

Smooth NIZK Arguments

Charanjit S. Jutla¹ and Arnab Roy²

¹ IBM T. J. Watson Research Center, NY, USA. Email: csjutla@us.ibm.com

² Fujitsu Laboratories of America, CA, USA. Email: aroy@us.fujitsu.com

Abstract. We introduce a novel notion of smooth (-verifier) non-interactive zero-knowledge proofs (NIZK) which parallels the familiar notion of smooth projective hash functions (SPHF). We also show that the single group element quasi-adaptive NIZK (QA-NIZK) of Jutla and Roy (CRYPTO 2014) and Kiltz and Wee (EuroCrypt 2015) for linear subspaces can be easily extended to be computationally smooth. One important distinction of the new notion from SPHFs is that in a smooth NIZK the public evaluation of the hash on a language member using the projection key does not require the witness of the language member, but instead just requires its NIZK proof.

This has the remarkable consequence that if one replaces the traditionally employed SPHFs with the novel smooth QA-NIZK in the Gennaro-Lindell paradigm of designing universally-composable password-authenticated key-exchange (UC-PAKE) protocols, one gets highly efficient UC-PAKE protocols that are secure even under adaptive corruption. This simpler and modular design methodology allows us to give the first single-round *asymmetric* UC-PAKE protocol, which is also secure under adaptive corruption in the erasure model. Previously, all asymmetric UC-PAKE protocols required at least two rounds. In fact, our protocol just requires each party to send a single message asynchronously. In addition, the protocol has short messages, with each party sending only four group elements. Moreover, the server password file needs to store only one group element per client. The protocol employs asymmetric bilinear pairing groups and is proven secure in the (limited programmability) random oracle model and under the standard bilinear pairing assumption SXDH.

Keywords: QA-NIZK, Bilinear pairings, SXDH, MDDH, UC-PAKE, online attack, server compromise, dual-system.

1 Introduction

Ever since the remarkably efficient non-interactive zero knowledge (NIZK) proofs [BFM88] for algebraic statements were developed by Groth and Sahai (GS-NIZK) [GS12], there have been significant efficiency improvements and innovations in the construction of cryptographic protocols. Jutla and Roy [JR13, JR14] and Libert, Peters, Joye and Yung [LPJY14] further improved the efficiency of algebraic NIZK proofs, culminating in *constant* size NIZK proofs for linear subspaces, independent of the number of equations and witnesses. This efficiency

improvement came in the weaker Quasi-Adaptive setting [JR13], which nevertheless proved sufficient for many applications.

Quasi-adaptive NIZK (QA-NIZK) proofs were further extended to provide simulation soundness [LPJY14, KW15] and dual-system simulation soundness [JR15], thus lending applicability to many more applications, such as structure preserving signatures, password authenticated key exchange in the UC model, and keyed homomorphic CCA-secure encryption.

In this paper, we further extend (QA-)NIZK proofs to provide an additional property called *smooth soundness*. The idea is to force the verification step to consist of computing hashes in two different ways and comparing the result. To this end, the verifier is split into three algorithms: a randomized hash-key generation algorithm, a public hashing algorithm and a private hashing algorithm. The verification step starts off by generating two hash-keys, the private key and the projection key. Next, the setting allows computation of a private hash given the private hash-key and the word, and computation of the public hash using the projection key and just a QA-NIZK proof for the word - the witness for the word is *not* required. Completeness states that the private hash is equal to the public hash for a language member and correct (QA-)NIZK proof. Computational soundness states that it is hard to come up with a proof such that a non-language word passes the same equality check. The new *smoothness* property states that for any non-language word, the private hash algorithm outputs a value (computationally) indistinguishable from uniformly random, even when the projection key is given to the adversary.

Comparison with SPHF's. The new primitive is modeled after smooth projective hash functions (SPHF [CS02]). An SPHF also generates private and projection hash-keys and defines a private hash and a public hash. Further, similar properties hold where (1) for a member word, private hash equals public hash, (2) for a non-member word, private hash is uniformly random (even given projection hash-key). The crucial difference is that³, whereas the SPHF public hash computation requires a witness of the member word, the smooth (QA-)NIZK public hash requires only a NIZK proof of the word. This allows for hiding of the witness, even when computing using the projection hash-key. In contrast, trapdoor-SPHF's as introduced by [BBC⁺13] allow a simulation world to have a trapdoor to evaluate a hash over a word without a witness and using only projection hash-key. As shown in Fig. 1, where trapdoor-SPHF's are compared with smooth QA-NIZK, their notion does not allow “erasure” of the witness w in the real world (if only projection hash-key is available).

While Fig. 1 is self-explanatory, it brings up an interesting alternative interpretation of smooth (QA-)NIZK. In the common reference string (CRS) setting, one can have a *composite*-SPHF, which is composed of two SPHF's: the first SPHF's projection hash-key is published in the CRS, that enables one to compute an intermediate-hash using the witness w of a language member x . At this

³ On the other hand, our constructions only allow computational smooth-soundness, while for SPHF's these properties hold information-theoretically.

point, the witness can be erased. Next, the real projection hash-key \mathbf{hp} is revealed (possibly, generated by another party along with the private hash-key \mathbf{hk}). Then, a final hash on input x can be computed using \mathbf{hp} and the intermediate-hash. If x is a language member then the final hash is same as that computed from x using \mathbf{hk} . The question then arises as to why the intermediate-hash is depicted as a (QA)-NIZK in Fig 1. However, we note that this first SPHF is already publicly verifiable, as the private hash-key of the first SPHF is generated by the CRS generator, and if it has to be used in any form in a private hash evaluation (and in the real world) it must be publicly available in the CRS. Indeed, in our construction the private hash-key k is given in the CRS as a commitment to k (this interpretation of QA-NIZK as a publicly-verifiable SPHF was given in [KW15]).

We remark that this interpretation of smooth (QA)-NIZK is similar to constructions of structure-preserving SPHF in [BC16a]. In that work, the intermediate -hash is a GS-NIZK proof in the commit and prove framework [GS12], and hence the (second) private-hash takes the commitments also as input. Although the first hash is not a SPHF, their construction can still be viewed as a smooth NIZK (as per our definition⁴). Since their construction also works only for linear subspaces, our smooth QA-NIZK construction turns out to be more efficient, namely that no commitments need to be given for private hash computation.

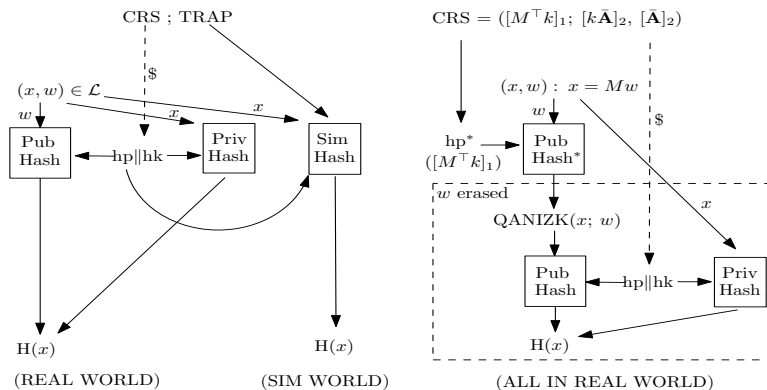


Fig. 1. Trapdoor-SPHF [BBC⁺13] vs Smooth QANIZK

Our Construction. In this work, we show that the single group element QA-NIZK arguments of [JR14, KW15] can be easily extended to be smooth. As a first application, we show that in the Gennaro-Lindell paradigm of designing

⁴ Since their construction uses the commit and prove paradigm of GS-NIZK proofs, their (composite-) construction is a smooth NIZK with a small tweak: they obtain information-theoretic smoothness, but computational zero-knowledge.

universally-composable password-authenticated key exchange (UC-PAKE) protocols, if one replaces the traditional SPHF with the novel smooth QA-NIZK, then one gets highly efficient single-round UC-PAKE protocols which are adaptively secure in the erasure model. At a high level, the UC simulator must emulate each party’s outgoing commitment to the password, without knowing the password. This is not difficult, as one can use ElGamal encryption to achieve a hiding commitment. However, if the party is corrupted after its message has been sent, the simulator is at a loss to produce a witness which each party must retain to eventually compute the SPHF public hash. In our new protocol, the parties need only save the QA-NIZK and not the witness, as that suffices to compute the public hash. This nicely captures the main idea behind the single-round adaptively secure UC-PAKE of [JR15]. In this work, the novel abstraction further allows us to obtain a single-round adaptively secure *asymmetric* UC-PAKE, which is a much more difficult notion to understand even from a definitional perspective, let alone constructing one.

(Asymmetric) Password-Authenticated Key-Exchange. The problem of setting up a secure channel between two parties that only share a human-memorizable password (or a low-entropy secret) was first studied by Bellare and Merritt [BM92], and later by Jiang and Gong [JG04]. Since then, this problem has been extensively studied and is called the *password-authenticated key-exchange* (PAKE) problem. One of the main challenges in designing such protocols is the intricacy in the natural security definition which requires that the protocol transcripts cannot be used to launch *offline dictionary attacks*. While an adversary can clearly try to guess the (low-entropy) password and impersonate one of the parties, its advantage from the fact that the password is of low entropy should be limited to such *online* impersonation attacks.

In a subsequent paper, Bellare and Merritt [BM93] also considered a stronger model of server compromise such that if a server’s password file is revealed to the adversary it cannot directly impersonate a client (cf. if the password was stored in the raw at the server). The adversary should be able to impersonate the client only if it succeeds in an offline dictionary attack on the revealed server password file. Clearly, this requires that the server does not store the password as it is (or in some reversibly-encrypted form), and protocols satisfying this stronger security requirement are referred to as *asymmetric PAKE* protocols.

Canetti et al [CHK⁺05] also considered designing (symmetric) UC-PAKE protocols in the universally-composable (UC) framework [Can01]. One of their main contributions was the definition of a natural UC-PAKE ideal functionality ($\mathcal{F}_{\text{PAKE}}$). Gentry et al [GMR06] extended the functionality of symmetric UC-PAKE [CHK⁺05] to the asymmetric setting ($\mathcal{F}_{\text{apwKE}}$) and gave a general method of extending any symmetric UC-PAKE protocol to an asymmetric UC-PAKE protocol (from now on referred to as UC-APAKE). Their general method adds an additional round to the UC-PAKE protocol.

Our Contributions. In this paper, we give the first single-round UC-APAKE protocol (realizing $\mathcal{F}_{\text{apwKE}}$). In fact, both parties just send a single message asyn-

chronously. The protocol is realized in the (limited programmability⁵) random-oracle (RO) [BR93] hybrid-model under standard static assumptions for bilinear groups, namely SXDH [BBS04] and the general MDDH [EHK⁺13] assumption. Our protocol is also secure against adaptive corruption (in the erasure model) and is very succinct, with each message consisting of only four group elements. Moreover, for each client the server need store only *one* group elements as a “password hash”. Many non-UC asymmetric PAKE protocols are at least two rounds [HK98, BPR00, BMP00, Mac01, Boy09]. Benhamouda and Pointcheval [BP13] proposed the first single round asymmetric PAKE protocol, but in a game-based model built on the BPR model [BPR00].

The first single-round UC-secure *symmetric* PAKE protocol was given in [KV11] (using bilinear pairings), which was then further improved (in the number of group elements) in subsequent papers [JR12, BBC⁺13]. Recently in [JR15], a single round UC-PAKE protocol (in the standard model and using bilinear pairings) was also proven secure against adaptive corruption using ideas from the dual-system IBE construction of Waters [Wat09]. However, the [JR15] construction did not employ their dual-system simulation-sound QA-NIZK proofs (DSS-QA-NIZK) in a black box manner. Instead, it used ideas from the DSS-QA-NIZK construction and properties as the underlying intuition for the proof.

In this paper, we show that the UC-PAKE of [JR15] can be built in a black-box manner using smooth QA-NIZK arguments. Next, we build on the verifier-based PAKE (VPAKE) construction of [BP13], to construct the first adaptively-secure UC-APAKE protocol, which in addition has a single (asynchronous) round. Since, in the UC framework, the simulator has to detect offline password guesses by an adversary that steals the server password file, for provable security this seems to inevitably require the RO model, and indeed our security proof is in the (limited programmability) RO model.

In our protocol, each party sends an ElGamal style encryption of the (hash of) the password pw to the other party, along with an SPHF of the underlying language and a projection verification hash-key of a smooth QA-NIZK of the underlying language (ElGamal augmented with the SPHF). If such a message is adversarially inserted, the simulator must have the capability to extract password pw' from it, so that it can feed the ideal functionality $\mathcal{F}_{\text{apwKE}}$ to test this guess of the password. Thus, the NIZK proof must have simulation-sound extractability. It was shown in [JR15] that dual-system simulation soundness suffices for this purpose (and that makes the protocol very simple). When using smooth QA-NIZK, this dual-system simulation-soundness can be attained by simply sending an SPHF.

Detailed explanations can be found in Section 5 with proof details in the full version [JR16], where we also explain how the random oracle is used to extract the password efficiently from the exponent. This leads to a security reduction which has an additive computational overhead of $n * m * \text{poly}(q)$, where n is the

⁵ Basically, the output values of the random oracle are all randomly chosen, but different inputs can be assigned dynamically to these outputs [FLR⁺10].

number of random oracle calls, m is the number of online attacks and q is the security parameter.

Recent Related Work. Recently, [JKX18] formulated a stronger UC-APAKE functionality that disallows use of pre-computation to attack a stolen password file. Using an ideal functionality for oblivious pseudo-random functions (OPRF), they give a compiler that converts any standard UC-APAKE realization (such as in this paper) into one that satisfies their stronger definition (in the OPRF-hybrid model). Thus, if the OPRF can be realized with adaptive corruption, then we obtain a adaptive corruption secure (strong) UC-APAKE realization. However, to the best of our knowledge, no adaptive corruption secure (UC-) OPRF realization is known, and hence the problem of realizing adaptive corruption secure strong UC-APAKE remains open.

Organization. The rest of the paper is organized as follows. In Section 2, we introduce the new notion of smooth QA-NIZK proofs. In Section 3, we recall the MDDH assumptions and establish a useful *boosting* theorem relating the assumptions. In Section 4, we give the single group element smooth QA-NIZK construction for linear subspaces. In Section 5, we describe the ideal functionality $\mathcal{F}_{\text{apwKE}}$ for asymmetric password-authenticated key-exchange and construct the new single-round UC-APAKE protocol. Preliminaries and proofs of many of the theorems are relegated to the Appendix.

2 Smooth Quasi-Adaptive NIZK Proofs

We start by reviewing the definition of Quasi-Adaptive computationally-sound NIZK proofs (QA-NIZK) [JR13]. A witness relation is a binary relation on pairs of inputs, the first called a (potential) language member and the second called a witness. Note that each witness relation R defines a corresponding language L which is the set of all x for which there exists a witness w , such that $R(x, w)$ holds.

We will consider QA-NIZK proofs for a probability distribution \mathcal{D} on a collection of (witness-) relations $\mathcal{R} = \{R_\rho\}$ (with corresponding languages L_ρ). Recall that in a QA-NIZK, the CRS can be set after the language parameter has been chosen according to \mathcal{D} . We recall the formal definition of Quasi-Adaptive NIZK below from [JR13].

Definition 1 (QA-NIZK [JR13]). *We call a tuple of efficient algorithms $(\text{pargen}, \text{crsgen}, \text{prover}, \text{ver})$ a quasi-adaptive non-interactive zero-knowledge (QA-NIZK) proof system for witness-relations $\mathcal{R}_\lambda = \{R_\rho\}$ with parameters sampled from a distribution \mathcal{D} over associated parameter language Lpar , if there exist simulators crssim and sim such that for all non-uniform PPT adversaries $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$, we have (in all of the following probabilistic experiments, the experiment starts by setting λ as $\lambda \leftarrow \text{pargen}(1^m)$, and choosing ρ as $\rho \leftarrow \mathcal{D}_\lambda$):*

Quasi-Adaptive Completeness:

$$\Pr \left[\begin{array}{l} \text{CRS} \leftarrow \text{crsgen}(\lambda, \rho) \\ (x, w) \leftarrow \mathcal{A}_1(\text{CRS}, \rho) \\ \pi \leftarrow \text{prover}(\text{CRS}, x, w) \end{array} : \begin{array}{l} \text{ver}(\text{CRS}, x, \pi) = 1 \text{ if} \\ R_\rho(x, w) \end{array} \right] = 1$$

Quasi-Adaptive Soundness:

$$\Pr \left[\begin{array}{l} \text{CRS} \leftarrow \text{crsgen}(\lambda, \rho) \\ (x, \pi) \leftarrow \mathcal{A}_2(\text{CRS}, \rho) \end{array} : \begin{array}{l} x \notin L_\rho \text{ and} \\ \text{ver}(\text{CRS}, x, \pi) = 1 \end{array} \right] \approx 0$$

Quasi-Adaptive Zero-Knowledge:

$$\Pr \left[\text{CRS} \leftarrow \text{crsgen}(\lambda, \rho) : \mathcal{A}_3^{\text{prover}(\text{CRS}, \cdot, \cdot)}(\text{CRS}, \rho) = 1 \right] \\ \approx \\ \Pr \left[(\text{CRS}, \text{trap}) \leftarrow \text{crssim}(\lambda, \rho) : \mathcal{A}_3^{\text{sim}^*(\text{CRS}, \text{trap}, \cdot, \cdot)}(\text{CRS}, \rho) = 1 \right],$$

where $\text{sim}^*(\text{CRS}, \text{trap}, x, w) = \text{sim}(\text{CRS}, \text{trap}, x)$ for $(x, w) \in R_\rho$ and both oracles (i.e. prover and sim^*) output failure if $(x, w) \notin R_\rho$.

We call a QA-NIZK **smooth (-verifier)** if the verifier ver consists of three efficient algorithms $\text{ver} = (\text{hkgen}, \text{pubH}, \text{privH})$, and it satisfies the following modified completeness and soundness conditions. Here, hkgen is a probabilistic algorithm that takes a CRS as input and outputs two keys, hp , a projection hash key, and hk , a private hash key. The algorithm privH takes as input a word (e.g. a potential language member), and a (private hash) key, and outputs a string. Similarly, the algorithm pubH takes as input a word, a proof (for instance generated by prover), and a (projection hash) key hp , and outputs a string.

The **completeness** property is now defined as:

$$\Pr \left[\begin{array}{l} \text{CRS} \leftarrow \text{crsgen}(\lambda, \rho) \\ (x, w) \leftarrow \mathcal{A}_1(\text{CRS}, \rho) \\ \pi \leftarrow \text{prover}(\text{CRS}, x, w) \\ (\text{hp}, \text{hk}) \leftarrow \text{hkgen}(\text{CRS}) \end{array} : \begin{array}{l} \text{privH}(\text{hk}, x) = \text{pubH}(\text{hp}, x, \pi) \\ \text{if } R_\rho(x, w) \end{array} \right] = 1$$

The QA-NIZK is said to satisfy **smooth-soundness** if for all words $x \notin L_\rho$, $\text{privH}(\text{hk}, x)$ is computationally indistinguishable to the Adversary from uniformly random, even when the Adversary is given hp , and even if it produces x after receiving hp .

More precisely, **Quasi-Adaptive Smooth-Soundness** is the following property (let \mathcal{U} be the uniform distribution on the range of privH , which is assumed to be of cardinality exponential in m): for every two-stage efficient oracle adversary \mathcal{A}

$$\Pr \left[\begin{array}{l} \text{CRS} \leftarrow \text{crsgen}(\lambda, \rho), (\text{hp}, \text{hk}) \leftarrow \text{hkgen}(\text{CRS}) \\ (x^*, \sigma) \leftarrow \mathcal{A}^\mathcal{O}(\text{CRS}, \rho, \text{hp}), u \leftarrow \mathcal{U} \end{array} : \mathcal{A}^\mathcal{O}(\text{privH}(\text{hk}, x^*), \sigma) = 1 \mid Q \right] \\ \approx \\ \Pr \left[\begin{array}{l} \text{CRS} \leftarrow \text{crsgen}(\lambda, \rho), (\text{hp}, \text{hk}) \leftarrow \text{hkgen}(\text{CRS}) \\ (x^*, \sigma) \leftarrow \mathcal{A}^\mathcal{O}(\text{CRS}, \rho, \text{hp}), u \leftarrow \mathcal{U} \end{array} : \mathcal{A}^\mathcal{O}(u, \sigma) = 1 \mid Q \right]$$

where the oracle \mathcal{O} is instantiated with $\text{privH}(\text{hk}, \cdot)$, and Q is the condition that x^* is not in the language L_ρ **and** all oracle calls by the adversary in both stages are with L_ρ -language members. Here, σ is a local state of \mathcal{A} .

Note that as opposed to the information-theoretic smoothness property of projective hash functions, one cannot argue here that $\text{privH}(\text{hk}, x)$ for $x \in L_\rho$ can instead just be computed using hp , as that would also require efficiently computing a witness for x . Hence, the need to provide oracle access to $\text{privH}(\text{hk}, \cdot)$ for language members.

Also, note that smooth-soundness implies the earlier definition of soundness [JR13] if verification of (x, π) is defined as $\text{privH}(\text{hk}, x) = \text{pubH}(\text{hp}, x, \pi)$.

To differentiate the functionalities of the verifier of a QA-NIZK from similar functionalities of an SPHF, we will prepend the SPHF functionalities with keyword `spfh` and the QA-NIZK verifier functionalities with the keyword `ver`.

3 Matrix Decisional Assumptions

We will consider bilinear groups that consist of three cyclic groups of prime order q , $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T with an efficient bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$. Group elements \mathbf{g}_1 and \mathbf{g}_2 will typically denote generators of the group \mathbb{G}_1 and \mathbb{G}_2 respectively. Following [EHK⁺13], in this section and the next we will use the notations $[a]_1, [a]_2$ and $[a]_T$ to denote $a\mathbf{g}_1, a\mathbf{g}_2$, and $a \cdot e(\mathbf{g}_1, \mathbf{g}_2)$ respectively and use additive notations for group operations. When talking about a general group \mathbb{G} with generator \mathbf{g} , we will just use the notation $[a]$ to denote $a\mathbf{g}$. However, in the UC-APAKE constructions, we will switch to multiplicative notation for easy readability.

For two vector or matrices A and B , we will denote the product $A^\top B$ as $A \cdot B$. The pairing product $e([A]_1, [B]_2)$ evaluates to the matrix product $[AB]_T$ in the target group with pairing as multiplication and target group operation as addition.

We recall the *Matrix Decisional Diffie Hellman* or MDDH assumptions from [EHK⁺13]. A matrix distribution $\mathcal{D}_{l,k}$, where $l > k$, is defined to be an efficiently samplable distribution on $\mathbb{Z}_q^{l \times k}$ which is full-ranked with overwhelming probability. The $\mathcal{D}_{l,k}$ -MDDH *assumption* in group \mathbb{G} states that with samples $\mathbf{A} \leftarrow \mathcal{D}_{l,k}$ and $(\mathbf{s}, \mathbf{s}') \leftarrow \mathbb{Z}_q^k \times \mathbb{Z}_q^l$, the tuple $([\mathbf{A}], [\mathbf{A}\mathbf{s}])$ is computationally indistinguishable from $([\mathbf{A}], [\mathbf{s}'])$. A matrix distribution $\mathcal{D}_{k+1,k}$ is simply denoted by \mathcal{D}_k .

Intuitively, a $\mathcal{D}_{l,k}$ -MDDH assumption allows us to generate l (computationally) independently random group elements from an initial k independently random exponents. A \mathcal{D}_k -MDDH assumption allows us to generate one extra random group element. In this section, we will establish that, in fact, a \mathcal{D}_k -MDDH assumption can be *boosted* to generate additional (computationally) independently random elements. This will be useful to us in the next section to prove the smoothness property of our construction.

We remark that boosting is different from the random self-reducibility of $\mathcal{D}_{l,k}$ -MDDH assumptions, as described by [EHK⁺13]. While the former aims to generate extra randomness from the same initial sample of vector of random

exponents, the latter talks about results from several independent samples of vector of random exponents. Boosting can be seen as an abstraction of the *switching lemma* of [JR14] and follows the same blueprint for the proof.

For an $l \times k$ matrix \mathbf{A} , we denote $\bar{\mathbf{A}}$ to be the top $k \times k$ square sub-matrix of \mathbf{A} and $\underline{\mathbf{A}}$ to be the bottom $(l - k) \times k$ sub-matrix of \mathbf{A} .

Theorem 1. *Let \mathcal{D}_k be a matrix distribution on $\mathbb{Z}_q^{(k+1) \times k}$. Define another matrix distribution $\mathcal{D}_{l,k}$ on $\mathbb{Z}_q^{l \times k}$ as follows: First sample matrices $\mathbf{A} \leftarrow \mathcal{D}_k$ and $\mathbf{R} \leftarrow \mathbb{Z}_q^{(l-k) \times k}$ and then output $\begin{pmatrix} \mathbf{A} \\ \mathbf{R} \end{pmatrix}$. Then the \mathcal{D}_k -MDDH assumption implies the $\mathcal{D}_{l,k}$ -MDDH assumption.*

We will call *boosting* to be the process of stretching \mathcal{D}_k to $\mathcal{D}_{l,k}$ as above. This theorem is proved as a corollary of an even more general theorem which we will describe after defining the notion of ‘boostable’-ity as follows.

Definition 2. *We say that a matrix distribution \mathcal{D}_k on $\mathbb{Z}_q^{(k+1) \times k}$ is boostable to a matrix distribution $\mathcal{D}_{l,k}$ on $\mathbb{Z}_q^{l \times k}$, where $l > k$, if there are efficiently samplable distributions \mathcal{E} on $\mathbb{Z}_q^{(l-k) \times k}$ and \mathcal{F} on $\mathbb{Z}_q^{(l-k) \times (k+1)}$, such that the following hold:*

– For $\mathbf{A} \leftarrow \mathcal{D}_k, \mathbf{B} \leftarrow \mathcal{D}_{l,k}, \mathbf{E} \leftarrow \mathcal{E}, \mathbf{F} \leftarrow \mathcal{F}$, we have:

$$\bar{\mathbf{B}} \approx \bar{\mathbf{A}}, \quad \underline{\mathbf{B}} \approx \mathbf{E}\bar{\mathbf{A}} \approx \mathbf{F}\mathbf{A}.$$

– For $\mathbf{F} \leftarrow \mathcal{F}$, with overwhelming probability, all entries of the rightmost column \mathbf{F}_r of \mathbf{F} are non-zero.

Theorem 2. *If a matrix distribution \mathcal{D}_k on $\mathbb{Z}_q^{(k+1) \times k}$ is boostable to a matrix distribution $\mathcal{D}_{l,k}$ on $\mathbb{Z}_q^{l \times k}$ then the \mathcal{D}_k -MDDH assumption implies the $\mathcal{D}_{l,k}$ -MDDH assumption.*

Proof. We prove this by a sequence of hybrids, where in the i -th hybrid we transform row $k + i$ from that of $[\mathbf{B}\mathbf{s}]$ to uniformly random. We start off with $i = 0$, where we have the real output $[\mathbf{B}\mathbf{s}]$ and end with $i = l - k$ where we have the fake output which is uniformly random in \mathbb{Z}_q^l .

The i -th hybrid ($[\mathbf{B}], [\mathbf{b}]$) is computed as follows. We sample $[\mathbf{A}]$ from \mathcal{D}_k and \mathbf{s} from \mathbb{Z}_q^k . We set $[\mathbf{B}]$ as $[\bar{\mathbf{A}}]$ and, if $i \neq 0$, the row i of $[\mathbf{B}]$ as the row i of $\mathbf{F}[\mathbf{A}]$. All other rows $j \neq i$ of $[\mathbf{B}]$ are set to the j -th row of $\mathbf{E}[\bar{\mathbf{A}}]$. We set the top k elements of $[\mathbf{b}]$ to be $[\bar{\mathbf{A}}\mathbf{s}]$ and choose all the $(k + j)$ -th elements, where $j < i$, of $[\mathbf{b}]$ uniformly at random from \mathbb{Z}_q . If $i \neq 0$, we set the $(k + i)$ -th element of $[\mathbf{b}]$ to be the i -th element of $\mathbf{F}[\bar{\mathbf{A}}\mathbf{s}]$. For all $j > i$, we set the $(k + j)$ -th element of $[\mathbf{b}]$ to be the j -th element of $\mathbf{E}[\bar{\mathbf{A}}\mathbf{s}]$. To summarize, $[\mathbf{b}]$ is computed as:

$$\begin{bmatrix} [\bar{\mathbf{A}}\mathbf{s}] \\ \$ \\ \vdots \\ \$ \\ (\mathbf{F}[\bar{\mathbf{A}}\mathbf{s}])_i \\ (\mathbf{E}[\bar{\mathbf{A}}\mathbf{s}])_{j=(i+1) \text{ to } (l-k)} \end{bmatrix}$$

We observe that the 0-th hybrid has the distribution of $([\mathbf{B}], [\mathbf{Bs}])$ and the $(l-k)$ -th hybrid has the distribution of $([\mathbf{B}], [s'])$, with s' uniform in \mathbb{Z}_q^l .

Now, $(\mathbf{F}[\mathbf{As}])_i = (\mathbf{F}_l)_i[\bar{\mathbf{A}}s] + (\mathbf{F}_r)_i[\mathbf{As}]$, where \mathbf{F}_l is the first k -column sub-matrix of \mathbf{F} and \mathbf{F}_r is the last column of \mathbf{F} . Suppose we are given a \mathcal{D}_k -MDDH challenge $([\mathbf{A}], \chi = [\mathbf{As}] \text{ or } [s'])$. If $\chi = [\mathbf{As}]$, then $(\mathbf{F}\chi)_i$ is distributed as $(\mathbf{F}[\mathbf{As}])_i$. Else, if $\chi = [s']$, then $(\mathbf{F}\chi)_i$ is distributed uniformly randomly in \mathbb{Z}_q , since $(\mathbf{F}_r)_i$ is overwhelmingly non-zero by design. Next we transition to an intermediate hybrid i' where $[\mathbf{b}]$ is computed as:

$$\begin{bmatrix} [\bar{\mathbf{A}}s] \\ \$ \\ \vdots \\ \$ \\ \$ \\ (\mathbf{E}[\bar{\mathbf{A}}s])_{j=(i+1) \text{ to } (l-k)} \end{bmatrix}$$

As shown above, the hybrid i' is indistinguishable from hybrid i by the \mathcal{D}_k -MDDH assumption. Next we transition to the hybrid $i+1$ where $[\mathbf{b}]$ is computed as:

$$\begin{bmatrix} [\bar{\mathbf{A}}s] \\ \$ \\ \vdots \\ \$ \\ \$ \\ (\mathbf{F}[\mathbf{As}])_{(i+1)} \\ (\mathbf{E}[\bar{\mathbf{A}}s])_{j=(i+2) \text{ to } (l-k)} \end{bmatrix}$$

The hybrid $i+1$ is indistinguishable from hybrid i' , as $\mathbf{E}\bar{\mathbf{A}}$ is identically distributed as $\mathbf{F}\mathbf{A}$. The theorem is thus established by chaining all the hybrids.

Corollary 1. *Any \mathcal{D}_k distribution can be boosted to a $\mathcal{D}_{l,k}$ distribution which inherits the distribution of the top $k \times k$ matrix of the samples.*

This can be seen by setting the top $k \times k$ matrix of a $\mathcal{D}_{l,k}$ sample to be the top $k \times k$ matrix of a \mathcal{D}_k sample and setting the bottom $(l-k) \times k$ sub-matrix of the $\mathcal{D}_{l,k}$ sample to be uniformly random in $\mathbb{Z}_q^{(l-k) \times k}$. The required distributions \mathcal{E} and \mathcal{F} are just the uniform distributions on their respective domains.

This corollary allows us to retain the *representation size* of the top square matrix of a \mathcal{D}_k distribution sample, while boosting it to an assumption required for security proofs. In particular, in applications such as this paper, this can lead to shorter public keys.

Finally, observe that Theorem 2 and the justification of Corollary 1 establishes Theorem 1.

4 Smooth Quasi-Adaptive NIZK Constructions

In this section we show that the single element QA-NIZK [JR14, KW15] for witness-samplable linear subspaces can easily be extended to be smooth QA-NIZK. Particularly, under SXDH, the public hash key hp generated by ver.hkgen consists of a single group element. We follow the construction of Kiltz and Wee [KW15] and prove the result under the more general MDDH assumption in bilinear groups.

We follow additive notation for group operations in this section. In later sections we will use product notation.

Linear Subspace Languages. We first consider languages that are linear subspaces of vectors of \mathbb{G}_1 elements. In other words, the languages we are interested in can be characterized as languages parametrized by $[\mathbf{M}]_1$ as below:

$$L_{[\mathbf{M}]_1} = \{[\mathbf{M}]_1 \mathbf{x} \in \mathbb{G}_1^n \mid \mathbf{x} \in \mathbb{Z}_q^t\}, \text{ where } [\mathbf{M}]_1 \text{ is an } n \times t \text{ matrix of } \mathbb{G}_1 \text{ elements.}$$

Here $[\mathbf{M}]_1$ is an element of the associated *parameter language* Lpar , which is all $n \times t$ matrices of \mathbb{G}_1 elements. The parameter language Lpar also has a corresponding witness relation \mathcal{R}_{par} , where the witness is a matrix of \mathbb{Z}_q elements : $\mathcal{R}_{\text{par}}([\mathbf{M}]_1, \mathbf{M}')$ iff $\mathbf{M} = \mathbf{M}'$.

Robust and Efficiently Witness-Samplable Distributions. Let the $t \times n$ dimensional matrix $[\mathbf{M}]_1$ be chosen according to a distribution \mathcal{D} on Lpar . The distribution \mathcal{D} is called *robust* if with probability close to one the left-most t columns of $[\mathbf{M}]_1$ are full-ranked. A distribution \mathcal{D} on Lpar is called *efficiently witness-samplable* if there is a probabilistic polynomial time algorithm such that it outputs a pair of matrices $([\mathbf{M}]_1, \mathbf{M}')$ that satisfy the relation \mathcal{R}_{par} (i.e., $\mathcal{R}_{\text{par}}([\mathbf{M}]_1, \mathbf{M}')$ holds), and further the resulting distribution of the output $[\mathbf{M}]_1$ is same as \mathcal{D} . For example, the uniform distribution on Lpar is efficiently witness-samplable, by first picking \mathbf{M} at random, and then computing $[\mathbf{M}]_1$.

Smooth QA-NIZK Construction. We now describe a smooth computationally-sound Quasi-Adaptive NIZK (pargen , crsgen , prover , ver) for linear subspace languages $\{L_{[\mathbf{M}]_1}\}$ with parameters sampled from a robust and efficiently witness-samplable distribution \mathcal{D} over the associated parameter language Lpar and given a \mathcal{D}_k -MDDH assumption.

crsgen : The crsgen algorithm generates the CRS as follows. Let $[\mathbf{M}^{n \times t}]_1$ be the parameter supplied to crsgen . It generates an $n \times k$ matrix \mathbf{K} with all elements chosen randomly from \mathbb{Z}_q and a $(k+1) \times k$ matrix \mathbf{A} from the MDDH distribution \mathcal{D}_k . Let $\bar{\mathbf{A}}$ be the top $k \times k$ square matrix of \mathbf{A} .

The common reference string (CRS) has two parts CRS_p and CRS_v which are to be used by the prover and the verifier respectively.

$$\text{CRS}_p^{t \times k} := ([\mathbf{P}]_1 = [\mathbf{M}^\top \mathbf{K}]_1) \quad \text{CRS}_v := ([\mathbf{C}]_2^{n \times k} = [\mathbf{K} \bar{\mathbf{A}}]_2, \quad [\bar{\mathbf{A}}]_2^{k \times k})$$

prover: Given candidate $[\mathbf{y}]_1 = [\mathbf{M}]_1 \mathbf{x}$ with witness vector $\mathbf{x}^{t \times 1}$, the prover generates the following proof consisting of k elements in \mathbb{G}_1 :

$$\pi := \mathbf{x}^\top \text{CRS}_p$$

ver: The algorithm `hkgen` is as follows: Sample $\mathbf{s} \leftarrow \mathbb{Z}_q^k$. Given CRS_v as above, compute `hk` and `hp` as follows:

$$\text{hk} := [\mathbf{C}]_2 \mathbf{s}, \quad \text{hp} := [\bar{\mathbf{A}}]_2 \mathbf{s}$$

The algorithms `pubH` and `privH` are as follows: Given candidate $[\mathbf{y}]_1$, and proof π , compute:

$$\text{privH}(\text{hk}, [\mathbf{y}]_1) := e([\mathbf{y}^\top]_1, \text{hk}) \quad \text{pubH}(\text{hp}, \pi) := e(\pi, \text{hp})$$

Theorem 3. *The above algorithms (`pargen`, `crsgen`, `prover`, `ver`) constitute a smooth computationally -sound Quasi-Adaptive NIZK proof system for linear subspace languages $\{L_{[\mathbf{M}]_1}\}$ with parameters $[\mathbf{M}]_1$ sampled from a robust and efficiently witness-samplable distribution \mathcal{D} over the associated parameter language Lpar , given any group generation algorithm for which the \mathcal{D}_k -MDDH assumption holds for group \mathbb{G}_2 .*

The proofs of completeness, zero knowledge and soundness are same as [KW15]. The proof of smooth soundness follows.

Proof (Smooth Soundness). First, note that the range of `privH` is exponential in the security parameter, for otherwise an adversarial circuit can compute discrete logarithms with non-negligible probability. We prove smoothness by transforming the system over a sequence of games. Game \mathbf{G}_0 just replicates the construction, but samples \mathbf{A} from a distribution $\mathcal{D}_{k+n-t,k}$ obtained by boosting the given distribution \mathcal{D}_k by Theorem 1. The construction only uses the top $k \times k$ sub-matrix $\bar{\mathbf{A}}$ of the sample which is distributed identically for both \mathcal{D}_k and $\mathcal{D}_{k+n-t,k}$. Let $\underline{\mathbf{A}}$ be the bottom $(n-t) \times k$ sub-matrix of \mathbf{A} .

In Game \mathbf{G}_1 , the challenger efficiently samples $[\mathbf{M}]_1$ according to distribution \mathcal{D} , along with witness \mathbf{M} (since \mathcal{D} is an efficiently witness samplable distribution). Since \mathbf{M} is an $n \times t$ dimensional rank t matrix, there is a rank $n-t$ matrix \mathbf{M}^\perp of dimension $n \times (n-t)$ whose columns form a complete basis for the kernel of \mathbf{M}^\top , which means $\mathbf{M}^\top \mathbf{M}^\perp = \mathbf{0}^{t \times (n-t)}$. In this game, the NIZK CRS is computed as follows: Generate matrix $\mathbf{K}'^{n \times k}$ and compute the matrix $\mathbf{T}^{(n-t) \times k}$, such that $\mathbf{T}\bar{\mathbf{A}} = \underline{\mathbf{A}}$. Implicitly set: $\mathbf{K} = \mathbf{K}' + \mathbf{M}^\perp \mathbf{T}$. Therefore we have,

$$\begin{aligned} \text{CRS}_p^{t \times k} &= [\mathbf{M}^\top \mathbf{K}]_1 = [\mathbf{M}^\top (\mathbf{K}' + \mathbf{M}^\perp \mathbf{T})]_1 = [\mathbf{M}^\top \mathbf{K}']_1 \\ [\mathbf{C}]_2^{n \times k} &= [(\mathbf{K}' + \mathbf{M}^\perp \mathbf{T})\bar{\mathbf{A}}]_2 = \mathbf{K}'[\bar{\mathbf{A}}]_2 + \mathbf{M}^\perp[\underline{\mathbf{A}}]_2, \\ \text{hk} &= [\mathbf{C}]_2 \mathbf{s}, \quad \text{hp} = [\bar{\mathbf{A}}]_2 \mathbf{s} \end{aligned}$$

In Game \mathbf{G}_2 , we sample fresh random vectors \mathbf{s}' in \mathbb{Z}_q^k and \mathbf{s}'' in \mathbb{Z}_q^{n-t} and modify the simulated computations as follows:

$$\text{CRS}_p^{t \times k} = [\mathbf{M}^\top \mathbf{K}']_1, \quad [\mathbf{C}]_2^{n \times k} = \mathbf{K}'[\bar{\mathbf{A}}]_2 + \mathbf{M}^\perp[\underline{\mathbf{A}}]_2,$$

$$\text{hk} = \mathbf{K}'[\mathbf{s}']_2 + \mathbf{M}^\perp[\mathbf{s}'']_2, \quad \text{hp} = [\mathbf{s}']_2$$

Given a $\mathcal{D}_{k+n-t,k}$ challenge which is either “real”: $([\mathbf{A}]_2, [\bar{\mathbf{A}}\mathbf{s}]_2, [\mathbf{A}\mathbf{s}]_2)$ or “fake”: $([\mathbf{A}]_2, [\mathbf{s}']_2, [\mathbf{s}'']_2)$, we observe that the real tuple can be used to simulate Game \mathbf{G}_1 , while the fake tuple can be used to simulate Game \mathbf{G}_2 . Thus the games \mathbf{G}_1 and \mathbf{G}_2 are indistinguishable by the $\mathcal{D}_{k+n-t,k}$ -MDDH assumption, which in turn is implied by the \mathcal{D}_k -MDDH assumption by Theorem 2.

Now in Game \mathbf{G}_2 we have,

$$\text{privH}(\text{hk}, [\mathbf{y}^*]_1) = e\left([\mathbf{y}^{*\top}]_1, \mathbf{K}'[\mathbf{s}']_2 + \mathbf{M}^\perp[\mathbf{s}'']_2\right)$$

For the oracle queries where $[\mathbf{y}^*]_1 \in L_{[\mathbf{M}]_1}$, we have $\mathbf{y}^{*\top}\mathbf{M}^\perp = 0^{1 \times (n-t)}$. Hence the simulator responds with $e([\mathbf{y}^{*\top}]_1, \mathbf{K}'[\mathbf{s}']_2)$. Note that \mathbf{s}'' does not appear in this response.

For the adversary supplied $[\mathbf{y}^*]_1 \notin L_{[\mathbf{M}]_1}$, we have $\mathbf{y}^{*\top}\mathbf{M}^\perp \neq 0^{1 \times (n-t)}$. Therefore $\text{privH}(\text{hk}, \mathbf{y}^*)$ is uniformly random, as \mathbf{s}'' is independently random of everything else given to the adversary.

Smooth Split-CRS QA-NIZK for Tagged Affine Languages. QA-NIZKs for linear subspaces were also extended by [JR13] to integer tag-based languages as well as provided split-CRS⁶ instantiation for affine languages. In [JR16], we combine all these extensions and describe a smooth computationally-sound Quasi-Adaptive NIZK (`pargen`, `crsgen`, `prover`, `ver`) for tagged affine linear subspace languages $\{L\}$, parametrized by $([\mathbf{M}_0]_1, [\mathbf{M}_1]_1, [\mathbf{M}_2]_1, [\mathbf{M}_3]_1, [\mathbf{a}]_1)$ and consisting of words of the form:

$$([\mathbf{M}_0\mathbf{x}]_1, [\mathbf{M}_1\mathbf{x} + \mathbf{a}]_1, [(\mathbf{M}_2 + \text{TAG} \cdot \mathbf{M}_3)\mathbf{x}]_1, \text{TAG} \in \mathbb{Z}_q),$$

with parameters sampled from a robust and efficiently witness-samplable distribution \mathcal{D} over $(\mathbf{M}_0, \mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_3, \mathbf{a})$ and given a \mathcal{D}_k -MDDH assumption. We assume that \mathbf{M}_0 is a square matrix and the robustness of \mathcal{D} is defined by \mathbf{M}_0 being non-singular. The smooth QA-NIZK will be split-CRS [JR13], so that CRS_v is independent of the language parameters.

5 Asymmetric UC-PAKE: UC-APAKE

Based on the UC-PAKE functionality of [CHK⁺05], Gentry et al [GMR06] gave another UC functionality for asymmetric PAKE (UC-APAKE). A salient feature of the UC-PAKE functionality [CHK⁺05] is that it models the security requirement that an adversary cannot perform efficient off-line computations on protocol transcripts to verifiably guess the low-entropy password. An adversary can only benefit from the low-entropy of the password by actually conducting an on-line attack (i.e. by impersonating one of the parties with a guessed password).

⁶ A split-CRS QA-NIZK allows the verifier CRS to be generated independent of the language.

This is modeled in the ideal world with a `TestPwd` capability available to the ideal world adversary: if `TestPwd` is called with the correct password, the ideal world adversary is allowed to set the session key. Moreover, in this functionality if any of the parties is corrupted, then the ideal world adversary is given the registered password.

5.1 The UC Ideal Functionality for Asymmetric PAKE

In asymmetric PAKE [GMR06], the ideal functionality also allows an adversary to steal the password file stored at the server (while not necessarily corrupting the server). However, this by itself does not directly provide the actual password to the adversary. However, after this point the adversary is allowed to perform `OfflineTestPwd` tests to mimic a similar capability in the real world (in fact, the ideal world adversary is even allowed to perform `OfflineTestPwd` tests before it steals the password file, but it does not get a confirmation of the guess being correct until after it steals the password file).

Moreover, after the “steal password file” event the adversary is also allowed to impersonate the server to a *correctly guessed* client, even without providing the actual password (as it can clearly do so in the real world). However, compromising impersonation of the client still requires providing a correct password. This differentiation in capabilities also becomes important when characterizing the complexity of a simulator in terms of the real world adversary, as we will see later.

The $\mathcal{F}_{\text{PAKE}}$ functionality for UC-PAKE was a single-session functionality. However, asymmetric PAKE requires that a password file be used across multiple sessions, so the $\mathcal{F}_{\text{apwKE}}$ functionality for UC-APAKE is defined as a multiple-session functionality. Note that this cannot be accomplished simply using composition with joint state [CR03] because the functionality itself requires shared state that needs to be maintained between sessions. The complete UC-APAKE functionality $\mathcal{F}_{\text{apwKE}}$ is described in detail in Fig. 2.

5.2 UC-APAKE based on VPAKE and Smooth-NIZK

We now design an asymmetric UC-PAKE based on Verifier-based PAKE or VPAKE of Benhamouda and Pointcheval [BP13] and the novel Smooth NIZK proofs. The essential idea of [BP13] is that while the Client holds the actual password, the Server does not hold password in the clear. Instead the Server stores a hard to invert function called `PHash` (password hash) evaluated over the password and a random “salt” (`PSalt`) published in the CRS. While executing a session, the client sends encryptions of the password or another function called `PPreHash` (password pre-hash) evaluated on the password. Correspondingly, the server sends encryptions of the stored `PHash`.

Of course, some kind of zero-knowledge proof must accompany these encryptions, and to that end [BP13] can utilize the new smooth projective hash functions (SPHF) for CCA2-encryption [BBC⁺13] such as Cramer-Shoup encryption [CS02]. In each session, both parties generate fresh SPHF private and

Functionality $\mathcal{F}_{\text{apwKE}}$

The functionality $\mathcal{F}_{\text{apwKE}}$ is parameterized by a security parameter k . It interacts with an adversary S and a set of parties via the following queries:

Password Storage and Authentication Sessions

Upon receiving a query (StorePwdFile, sid, P_i, pw) **from party** P_j :

If this is the first StorePwdFile query, store password data record $(\text{file}, P_i, P_j, pw)$ and mark it uncompromised.

Upon receiving a query (CltSession, $sid, ssid, P_i, P_j, pw$) **from party** P_i :

Send $(\text{CltSession}, sid, ssid, P_i, P_j)$ to S . In addition, if this is the first CltSession query for $ssid$, then store session record $(\text{Clt}, ssid, P_i, P_j, pw)$ and mark this record fresh.

Upon receiving a query (SvrSession, $sid, ssid, P_i$) **from party** P_j :

If there is a password data record $(\text{file}, P_i, P_j, pw)$, then send $(\text{SvrSession}, sid, ssid, P_j, P_i)$ to S , and if this is the first SvrSession query for $ssid$, store session record $(\text{Svr}, ssid, P_j, P_i, pw)$, and mark it fresh.

Stealing Password Data

Upon receiving a query (StealPwdFile, sid) **from adversary** S :

If there is no password data record reply to S with 'no password file'. Otherwise, do the following: If the password data record $(\text{file}, P_i, P_j, pw)$ is marked uncompromised, mark it compromised. If there is a tuple $(\text{offline}, pw')$ stored with $pw' = pw$ then send pw to S , otherwise reply to S with 'password file stolen'.

Upon receiving a query (OfflineTestPwd, sid, pw') **from Adversary** S :

If there is no password data record, or if there is a password data record $(\text{file}, P_i, P_j, pw)$ that is marked uncompromised, then store $(\text{offline}, pw')$. Otherwise do: if $pw = pw'$, send pw back to S . If $pw \neq pw'$, reply with 'wrong guess'.

Active Session Attacks

Upon receiving a query (TestPwd, $sid, ssid, P_i, pw'$) **from the adversary** S :

If there is a session record of the form $(\text{role}, ssid, P_i, P_j, pw)$ which is fresh, then do: If $pw = pw'$, mark the record compromised and reply to S with "correct guess". If $pw \neq pw'$, mark the record interrupted and reply with "wrong guess".

Upon receiving a query (Impersonate, $sid, ssid$)

If there is a session record of the form $(\text{Clt}, ssid, P_i, P_j, pw)$ which is fresh, then do: then if there is a password data record file $(\text{file}, P_i, P_j, pw)$ that is marked compromised, mark the session record compromised and reply to S with 'correct guess', else mark the session record interrupted and reply with wrong guess.

Key Generation and Authentication

Upon receiving a query (NewKey, $sid, ssid, P_i, sk$) **from** S , **where** $|sk| = k$:

If there is a session record of the form $(\text{role}, ssid, P_i, P_j, pw)$ that is not marked completed,

- If this record is compromised, or either P_i or P_j is corrupted, then output $(sid, ssid, sk)$ to player P_i .
- If this record is fresh, and there is a session record $(\text{role}, ssid, P_j, P_i, pw')$ with $pw' = pw$, and a key sk' was sent to P_j , and $(\text{role}, ssid, P_j, P_i, pw)$ was fresh at the time, then output $(sid, ssid, sk')$ to P_i .
- In any other case, pick a new random key sk' of length k and send $(sid, ssid, sk')$ to P_i .

Either way, mark the record (P_i, P_j, pw) as completed.

Upon receiving (Corrupt, sid, P) **from** S : if there is a $(\text{Clt}, sid, P, P', pw)$ recorded, return pw to S , and mark P_i corrupted. If there is a $(\text{Svr}, sid, P, P', pw)$ recorded, then mark P corrupted and (internally) call (StealPwdFile, sid).

Fig. 2. The password-based key-exchange functionality $\mathcal{F}_{\text{apwKE}}$

projection keys (to be employed on incoming messages). The projection key is sent (piggy-backed) along with the encrypted message. If the encrypted messages use the correct password (meaning both parties have the same password

<p>Generate $\mathbf{g} \leftarrow \mathbb{G}_1$; $a_1, a_2, bc, bs \leftarrow \mathbb{Z}_q$ and let $\rho = \{\mathbf{a}_1 = \mathbf{g}^{a_1}, \mathbf{a}_2 = \mathbf{g}^{a_2}, \mathbf{b}_c = \mathbf{g}^{bc}, \mathbf{b}_s = \mathbf{g}^{bs}\}$.</p> <p>Define languages $\left[\begin{array}{l} L_C = \{(R, S, H) \mid \exists r, p : R = \mathbf{g}^r, S = \mathbf{a}_1^r \mathbf{b}_c^p, H = \mathbf{b}_s^p\} \\ L_S = \{(R, S) \mid \exists r : R = \mathbf{g}^r, S = \mathbf{a}_2^r\} \end{array} \right]$</p> <p>Let $(\text{hk}_c, \text{hp}_c) \leftarrow \text{sphf}(L_C).\text{hkgen}$ and $(\text{hp}_s, \text{hks}) \leftarrow \text{sphf}(L_S).\text{hkgen}$.</p> <p>Define languages:</p> $\left[\begin{array}{l} L_C^+ = \{(R, S, H, T, l) \mid \exists r, p : R = \mathbf{g}^r, S = \mathbf{a}_1^r \mathbf{b}_c^p, H = \mathbf{b}_s^p, T = \text{sphf}.\text{pubH}(\text{hp}_c, \langle R, S, H \rangle, l; r, p)\} \\ L_S^+ = \{(R, S, T, l) \mid \exists r : R = \mathbf{g}^r, S = \mathbf{a}_2^r, T = \text{sphf}.\text{pubH}(\text{hp}_s, \langle R, S \rangle, l; r)\} \end{array} \right]$ <p>Let $(\text{pargen}_P, \text{crsgen}_P, \text{prover}_P, \text{ver}_P)$ be Smooth QA-NIZKs for languages L_P^+, with $P \in \{C, S\}$.</p> <p>Let $\text{CRS}_P \leftarrow \text{crsgen}_P(\rho)$ and \mathcal{H} be a collision resistant hash function.</p> <p>Let \mathcal{RO} be a random oracle and let $\text{phash} = \mathcal{RO}(\text{sid}, P_i, P_j, \text{pwd})$.</p> <p>Note that there are several sessions, designated by unique ssid-s, within the scope of a single sid. Thus, phash is the same across all these sessions.</p> <p style="text-align: center;">$\text{CRS} := (\rho, \text{hp}_c, \text{hp}_s, \text{CRS}_C, \text{CRS}_S, \mathcal{H})$.</p> <p style="text-align: center;">Server Persistent State := $\mathbf{b}_s^{\text{phash}}$.</p>	
Client P_i	Network
<p>Input $(\text{CltSession}, \text{sid}, \text{ssid}, P_i, P_j, \text{pwd})$.</p> <p>Choose $r_1 \leftarrow \mathbb{Z}_q$ and $(\text{HK}_1, \text{HP}_1) \leftarrow \text{vers}.\text{hkgen}(\text{CRS}_S)$.</p> <p>Set $R_1 = \mathbf{g}^{r_1}$, $S_1 = \mathbf{a}_1^{r_1} \mathbf{b}_c^{\text{phash}}$,</p> <p style="padding-left: 2em;">$T_1 = \text{sphf}_c.\text{pubH}(\text{hp}_c, \langle R_1, S_1, \mathbf{b}_s^{\text{phash}} \rangle, i_1; r_1, \text{phash})$,</p> <p style="padding-left: 2em;">$W_1 = \text{prover}_C(\text{CRS}_C, \langle R_1, S_1, \mathbf{b}_s^{\text{phash}} \rangle, T_1, i_1; r_1, \text{phash})$,</p> <p style="padding-left: 2em;">where $i_1 = \mathcal{H}(\text{sid}, \text{ssid}, P_i, P_j, R_1, S_1, \text{HP}_1)$.</p> <p>Erase r_1, send $(R_1, S_1, T_1, \text{HP}_1)$ and retain (W_1, HK_1).</p>	$\xrightarrow{R_1, S_1, T_1, \text{HP}_1} P_j$
<p>Receive $(R'_2, S'_2, T'_2, \text{HP}'_2)$.</p> <p>If any of $R'_2, S'_2, T'_2, \text{HP}'_2$ is not in their respective group or is 1, set $\text{sk}_1 \xleftarrow{\\$} \mathbb{G}_T$,</p> <p style="padding-left: 2em;">else compute $i'_2 = \mathcal{H}(\text{sid}, \text{ssid}, P_j, P_i, R'_2, S'_2, \text{HP}'_2)$,</p> <p style="padding-left: 2em;">and $\text{sk}_1 = \text{vers}.\text{privH}(\text{HK}_1, \langle R'_2, S'_2 / \mathbf{b}_s^{\text{phash}}, T'_2, i'_2 \rangle) \cdot \text{verc}.\text{pubH}(\text{HP}'_2, W_1)$.</p> <p>Output $(\text{sid}, \text{ssid}, \text{sk}_1)$.</p>	$\xleftarrow{R'_2, S'_2, T'_2, \text{HP}'_2} P_j$
Server P_j	Network
<p>Input $(\text{SvrSession}, \text{sid}, \text{ssid}, P_j, P_i, \text{Server Persistent State})$.</p> <p>Choose $r_2 \leftarrow \mathbb{Z}_q$ and $(\text{HK}_2, \text{HP}_2) \leftarrow \text{verc}.\text{hkgen}(\text{CRS}_C)$.</p> <p>Set $R_2 = \mathbf{g}^{r_2}$, $S_2 = \mathbf{a}_2^{r_2} \mathbf{b}_s^{\text{phash}}$,</p> <p style="padding-left: 2em;">$T_2 = \text{sphf}_s.\text{pubH}(\text{hp}_s, \langle R_2, S_2 / \mathbf{b}_s^{\text{phash}} \rangle, i_2; r_2)$,</p> <p style="padding-left: 2em;">$W_2 = \text{prover}_S(\text{CRS}_S, \langle R_2, S_2 / \mathbf{b}_s^{\text{phash}} \rangle, T_2, i_2; r_2)$,</p> <p style="padding-left: 2em;">where $i_2 = \mathcal{H}(\text{sid}, \text{ssid}, P_j, P_i, R_2, S_2, \text{HP}_2)$.</p> <p>Erase r_2, send $(R_2, S_2, T_2, \text{HP}_2)$ and retain (W_2, HK_2).</p>	$\xrightarrow{R_2, S_2, T_2, \text{HP}_2} P_i$
<p>Receive $(R'_1, S'_1, T'_1, \text{HP}'_1)$.</p> <p>If any of $R'_1, S'_1, T'_1, \text{HP}'_1$ is not in their respective group or is 1, set $\text{sk}_2 \xleftarrow{\\$} \mathbb{G}_T$,</p> <p style="padding-left: 2em;">else compute $i'_1 = \mathcal{H}(\text{sid}, \text{ssid}, P_i, P_j, R'_1, S'_1, \text{HP}'_1)$,</p> <p style="padding-left: 2em;">and $\text{sk}_2 = \text{verc}.\text{privH}(\text{HK}_2, \langle R'_1, S'_1, \mathbf{b}_s^{\text{phash}} \rangle, T'_1, i'_1) \cdot \text{vers}.\text{pubH}(\text{HP}'_1, W_2)$.</p> <p>Output $(\text{sid}, \text{ssid}, \text{sk}_2)$.</p>	$\xleftarrow{R'_1, S'_1, T'_1, \text{HP}'_1} P_i$

Fig. 3. Single round RO-hybrid UC-APAKE protocol under SXDH assumption.

or its PHash), then SPHF computed on the message by the receiving party using the SPHF hash key it generated equals the SPHF computed on the message by the sending party using the SPHF projection key it received. Thus, these SPHF hashes can be used to compute the session key. Smoothness property of the SPHF guarantees security of the VPAKE scheme.

Unfortunately, each party must retain the witness used in the CCA2 encryption, as computing the SPHF projection-hash of its outgoing encrypted message using the received projection key requires this witness. In the strong simulation paradigm of universally composable security, this leads to a problem if an Adversary can corrupt a session dynamically after the outgoing message has been sent and the incoming message has not yet been received. Thus, this SPHF methodology can only handle static corruption. While Jutla and Roy [JR15] have recently given an efficient UC-PAKE protocol which can handle dynamic corruption, the construction uses ideas from dual-system simulation-sound QA-NIZK that they introduce there. These ideas are rather intricate and do not seem to allow a modular or generic design of such UC password-authenticated protocols.

In this paper, we show that the new notion of Smooth QA-NIZK allows easy to understand (and equally efficient) modular or generic design. Just as QA-NIZK proofs can be seen as generalization of projective hash proof systems to public verifiability (and also assuring zero-knowledge), the novel notion of Smooth QA-NIZK naturally generalizes the notion of smooth projective hash functions where instead of the witness, the publicly verifiable proof can be used to evaluate the projection-hash. The zero-knowledge property of this publicly verifiable proof assures that this proof and hence the projection-hash can be generated by a simulator with no access to the witness. In particular, each party in the UC-PAKE protocol can generate an encryption of the password and generate this publicly verifiable QA-NIZK proof, send the encryption to the other party, erase the witness and retain just the proof for later generation of session key.

The natural question that arises is whether one needs a notion of smooth-soundness under simulation. Indeed, one does need some form of unbounded simulation-soundness as the UC simulator generates QA-NIZK proofs on non-language members without access to the password. Unfortunately, the recent efficient unbounded simulation sound QA-NIZK construction of [KW15] does not extend to be smooth under unbounded simulation (or at least current techniques do not seem to allow one to prove so). The dual-system simulation sound QA-NIZK [JR15] does satisfy smoothness property, but it would need introduction of various new intricate definitions and complicated proofs. One may also ask whether CCA2 encryption by itself provides the required simulation soundness, but that is also not the case, as CCA2 encryption by itself does not give a privately-verifiable (say, via its underlying SPHF as in Cramer-Shoup encryption) proof that it is the password that is being encrypted.

In light of this, it turns out that the simplest way to design the UC-APAKE (or UC-PAKE) protocol is to use an ElGamal encryption of the password (or its PPreHash or PHash) and augment it with an SPHF proof of its consistency,

and finally a Smooth QA-NIZK on this augmented ElGamal encryption. (If the reader is interested in the simpler UC-PAKE protocol secure under dynamic corruption in the new Smooth QA-NIZK framework, the UC-PAKE definition and protocol are provided in [JR16]).

We will also need the random oracle hybrid model to achieve the goal of a UC-APAKE protocol, as explained next. The focus of [BP13] was to design protocols which can be proven secure in the standard model. They formalized a security notion for APAKEs modifying the game-based BPR model [BPR00]. However, our focus is to construct an APAKE protocol in the UC model. In the UC model of [GMR06], the UC simulator must be able to detect offline password guess attempts of the adversary. This is not possible in the standard model as offline tests can be internally performed by the adversary. In order to intercept offline tests by the adversary, it thus becomes inevitable to use an idealized model, such as the random oracle model.

So in particular, we adapt the random oracle-based password hashing scheme of [BP13]. In the scheme, the public parameters are $param = \mathbf{b}_C, \mathbf{b}_S$ randomly sampled from \mathbb{G}_1 and a random oracle \mathcal{RO} . Define $\mathbf{phash} = \mathcal{RO}(sid, \text{Client-id}, \text{Server-id}, \text{pwd})$, where Client-id, Server-id are the ids of the participating parties, sid is the common session-id for all sessions between these parties and pwd is the password of the client. Note that there are several sessions designated by unique $ssid$ -s within the scope of a single sid . Thus \mathbf{phash} is the same across all these sessions. We set:

$$\begin{aligned} \text{PPreHash}(param, \text{pwd}) &= \mathbf{b}_C^{\mathbf{phash}} \\ \text{PSalt}(param) &= \mathbf{b}_S \\ \text{PHash}(param, \text{pwd}) &= \mathbf{b}_S^{\mathbf{phash}} \end{aligned}$$

Corresponding to the asymmetric storages of the client and the server, we define the following languages, one for each party, which implicitly check the consistency of correct elements being used (\mathbf{a}_1 and \mathbf{a}_2 are essentially public keys for ElGamal encryption):

$$\begin{aligned} L_C &= \{(R, S, H) \mid \exists r, p : R = \mathbf{g}^r, S = \mathbf{a}_1^r \mathbf{b}_C^p, H = \mathbf{b}_S^p\} \\ L_S &= \{(R, S) \mid \exists r : R = \mathbf{g}^r, S = \mathbf{a}_2^r\} \end{aligned}$$

We now plug these languages into UC-PAKE methodology described above. The client sends ElGamal encryption of \mathbf{b}_C^p , as in (R, S) of L_C , while the server supplies the last element H for forming a word of L_C . The server sends ElGamal encryption of \mathbf{b}_S^p , while the client divides out \mathbf{b}_S^p from the second component to form a word of L_S .

The CRS provides public smooth₂ SPHF keys for the languages L_C and L_S , which are used by the client and server respectively to compute T_1 and T_2 for their flows.

Lastly, we use Smooth QA-NIZK proofs for generating a public hash key and a private hash key over the above languages augmented with the SPHF's as

below:

$$L_C^+ = \left\{ (R, S, H, T, l) \mid \exists r, p : \begin{array}{l} R = \mathbf{g}^r, S = \mathbf{a}_1^r \mathbf{b}_C^p, H = \mathbf{b}_S^p, \\ T = \text{sphf.pubH}(\text{hp}_C, \langle R, S, H \rangle, l; r, p) \end{array} \right\}$$

$$L_S^+ = \{(R, S, T, l) \mid \exists r : R = \mathbf{g}^r, S = \mathbf{a}_2^r, T = \text{sphf.pubH}(\text{hp}_S, \langle R, S \rangle, l; r)\}$$

The client generates a Smooth QA-NIZK verification key pair for the server language L_S^+ , retains the private key HK_1 and sends the public key HP_1 along with the ElGamal encryption and the SPHF. The client computes a QA-NIZK proof W_1 of $(R_1, S_1, \mathbf{b}_S^{\text{hash}}, T_1) \in L_C^+$ with label $i_1 = \mathcal{H}(\text{sid}, \text{ssid}, P_i, P_j, R_1, S_1, T_1, \text{HP}_1)$ and retains that for later key computation.

Similarly, the server generates a Smooth QA-NIZK verification key pair for the client language L_C^+ , retains the private key HK_2 and sends the public key HP_2 along with the ElGamal encryption and the SPHF. The server computes a QA-NIZK proof W_2 of $(R_2, S_2/\mathbf{b}_S^{\text{hash}}, T_2) \in L_S^+$ with label $i_2 = \mathcal{H}(\text{sid}, \text{ssid}, P_j, P_i, R_2, S_2, T_2, \text{HP}_2)$ and retains that for later key computation.

In the second part of the protocol, after receiving the peer flow, each party computes the final secret key as the product of the private Smooth QA-NIZK hash of the peer flow with own private Smooth QA-NIZK key and the public Smooth QA-NIZK hash of the (retained) QA-NIZK proof of own flow with the peer public Smooth QA-NIZK hash key. Formally the client computes:

$$\text{ver}_S.\text{privH}(\text{HK}_1, \langle R'_2, S'_2/\mathbf{b}_S^{\text{hash}}, T'_2, i'_2 \rangle) \cdot \text{ver}_C.\text{pubH}(\text{HP}'_2, W_1).$$

Similarly, the server computes:

$$\text{ver}_C.\text{privH}(\text{HK}_2, \langle R'_1, S'_1, \mathbf{b}_S^{\text{hash}}, T'_1, i'_1 \rangle) \cdot \text{ver}_S.\text{pubH}(\text{HP}'_1, W_2).$$

Given the completeness property of the Smooth QA-NIZK, it is not difficult to see that legitimately completed peer sessions end up with equal keys. In the next section, we prove that this protocol securely realizes $\mathcal{F}_{\text{apwKE}}$, as stated in the theorem below.

The complete protocol is described in detail in Figure 3. The SPHF sphf is required to be a smooth_2 projective hash function (see [JR16] for definitions). For simplicity, in this paper we focus on constructions based on \mathcal{D}_1 -MDDH assumptions, and in particular the SXDH assumption.

Theorem 4. *Under the \mathcal{D}_1 -MDDH assumption SXDH, the protocol in Fig. 3 securely realizes the $\mathcal{F}_{\text{apwKE}}$ functionality in the $(\mathcal{F}_{\text{CRS}}, \mathcal{F}_{\text{RO}})$ -hybrid model, in the presence of adaptive corruption adversaries. The number of unique password arguments passed to TestPw and OfflineTestPw of $\mathcal{F}_{\text{apwKE}}$ combined in the ideal world is at most the number of random oracle calls in the $(\mathcal{F}_{\text{CRS}}, \mathcal{F}_{\text{RO}})$ -hybrid world.*

We describe the intuition of the proof below and describe the UC simulator, while detailed formal steps proving indistinguishability of the real and the ideal world are relegated to the full version [JR16].

5.3 Main Idea of the UC Simulator

The UC simulator \mathcal{S} works as follows: It simulates the random oracle calls and records all the query response pairs. It will generate the CRS for $\widehat{\mathcal{F}}_{\text{PAKE}}$ using the real world algorithms, except for the Smooth QA-NIZK, for which it uses the simulated CRS generator. It also retains the private hash keys of the SPHF's. The next main difference is in the simulation of the outgoing message of the real world parties: \mathcal{S} uses a dummy message μ instead of the real password which it does not have access to. Further, it postpones computation of W till the session-key generation time. Finally, another difference is in the processing of the incoming message, where \mathcal{S} decrypts the incoming message R'_2, S'_2 and runs through the list of random oracle queries to search for a pwd' , such that the decryption is $\mathbf{b}_S^{\mathcal{R}\mathcal{O}(sid, P_i, P_j, \text{pwd}')}$, which it uses to call the ideal functionality's test function. It next generates an sk similar to how it is generated in the real-world. It sends sk to the ideal functionality to be output to the party concerned.

Since the (R_1, S_1) that it sends out is no longer such that $(R_1, S_1, \mathbf{b}_S^{\text{phash}})$ in the language L_C , it has to use the private key of the SPHF in order to compute T_1 on $(R_1, S_1, \mathbf{b}_S^{\text{phash}})$ and the QA-NIZK proof simulator to compute W_1 .

There are other special steps designed to simulate stealing the password file and then impersonating the server to the client. Specifically, when the password file is stolen, the simulator still may not know pwd . It then preemptively sets phash to a random value and pretends that this is the random oracle response with the correct pwd query. Later on when there is a successful pwd query, which the simulator can find out by the online or offline testpwd ideal functionality calls, it sets the record accordingly.

In case of a stolen password file, the simulator includes a ‘‘Client Only Step’’ which lets it test (modified) server flows for consistency and call the Impersonate functionality if consistency checks out. The server simulation steps do not include such a step to model the security notion that even if the password file is stolen, the adversary should still not be able to impersonate the client.

5.4 Main Idea of the Proof of UC Realization

The proof that the simulator \mathcal{S} described above simulates the Adversary in the real-world protocol, follows essentially from the properties of the Smooth QA-NIZK and smooth_2 SPHF, and we give a broad outline here. The proof will describe various experiments between a challenger \mathcal{C} and the adversary, which we will just assume to be the environment \mathcal{Z} (as the adversary \mathcal{A} can be assumed to be just dummy and following \mathcal{Z} 's commands). In the first experiment the challenger \mathcal{C} will just be the combination of the code of the simulator \mathcal{S} above and $\widehat{\mathcal{F}}_{\text{PAKE}}$. In particular, after the environment issues a CltSession request with a password pwd , the challenger gets that password. So, while in the first experiment, the challenger (copying \mathcal{S}) does not use pwd directly, from the next experiment onwards, it can use pwd . Thus, the main goal of the ensuing experiments is to modify the fake tuples $\mathbf{g}^{r_1}, \mathbf{g}^{r'}$ by real tuples (as in real-world)

$\mathbf{g}^{r_1}, \mathbf{a}_1^{r_1} \mathbf{b}_C^{\text{phash}}$, since the challenger has access to `pwd`, and hence `phash`. This is accomplished by a hybrid argument, modifying one instance at a time using DDH assumption in group \mathbb{G}_1 .

The guarantee that the client cannot be impersonated by the adversary, even when the password file is stolen is established by noting that $\mathbf{b}_C^{\text{phash}}$, which is what the client encrypts in its flows, is hard to compute given the server persistent state $\mathbf{b}_S^{\text{phash}}$. This is formally captured in the proof by using a DDH transition from $(\mathbf{b}_S, \mathbf{b}_C, \mathbf{b}_S^{\text{phash}}, \mathbf{b}_C^{\text{phash}})$ to $(\mathbf{b}_S, \mathbf{b}_C, \mathbf{b}_S^{\text{phash}}, \mathbf{b}_C^z)$, where z is independently random from `phash`.

Once all the instances are corrected, i.e. R_1, S_1 are generated as $\mathbf{g}^{r_1}, \mathbf{a}_1^{r_1} \mathbf{b}_C^{\text{phash}}$, the challenger can switch to the real-world because the tuples $R_1, S_1, \mathbf{b}_S^{\text{phash}}$ are now in the language L_C . This implies that the session keys are generated exactly as in the real-world.

5.5 Adaptive Corruption

The UC protocol described above is also UC-secure against adaptive corruption of parties by the Adversary in the erasure model. In the real-world when the adversary corrupts a client (with a `Corrupt` command), it gets the internal state of the client. Clearly, if the party has already been invoked with a `CltSession` command then the password `pwd` is leaked at the minimum, and hence the ideal functionality $\mathcal{F}_{\text{PAKE}}$ leaks the password to the Adversary in the ideal world. In the protocol described above, the Adversary also gets W_1 and HK_1 , as this is the only state maintained by each client between sending $R_1, S_1, T_1, \text{HP}_1$, and the final issuance of session-key. Simulation of HK_1 is easy for the simulator \mathcal{S} since \mathcal{S} generates HK_1 exactly as in the real world. For generating W_1 , which \mathcal{S} had postponed to computing till it received an incoming message from the adversary, it can now use the `pwd` which it gets from $\widehat{\mathcal{F}}_{\text{PAKE}}$ by issuing a `Corrupt` call to $\widehat{\mathcal{F}}_{\text{PAKE}}$. More precisely, it issues the `Corrupt` call, and gets `pwd`, and then calls the QA-NIZK simulator with the tuple $(R_1, S_1, \mathbf{b}_S^{\text{phash}}, T_1, i_1)$ to get W_1 . Note that this computation of W_1 is identical to the postponed computation of W_1 in the computation of client factor of sk_1 (which is really used in the output to the environment when `pwd' = pwd`).

In case of server corruption, the simulator does not get `pwd`, but is able to set `phash` which also enables it to compute W_2 using the QA-NIZK simulator on $(R_2, S_2, \mathbf{b}_S^{\text{phash}}, T_2, i_2)$.

We first define a simulator which interfaces with the ideal functionality and the adversary and then through a series of experiments convert it to just the real world protocol interacting with the same adversary.

5.6 Simulator for the Protocol

We will assume that the adversary \mathcal{A} in the UC protocol is dummy, and essentially passes back and forth commands and messages from the environment \mathcal{Z} .

Thus, from now on we will use environment \mathcal{Z} as the real adversary, which outputs a single bit. The simulator \mathcal{S} will be the ideal world adversary for $\mathcal{F}_{\text{apwKE}}$. It is a universal simulator that uses \mathcal{A} as a black box. For each instance (session and a party), we will use a prime, to refer to variables received in the message from \mathcal{Z} (i.e. \mathcal{A}). We will call a message *legitimate* if it was not altered by \mathcal{Z} , and delivered in the correct session and to the correct party.

Responding to random oracle queries. Let the input be m . If there is a record of the form (m, r) , that is, m was queried before and was responded with r , then just return r .

Otherwise, if m is of the form (sid, P_i, P_j, x) , for some x and the password file has been stolen then call `OfflineTestPwd` with x . If the test succeeds then return `phash`, which must already have been set (see Stealing Password File below), and record (m, phash) .

In all other cases, generate $r \leftarrow \mathbb{Z}_q$, record (m, r) and return r .

Setting the CRS. The simulator \mathcal{S} picks the CRS just as in the real world, except the QA-NIZK CRS-es are generated using the crs-simulators, which also generate simulator trapdoors $\text{trap}_{\mathcal{C}}, \text{trap}_{\mathcal{S}}$. It retains $a_1, a_2, \text{trap}_{\mathcal{C}}, \text{trap}_{\mathcal{S}}, \text{hk}_{\mathcal{C}}, \text{hk}_{\mathcal{S}}$ as trapdoors.

New Client Session: Sending a message to \mathcal{Z} . On message $(\text{CltSession}, sid, \text{ssid}, P_i, P_j)$ from $\mathcal{F}_{\text{apwKE}}$, \mathcal{S} starts simulating a new instance of the protocol for client P_i , server P_j , session identifier `ssid`, and CRS set as above. We will denote this instance by (P_i, ssid) and call it a *client instance*.

To simulate this instance, \mathcal{S} chooses r_1, r'_1, r''_1 at random, and sets $R_1 = \mathbf{g}^{r_1}$, $S_1 = \mathbf{g}^{r'_1}$ and $T_1 = \mathbf{g}^{r''_1}$. Next, \mathcal{S} generates $(\text{HK}_1, \text{HP}_1) \leftarrow \text{ver}_{\mathcal{S}}.\text{hkgen}(\text{CRS}_{\mathcal{S}})$ and sets $i_1 = \mathcal{H}(sid, \text{ssid}, P_i, P_j, R_1, S_1, \text{HP}_1)$. It retains (i_1, HK_1) . It then hands $(R_1, S_1, T_1, \text{HP}_1)$ to \mathcal{Z} on behalf of this instance.

New Server Session: Sending a message to \mathcal{Z} . On message $(\text{SvrSession}, sid, \text{ssid}, P_j, P_i)$ from $\mathcal{F}_{\text{apwKE}}$, \mathcal{S} starts simulating a new instance of the protocol for client P_i , server P_j , session identifier `ssid`, and CRS set as above. We will denote this instance by (P_j, ssid) and call it a *server instance*.

To simulate this instance, \mathcal{S} chooses r_2, r'_2, r''_2 at random, and sets $R_2 = \mathbf{g}^{r_2}$, $S_2 = \mathbf{g}^{r'_2}$ and $T_2 = \mathbf{g}^{r''_2}$. Next, \mathcal{S} generates $(\text{HK}_2, \text{HP}_2) \leftarrow \text{ver}_{\mathcal{C}}.\text{hkgen}(\text{CRS}_{\mathcal{C}})$ and sets $i_2 = \mathcal{H}(sid, \text{ssid}, P_j, P_i, R_2, S_2, \text{HP}_2)$. It retains (i_2, HK_2) . It then hands $(R_2, S_2, T_2, \text{HP}_2)$ to \mathcal{Z} on behalf of this instance.

On Receiving a Message from \mathcal{Z} . On receiving a message $R'_2, S'_2, T'_2, \text{HP}'_2$ from \mathcal{Z} intended for a **client instance** (P, ssid) , the simulator \mathcal{S} does the following:

1. If any of the the real world protocol checks, namely group membership and non-triviality fail it goes to the step “Other Cases” below.
2. If the message received from \mathcal{Z} is same as message sent by \mathcal{S} on behalf of peer P' in session `ssid`, then \mathcal{S} just issues a `NewKey` call for P .

3. (“Client Only Step”): If StealPwdfFile has already taken place then do the following: If $S'_2 = R_2'^{a_2} \mathbf{b}_S^{\text{phash}}$, then \mathcal{S} calls $\mathcal{F}_{\text{apwKE}}$ with (Impersonate, P , $ssid$, $ssid$) and skips to the “Key Setting” step below, and otherwise go to the step “Other Cases”.
4. It searches its random oracle query response pairs $\{(m_k, h_k)\}_k$ and checks whether for some $k = x$ we have $S'_2 = R_2'^{a_2} \mathbf{b}_S^{h_x}$ and m_x is of the form $(ssid, P_i, P_j, \text{pwd}')$. If so, then \mathcal{S} calls $\mathcal{F}_{\text{apwKE}}$ with (TestPwdf, $ssid$, P , pwd') else it goes to the step “Other Cases” below. If the test passes, it sets $\text{phash} = h_x$ and goes to the “Key Setting” step below, else it goes to the step “Other Cases” below.
5. (“Key Setting Step”): Compute $i'_2 = \mathcal{H}(ssid, ssid, P_j, P_i, R_2', S'_2, \text{HP}'_2)$. If $T'_2 \neq \text{sphf}_S.\text{privH}(h_{\text{ks}}, \langle R_2', S'_2 / \mathbf{b}_S^{\text{phash}} \rangle, i'_2)$ then goto the step “Other Cases”. Else, compute $W_1 = \text{sim}(\text{CRS}_C, \text{trap}_C, \langle R_1, S_1, \mathbf{b}_S^{\text{phash}}, T_1, i_1 \rangle)$. Issue a NewKey call to $\widehat{\mathcal{F}}_{\text{PAKE}}$ with key

$$\text{ver}_S.\text{privH}(\text{HK}_1, \langle R_2', S'_2 / \mathbf{b}_S^{\text{phash}}, T'_2, i'_2 \rangle) \cdot \text{ver}_C.\text{pubH}(\text{HP}'_2, W_1)$$

6. (“Other Cases”): \mathcal{S} issues a TestPwdf call to $\widehat{\mathcal{F}}_{\text{PAKE}}$ with the dummy password μ , followed by a NewKey call with a random session key, which leads to the functionality issuing a random and independent session key to the party P .

On receiving a message $R'_1, S'_1, T'_1, \text{HP}'_1$ from \mathcal{Z} intended for a **server instance** $(P, ssid)$, the response of the simulator \mathcal{S} is symmetric to the response described above for client instances, except the above step “Client Only Step” is skipped.

Stealing Password File. If there was a successful online TestPwdf call by the simulator, before this StealPwdfFile call, the corresponding random oracle response h_k was already assigned to the variable phash . Otherwise, the simulator runs through the set of random oracle query response set of the adversary $\{(m_k, h_k)\}_k$, which were not used for an online TestPwdf call. For all the m_k 's of the form $(ssid, P_i, P_j, \text{pwd}')$, it calls (OfflineTestPwdf, $ssid, \text{pwd}'$). Next, \mathcal{S} calls StealPwdfFile. If StealPwdfFile returns pwd then it must equal pwd' in some m_k . Assign to the variable phash the value h_k from the earlier recorded random oracle response to m_k . Otherwise, phash is assigned a fresh random value. The Server Persistent State $\mathbf{b}_S^{\text{phash}}$ is computed accordingly and given to the adversary.

Client Corruption. On receiving a Corrupt call from \mathcal{Z} for client instance P_i in session $ssid$, the simulator \mathcal{S} calls the Corrupt routine of $\mathcal{F}_{\text{apwKE}}$ to obtain pwd . If \mathcal{S} had already output a message to \mathcal{Z} , and not output sk_1 it computes

$$W_1 = \text{sim}_C(\text{CRS}_C, \text{trap}_C, \langle R_1, S_1, \mathbf{b}_S^{\text{phash}}, T_1, i_1 \rangle).$$

and outputs this W_1 along with pwd , and HK_1 as internal state of P_i . Note that this computation of W_1 is identical to the computation of W_1 in the computation of sk_1 (which is really output to \mathcal{Z} only when $\text{pwd}' = \text{pwd}$).

Without loss of generality, we can assume that in the real-world if the Adversary (or Environment \mathcal{Z}) corrupts an instance before the session key is output then the instance does not output any session key. This is so because the Adversary (or \mathcal{Z}) either sets the key for that session or can compute it from the internal state it broke into.

Server Corruption. On receiving a `Corrupt` call from \mathcal{Z} for server instance P_j in session `ssid`, the simulator \mathcal{S} first performs the steps in the section on Stealing Password File above. In particular this sets the value of `phash`. It then calls the `Corrupt` routine of $\mathcal{F}_{\text{apwKE}}$. If \mathcal{S} had already output a message to \mathcal{Z} , and not output `sk1` it computes

$$W_2 = \text{sim}_5(\text{CRS}_S, \text{trap}_S, \langle R_2, S_2/\mathbf{b}_S^{\text{phash}}, T_2, i_2 \rangle).$$

and outputs this W_2 along with `HK2` as internal state of P_j . Note that `pwd` is not given out.

Complexity of the simulator. Observe that on stealing the password file, the function `OfflineTestPwd` is only called once for each random oracle input, which was not already tested by calling `TestPwd`. Hence the number of unique password arguments passed to `TestPwd` and `OfflineTestPwd` of $\mathcal{F}_{\text{apwKE}}$ combined in the ideal world is at most the number of random oracle calls in the hybrid model.

Time complexity-wise, most of the simulator steps are $\log q$ -time, where q is the security parameter. Due to Step 4 of the simulator code, where for each of the m sessions, in the worst case, it might go through all the n random oracle calls, there is an additive component of $m * n * \log q$ time. So the simulator runs in $O(mn \log q)$ -time.

5.7 Proof of Indistinguishability

We now describe a series of experiments between a probabilistic polynomial time challenger \mathcal{C} and the environment \mathcal{Z} , starting with `Expt0` which we describe next. We will show that the view of \mathcal{Z} in `Expt0` is same as its view in UC-APAKE ideal-world setting with \mathcal{Z} interacting with $\mathcal{F}_{\text{apwKE}}$ and the UC-APAKE simulator \mathcal{S} described above in Section 5.6. We end with an experiment which is identical to the real world execution of the protocol in Figure 3. We prove that environment has negligible advantage in distinguishing between these series of experiments, leading to a proof of realization of $\mathcal{F}_{\text{apwKE}}$ by the protocol Π . Due to space limitations, in this version we only describe `Expt0`, and rest of the experiments and related proofs of indistinguishability can be found in the full version [JR16].

Here is the complete code in `Expt0` (stated as it's overall experiment with \mathcal{Z}):

1. Responding to a random oracle query on input m : If there is a record of the form (m, r) , then just return r . Otherwise, generate $r \leftarrow \mathbb{Z}_q$, record (m, r) and return r .

2. The challenger \mathcal{C} picks the CRS just as in the real world, except the QA-NIZK CRS-es are generated using the crs-simulators, which also generate simulator trapdoors $\text{trap}_{\mathcal{C}}, \text{trap}_{\mathcal{S}}$. It retains $a_1, a_2, \text{trap}_{\mathcal{C}}, \text{trap}_{\mathcal{S}}, \text{hk}_{\mathcal{C}}, \text{hk}_{\mathcal{S}}$ as trapdoors. Next, (on `StorePwdFile`) the challenger calls the random oracle with query $(\text{sid}, P_i, P_j, \text{pwd})$. It sets `phash` equal to the random oracle response and sets the server persistent state as $\mathbf{b}_{\mathcal{S}}^{\text{phash}}$. Define `PHASHISSET` to be true after either `StealPwdFile` has been called or the random oracle has been called with $(\text{sid}, P_i, P_j, \text{pwd})$ by the adversary, and false before. Define `PWDCALLED` to be true after the random oracle has been called with $(\text{sid}, P_i, P_j, \text{pwd})$ by the adversary, and false before.
3. On receiving $(\text{ClSession}, \text{sid}, \text{ssid}, P_i, P_j)$ from \mathcal{Z} , \mathcal{C} generates $(\text{HK}_1, \text{HP}_1) \leftarrow \text{ver}_{\mathcal{S}}.\text{hkgen}(\text{CRS}_{\mathcal{S}})$. Next, \mathcal{C} chooses r_1, r'_1, r''_1 at random, and sets $R_1 = \mathbf{g}^{r_1}$, $S_1 = \mathbf{g}^{r'_1}$ and $T_1 = \mathbf{g}^{r''_1}$. It then hands $(R_1, S_1, T_1, \text{HP}_1)$ to \mathcal{Z} on behalf of this instance.
4. On receiving $(R'_2, S'_2, T'_2, \text{HP}'_2)$ from \mathcal{Z} , intended for client session (P_i, ssid) (and assuming no corruption of this instance):
 - (a) If the received elements are either not in their respective groups, or are trivially 1, output $\text{sk}_1 \leftarrow \mathbb{G}_T$.
 - (b) If the message received is identical to message sent by \mathcal{C} in the same session (i.e. same `ssid`) on behalf of the peer, then output $\text{sk}_1 \leftarrow \mathbb{G}_T$ (unless the simulation of peer also received a legitimate message and its key has already been set, in which case the same key is used to output sk_1 here).
 - (c) If `PHASHISSET` is false, then output $\text{sk}_1 \leftarrow \mathbb{G}_T$.
 - (d) Compute: $i'_2 = \mathcal{H}(\text{sid}, \text{ssid}, P_j, P_i, R'_2, S'_2, \text{HP}'_2)$. If $S'_2 = R_2^{a_2} \mathbf{b}_{\mathcal{S}}^{\text{phash}}$ and $T'_2 = \text{sphf}_{\mathcal{S}}.\text{privH}(\text{hk}_{\mathcal{S}}, \langle R'_2, S'_2 / \mathbf{b}_{\mathcal{S}}^{\text{phash}} \rangle, i'_2)$, compute:

$$W_1 = \text{sim}_{\mathcal{C}}(\text{CRS}_{\mathcal{C}}, \text{trap}_{\mathcal{C}}, \langle R_1, S_1, \mathbf{b}_{\mathcal{S}}^{\text{phash}}, T_1, i_1 \rangle).$$

Output:

$$\text{sk}_1 = \text{ver}_{\mathcal{S}}.\text{privH}(\text{HK}_1, \langle R'_2, S'_2 / \mathbf{b}_{\mathcal{S}}^{\text{phash}}, T'_2, i'_2 \rangle) \cdot \text{ver}_{\mathcal{C}}.\text{pubH}(\text{HP}'_2, W_1)$$

- (e) If the above check failed then output $\text{sk}_1 \leftarrow \mathbb{G}_T$.
5. On a `Corrupt` call for client P_i , output `pwd`. If Step 3 has already happened then also output HK_1 and $W_1 = \text{sim}_{\mathcal{C}}(\text{CRS}_{\mathcal{C}}, \text{trap}_{\mathcal{C}}, \langle R_1, S_1, \mathbf{b}_{\mathcal{S}}^{\text{phash}}, T_1, i_1 \rangle)$.
6. On receiving $(\text{SrvSession}, \text{sid}, \text{ssid}, P_j, P_i)$ from \mathcal{Z} , follow steps symmetric to Step 4, swapping subscripts and languages accordingly and replacing the condition `PHASHISSET` by `PWDCALLED` in Step 4c. Also, in step 4d, the condition becomes: if $S'_1 = R_1^{a_1} \mathbf{b}_{\mathcal{C}}^{\text{phash}}$ and $T'_1 = \text{sphf}_{\mathcal{C}}.\text{privH}(\text{hk}_{\mathcal{C}}, \langle R'_1, S'_1, \mathbf{b}_{\mathcal{S}}^{\text{phash}} \rangle, i'_1)$,
7. On a `Corrupt` call for server P_j , if Step 3 has already happened then output HK_2 , and $W_2 = \text{sim}_{\mathcal{S}}(\text{CRS}_{\mathcal{S}}, \text{trap}_{\mathcal{S}}, \langle R_2, S_2 / \mathbf{b}_{\mathcal{S}}^{\text{phash}}, T_2, i_2 \rangle)$. Finally, execute a `StealPwdFile` call, as described below.

8. On a `StealPwdFile` call, return $\mathbf{b}_S^{\text{phash}}$ as the Server Persistent State to the adversary.

All outputs of sk_1 are also accompanied with sid, ssid (but are not mentioned above for ease of exposition).

Note that each instance has two asynchronous phases: a phase in which \mathcal{C} outputs R_1, S_1, \dots to \mathcal{Z} , and a phase where it receives a message from \mathcal{Z} . However, \mathcal{C} cannot output sk_1 until it has completed both phases. These orderings are dictated by \mathcal{Z} . We will consider two different kinds of temporal orderings. A temporal ordering of different instances based on the order in which \mathcal{C} outputs sk_1 in an instance will be called **temporal ordering by key output**. A temporal ordering of different instances based on the order in which \mathcal{C} outputs its first message (i.e. R_1, S_1, \dots) will be called **temporal ordering by message output**. It is easy to see that \mathcal{C} can dynamically compute both these orderings by maintaining a counter (for each ordering).

We now claim that the view of \mathcal{Z} in Expt_0 is statistically indistinguishable from its view in its combined interaction with $\mathcal{F}_{\text{apwKE}}$ and \mathcal{S} . The CRS is set identically by both \mathcal{C} and \mathcal{S} . While \mathcal{C} has access to `pwd` from the outset and sets up the random oracle output `phash` corresponding to $(\text{sid}, P_i, P_j, \text{ssid})$ at the beginning, \mathcal{S} doesn't have access to `pwd` at the beginning and hence defers this step till the point where either (1) a correct online guess has been made, (2) the password file was stolen and a correct offline guess was made, (3) the client was corrupted. In all these three cases the simulator gets to know `pwd` and has the chance to set `phash`. At the point when password file is stolen, the correct `pwd` may not have been guessed, but `phash` has to be set in order to output the server persistent state. In that case \mathcal{S} generates a random `phash`, remembers it and assigns it to the correct input when the actual password is queried. At all points, although their algorithms differ, we can see that \mathcal{C} and \mathcal{S} respond to random oracle queries identically.

Both \mathcal{C} and \mathcal{S} generate the client and server flows identically. In particular, observe that the condition `PHASHISSET` exactly captures the state of \mathcal{S} for a client session where it knows `phash` and can compute the relevant elements and keys. \mathcal{C} uses the condition `PHASHISSET` to do the same computations. Similarly for the server sessions with the condition `PWDCALLED`. The stronger condition for the server reflects the absence of the ‘‘Client Only Step’’ in the server sessions simulation. In the steps where a party receives a message from the adversary, both \mathcal{C} and \mathcal{S} end up computing keys identically. While \mathcal{C} directly checks by exponentiation with `phash` in the case that `pwd` was guessed correctly, \mathcal{S} goes through the list of random oracle calls to see which response was used for exponentiation as it may not know `pwd` or `phash` at this point.

Due to space limitations, the rest of the experiments and related proofs of indistinguishability are relegated to the full version [JR16].

References

- [BBC⁺13] Fabrice Benhamouda, Olivier Blazy, Céline Chevalier, David Pointcheval, and Damien Vergnaud. New techniques for SPHF's and efficient one-round PAKE protocols. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 449–475. Springer, Heidelberg, August 2013.
- [BBS04] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 41–55. Springer, Heidelberg, August 2004.
- [BC16a] Olivier Blazy and Céline Chevalier. Structure-preserving smooth projective hashing. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 339–369. Springer, Heidelberg, December 2016.
- [BC16b] Olivier Blazy and Céline Chevalier. Structure-preserving smooth projective hashing. Cryptology ePrint Archive, Report 2016/258, 2016. <http://eprint.iacr.org/2016/258>.
- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th ACM STOC*, pages 103–112. ACM Press, May 1988.
- [BM92] Steven M. Bellovin and Michael Merritt. Encrypted key exchange: Password-based protocols secure against dictionary attacks. In *1992 IEEE Symposium on Security and Privacy*, pages 72–84. IEEE Computer Society Press, May 1992.
- [BM93] Steven M. Bellovin and Michael Merritt. Augmented encrypted key exchange: A password-based protocol secure against dictionary attacks and password file compromise. In V. Ashby, editor, *ACM CCS 93*, pages 244–250. ACM Press, November 1993.
- [BMP00] Victor Boyko, Philip D. MacKenzie, and Sarvar Patel. Provably secure password-authenticated key exchange using Diffie-Hellman. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 156–171. Springer, Heidelberg, May 2000.
- [Boy09] Xavier Boyen. HPAKE: Password authentication secure against cross-site user impersonation. In Juan A. Garay, Atsuko Miyaji, and Akira Otsuka, editors, *CANS 09*, volume 5888 of *LNCS*, pages 279–298. Springer, Heidelberg, December 2009.
- [BP13] Fabrice Benhamouda and David Pointcheval. Verifier-based password-authenticated key exchange: New models and constructions. Cryptology ePrint Archive, Report 2013/833, 2013. <http://eprint.iacr.org/2013/833>.
- [BPR00] Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 139–155. Springer, Heidelberg, May 2000.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.

- [CHK⁺05] Ran Canetti, Shai Halevi, Jonathan Katz, Yehuda Lindell, and Philip D. MacKenzie. Universally composable password-based key exchange. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 404–421. Springer, Heidelberg, May 2005.
- [CR03] Ran Canetti and Tal Rabin. Universal composition with joint state. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 265–281. Springer, Heidelberg, August 2003.
- [CS02] Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 45–64. Springer, Heidelberg, April / May 2002.
- [EHK⁺13] Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge Villar. An algebraic framework for Diffie-Hellman assumptions. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147. Springer, Heidelberg, August 2013.
- [FLR⁺10] Marc Fischlin, Anja Lehmann, Thomas Ristenpart, Thomas Shrimpton, Martijn Stam, and Stefano Tessaro. Random oracles with(out) programmability. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 303–320. Springer, Heidelberg, December 2010.
- [GMR06] Craig Gentry, Philip MacKenzie, and Zulfikar Ramzan. A method for making password-based key exchange resilient to server compromise. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 142–159. Springer, Heidelberg, August 2006.
- [GS12] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. *SIAM J. Comput.*, vol.41(5), pages 1193–1232, 2012.
- [HK98] Shai Halevi and Hugo Krawczyk. Public-key cryptography and password protocols. In *ACM CCS 98*, pages 122–131. ACM Press, November 1998.
- [JG04] Shaoquan Jiang and Guang Gong. Password based key exchange with mutual authentication. In Helena Handschuh and Anwar Hasan, editors, *SAC 2004*, volume 3357 of *LNCS*, pages 267–279. Springer, Heidelberg, August 2004.
- [JKX18] Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 456–486. Springer, Heidelberg, April / May 2018.
- [JR12] Charanjit S. Jutla and Arnab Roy. Relatively-sound NIZKs and password-based key-exchange. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 485–503. Springer, Heidelberg, May 2012.
- [JR13] Charanjit S. Jutla and Arnab Roy. Shorter quasi-adaptive NIZK proofs for linear subspaces. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 1–20. Springer, Heidelberg, December 2013.
- [JR14] Charanjit S. Jutla and Arnab Roy. Switching lemma for bilinear tests and constant-size NIZK proofs for linear subspaces. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 295–312. Springer, Heidelberg, August 2014.
- [JR15] Charanjit S. Jutla and Arnab Roy. Dual-system simulation-soundness with applications to UC-PAKE and more. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 630–655. Springer, Heidelberg, November / December 2015.

- [JR16] Charanjit Jutla and Arnab Roy. Smooth NIZK arguments with applications to asymmetric UC-PAKE. Cryptology ePrint Archive, Report 2016/233, 2016. <http://eprint.iacr.org/2016/233>.
- [KV11] Jonathan Katz and Vinod Vaikuntanathan. Round-optimal password-based authenticated key exchange. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 293–310. Springer, Heidelberg, March 2011.
- [KW15] Eike Kiltz and Hoeteck Wee. Quasi-adaptive NIZK for linear subspaces revisited. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 101–128. Springer, Heidelberg, April 2015.
- [LPJY14] Benoît Libert, Thomas Peters, Marc Joye, and Moti Yung. Non-malleability from malleability: Simulation-sound quasi-adaptive NIZK proofs and CCA2-secure encryption from homomorphic signatures. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 514–532. Springer, Heidelberg, May 2014.
- [Mac01] Philip D. MacKenzie. More efficient password-authenticated key exchange. In David Naccache, editor, *CT-RSA 2001*, volume 2020 of *LNCS*, pages 361–377. Springer, Heidelberg, April 2001.
- [Wat09] Brent Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 619–636. Springer, Heidelberg, August 2009.