1$  times to  $\text{BLINDEV}$  on the first value  $t_1$  queried in the course of the game, at most  $\bar{q}_2$  times for the second value  $t_2$  queried in the course of the game, and so on.

In  $\text{RO}_2$ , the form of  $(t, x, Y)$  is checked using  $\text{SDDH}$  to determine if a repair between  $R$  and  $R_2$  needs to be performed. However, before the repair is done, if there have been more valid  $(t, x, Y)$  tuples queried to  $\text{PRIM}_2$  than queries to  $\text{BLINDEV}$ ,  $\mathcal{A}_{\text{s d h i}}$  then halts execution of  $\mathcal{A}_{\text{prf}}$  and concludes by outputting  $(\gamma, \hat{Z}[\gamma])$ . In this case, the adversary has found “one more” valid strong Diffie-Hellman tuple (one corresponding to each valid  $(t, x, Y)$  tuple since  $Y = Y_i^{1/(sk+c_j)}$  for some  $Y_i, c_j$ ) than calls made to  $\text{SDH}$ . The adversary therefore wins the  $(m, n)$ -OM-GAP-SDHI game.

*Claim 5:* The final game transition restores the possibility of collisions to  $\text{RO}_3$ , and a birthday analysis gives the upper bound on this transition.

■

## C Security with Restricted Tag Space

**Corollary 3.** *For any adversary  $\mathcal{A}_{\text{prf}}$  against the partially-oblivious pseudorandomness of  $3\text{H}$  with restricted tag space of size  $|\mathbb{T}|$ , we give adversaries  $\mathcal{A}_{\text{zk}}$  and  $\mathcal{A}_{\text{sdhi}}$  such that*

$$\text{Adv}_{3\text{H}, \mathbb{P}, \mathbb{S}[\mathbb{S}_\Sigma], \mathcal{A}_{\text{prf}}}^{\text{po-prf}}(\lambda) \leq \text{Adv}_{\Sigma_{\mathcal{R}}, \mathcal{R}, \text{RO}_4, \mathbb{S}_\Sigma, \mathcal{A}_{\text{zk}}}^{\text{zk}}(\lambda) + \text{Adv}_{\text{GGen}, \mathcal{A}_{\text{sdhi}}}^{(q_{\text{H}_1}, |\mathbb{T}|)\text{-om-gap-sdhi}}(\lambda),$$

where  $\mathbb{S}$  is the simulator defined in Figure 10, the ideal primitive  $\mathbb{P} = \text{RO}_1 \times \text{RO}_2 \times \text{RO}_3 \times \text{RO}_4$  for random oracles over  $\text{RO}_1 : * \rightarrow \mathbb{G}$ ,  $\text{RO}_2 : ** \times \mathbb{G} \rightarrow \{0, 1\}^\lambda$ ,  $\text{RO}_3 : * \rightarrow \mathbb{Z}_p$ ,  $\text{RO}_4 : \mathbb{G}^6 \rightarrow \mathbb{Z}_p$  for  $(p, \mathbb{G})$  determined by  $\text{GGen}(\lambda)$ . The running time  $T(\mathcal{A}_{\text{prf}}) \approx T(\mathcal{A}_{\text{zk}}) \approx T(\mathcal{A}_{\text{sdhi}})$  and  $\mathcal{A}_{\text{prf}}$  makes at most  $q_{\text{H}_1}$  queries to  $\text{RO}_1$ .

*Proof.* The proof follows the same as above except only queries to  $\text{PRIM}_3$  that fall within the tag space  $\mathbb{T}$  need to be programmed with a strong Diffie-Hellman constant  $c$ ; otherwise, a random value can be sampled.