

Generalized Fuzzy Password-Authenticated Key Exchange from Error Correcting Codes

Jonathan Bootle¹, Sebastian Faller^{1,2}, Julia Hesse^{1*}, Kristina Hostáková², and Johannes Ottenhues^{3**}

¹ IBM Research Europe – Zurich

² ETH Zurich

³ University of St. Gallen

Abstract. Fuzzy Password-Authenticated Key Exchange (fuzzy PAKE) allows cryptographic keys to be generated from authentication data that is both fuzzy and of low entropy. The strong protection against offline attacks offered by fuzzy PAKE opens an interesting avenue towards secure biometric authentication, typo-tolerant password authentication, and automated IoT device pairing. Previous constructions of fuzzy PAKE are either based on Error Correcting Codes (ECC) or generic multi-party computation techniques such as Garbled Circuits. While ECC-based constructions are significantly more efficient, they rely on multiple special properties of error correcting codes such as maximum distance separability and smoothness.

We contribute to the line of research on fuzzy PAKE in two ways. First, we identify a subtle but devastating gap in the security analysis of the currently most efficient fuzzy PAKE construction (Dupont et al., Eurocrypt 2018), allowing a man-in-the-middle attacker to test individual password characters. Second, we provide a new fuzzy PAKE scheme based on ECC and PAKE that provides a built-in protection against individual password character guesses and requires fewer, more standard properties of the underlying ECC. Additionally, our construction offers better error correction capabilities than previous ECC-based fuzzy PAKEs.

Keywords: Attacks on Public-Key Constructions · Key Exchange Protocols · Password-Based Cryptography · UC Framework.

1 Introduction

Password-authenticated key exchange (PAKE) protocols allow two users to exchange symmetric keys from plaintext passwords only. PAKEs are a useful tool in unlocking smartcards (e.g., German ID card [BFK09,BDFK12,BFK13], FIDO2

* Author supported by the Swiss National Science Foundation (SNSF) under the AMBIZIONE grant “Cryptographic Protocols for Human Authentication and the IoT”

** Author partially funded by the EU-funded Marie Curie ITN TReSPAsS-ETN project under the grant agreement 860813

[BBCW21]), securing wireless networks [All22,CNPR22], and pairing IoT devices [HL19]. First formalized by Bellare and Merrit [BM92], PAKEs provide optimal protection of potentially low-entropy passwords, a feature that is called *resistance against offline dictionary attacks*. Essentially, a PAKE does not leak any information about the password used by the counterparty in case of password mismatch. This guarantee needs to hold against network adversaries observing the protocol execution as well as active attackers playing one user in the protocol. Such strong protection is vital for use cases of PAKE. Without it, any typing of, e.g., the 6-digit PIN restricting access to a German national ID card would already expose that PIN to brute-force attacks mounted through malicious card reader hardware or software. Building secure and efficient PAKEs that even feature strong universal composability guarantees is fairly simple. The general idea is to run a Diffie-Hellman key exchange, and either encrypt the Diffie-Hellman public keys with the password [BM92,BPR00], or derive the group generator from the password [Jab96,BMP00,Mac01,HS14,HL19]⁴.

Unfortunately, in many application scenarios of PAKE protocols, usability can hinder their adoption. Manually entering PINs in stressful situations, such as unlocking one’s smart ID card next to a waiting officer, or pairing a wearable IoT device while walking, is prone to repeated mistyping. As another example, consider IoT devices that want to automatically exchange a key with devices located in their close proximity. These devices can automatically derive passwords from sensor readings of their direct environment, but they likely end up with similar but not exactly matching bitstrings. In these situations, key exchange through PAKE fails even though the legitimate owner of the ID card was present, and even though the IoT devices were actually sitting next to each other.

In 2018, Dupont et al. [DHP⁺18] proposed a new cryptographic primitive called *fuzzy* PAKE (fPAKE), which allows two users to exchange a symmetric key from *similar* passwords. In addition to being a candidate to resolve the usability issues with password mistyping/mismatching in the aforementioned scenarios, fuzzy PAKE opens an interesting avenue towards secure biometric authentication, as it enables secure comparison of consecutive biometric readings *without leaking any other information about the biometric* than “match” or “no match”. To date, fuzzy PAKE is the only cryptographic primitive that can provide such optimal protection of data that is both noisy *and* of low entropy.

Dupont et al. [DHP⁺18] gave two constructions of fuzzy PAKE, one relying on Garbled Circuits and one relying on error-correcting codes (ECC)⁵ and PAKE, using Hamming distance as a metric for similarity of passwords. The ECC based construction is to the best of our knowledge the most efficient fuzzy PAKE construction in the literature [DHP⁺18] so far. At a high level, the proto-

⁴ [Hv22] provides an excellent overview of PAKEs in the literature, also mentioning other approaches than building PAKEs through Diffie-Hellman.

⁵ Dupont et al. [DHP⁺18] use the terminology of *robust secret sharing* (RSS) instead of error-correcting codes, and show how to instantiate RSS with ECC. In this work, we state their construction in terms of an ECC, as it enables better comparison with our protocol.

col works as follows. One party takes the role of the sender, chooses a symmetric key s , encodes it with the error-correcting code to get $C \leftarrow \text{Enc}(s)$, XORs the resulting codeword with the password to get $E \leftarrow C \oplus \text{pw}$, and sends E to the receiver. The receiver XORs the received E with their own password to get $C' \leftarrow E \oplus \text{pw}'$, and attempts to decode C' to retrieve s . Depending on whether pw and pw' were similar enough, the codeword C' is close enough to C to uniquely decode to s . The error correction capability of the fuzzy PAKE is hence directly related to the error correction threshold of the code. However, the protocol as described above would be prone to offline attacks: the receiver can repeatedly decode with many password guesses pw' , yielding a list of candidate keys. To prevent such a dictionary attack, [DHP⁺18] introduce a “password expansion” step prior to the encoding of the key. Here, the two parties run a PAKE on *every character* of their passwords pw and pw' , yielding key vectors K and K' that match in exactly the same entries as pw and pw' , while the other entries are uniformly distributed. Because this expansion step requires active participation of both parties and the vector entries are PAKE keys that are hard to guess, one party alone cannot turn a password into a key vector. Hence, parties – malicious or honest – are limited to only one password guess per protocol run.

In fact, to prevent C from leaking information about the key s in case the receiver is malicious, C is computed by encoding s concatenated with an extra random input r . The error correcting code must satisfy multiple special properties, such as a *uniformity* property, whereby small subsets of codeword entries appear uniformly random, independently of s , and a *smoothness* property, whereby decoding an extremely noisy codeword results in a random key. Uniformity is a well-known property, satisfied by the popular Reed-Solomon codes, which has been extensively studied [CCG⁺07,CDBN15,Wei16,BCL22], and can be obtained via various transformations including the method presented in [DHP⁺18]. Smoothness is less standard, and [DHP⁺18] rely on maximum distance separable codes in order to achieve it.

1.1 Our contribution

First, we identify a subtle but devastating issue in the ECC-based fuzzy PAKE construction of [DHP⁺18]. In a bit more detail, we give an attack that allows the adversary to make guesses on individual password bits (or characters), allowing them to extract a user’s, e.g., n -bit password from only n executions of the protocol with that user. Our attack demonstrates that the measures taken by [DHP⁺18] to protect against such individual guesses, namely ensuring that an attacker either guesses “all-or-zero” bits of the password, subtly fails due to incorrect binding of the different protocol parts. Here, binding refers to ensuring that all messages come from the same sender. Although there exist generic techniques for binding together messages sent over unauthenticated channels [BCL⁺05], they cannot be applied to the construction of [DHP⁺18] in black-box way due to its modular layout.⁶

⁶ More precisely, and for the reader who is familiar with the Universal Composability framework [Can01]: In Section 3.2, we argue that the split transformation of Barak

The effect of our attack is that an active attacker learns one bit of an honest user’s password by simply messing with an honest protocol run between two users. While the security notion of fuzzy PAKE [DHP⁺18] does leak information about the password, i.e., it leaks the exact password of an honest user to an attacker, we stress that this attacker must be actively running a session with the honest user *and* use close password to the one of the attacked honest user. Our attack requires much less knowledge from the adversary. In more detail, it can be executed by a network adversary essentially flipping some bits of the exchanged messages between honest clients. The attacker only needs to make an assumption about the distance of the passwords used in the execution. For example, in settings where fuzzy PAKE is usually executed without the proper passwords, e.g, mistyping happens only rarely, every honest protocol run would reveal one bit about the secret password of an honest party to a network attacker. Deriving such information from honest protocol runs is prohibited by the security model of [DHP⁺18].

Reference	error corr. capability	rounds	msg size	Sender	Receiver	model
[DHP ⁺ 18]	$\lfloor \frac{n-k}{2} \rfloor$	1	$\mathcal{O}(n)$	$2n \text{ exp}$	$\text{Dec} + 2n \text{ exp}$	RO, IC, CRS
Our work , unique dec.	$\lfloor \frac{n-k}{2} \rfloor$	1	$\mathcal{O}(n)$	$2n \text{ exp}$	$\text{Dec} + 2n \text{ exp}$	RO, IC, CRS, SA
Our work , list dec.	$n - 1 - \lfloor \sqrt{(k-1)n} \rfloor$	1	$\mathcal{O}(n)$	$2n \text{ exp}$	$\text{LDec} + 2n \text{ exp}$	RO, IC, CRS, SA

Fig. 1. Comparison of ECC-based fuzzy PAKE constructions, using a Reed-Solomon code of rank k , and block length n equal to the size of the password. Sender and Receiver columns indicate computation complexity. Both protocols are instantiated with EKE2 [DHP⁺18] to derive concrete performances. RO, IC, and CRS are required for the security of EKE2. SA is split authentication. We stress that the ECC-based [DHP⁺18] cannot be considered a secure fuzzy PAKE (Section 3.2).

Our second contribution is a new fuzzy PAKE protocol based on error-correcting codes, which directly builds upon the ideas of [DHP⁺18] but (i) fixes the previous insecurities, (ii) provides better error-correction capabilities, and (iii) relies on fewer properties of the underlying code. These improvements come at the cost of less efficient computation and communication. We provide a high-level comparison in Figure 1. For (i), the idea is to ensure that the transformation of Barak et al. [BCL⁺05] applies to the whole protocol, simply by instantiating all building blocks that require interaction between the parties *before* applying the transformation. For (ii) we investigate how the (fixed) fuzzy PAKE protocol of [DHP⁺18], which applies unique decoding of the ECC, benefits from applying a non-unique decoding technique. Namely, we make use of *list decoding* which is a decoding technique that takes a codeword with errors as input, and that can produce a list of candidate codewords including the original one. When the

et al. [BCL⁺05] cannot be meaningfully applied to transform a *hybrid* protocol that assumes authenticated channels, to a version that is secure with unauthenticated channels. In a nutshell, the reason is that the hybrid building blocks are unaffected by the transformation and do not carry any authentication guarantees.

ECC threshold $n - k + 1$ (as % of n)	Password Length n									
	$n = 8$		$n = 32$		$n = 64$		$n = 128$		$n = 256$	
	[DHP ⁺ 18]	Ours	[DHP ⁺ 18]	Ours	[DHP ⁺ 18]	Ours	[DHP ⁺ 18]	Ours	[DHP ⁺ 18]	Ours
10%			1	1	2	3	5	6	12	12
25%	0	1	3	4	7	8	15	17	31	34
30%	0	1	4	4	9	10	18	20	37	41
50%	1	2	7	9	15	18	31	37	63	74

Fig. 2. Error correction capability for the ECC-based PAKE constructions of [DHP⁺18] and in this work using list decoding, for various password lengths n and a Reed-Solomon code with rank k and block length n . The table shows the number of password errors below which key exchange will succeed. The ECC threshold refers to the percentage of errors above which key exchange will fail without leaking information about the password. Gray colored cells show parameter settings for which our construction improves upon previous works.

number of errors is too large for unique decoding to work, list decoding can still produce reasonable-sized lists of codewords that still contain the original one. We give examples of the improved error correction capability of our scheme over previous works in Figure 2. For (iii) we leverage new ingredients in our construction in order to bypass the smoothness property completely.

Intuitively, such non-unique decoding introduces a “correctness error” into the fuzzy PAKE protocol of [DHP⁺18], since the receiver decodes a list of key candidates and has to guess which one is the key of the sender. Moreover, it is unclear how to generalize the smoothness property and its proof in [DHP⁺18] to the list decoding regime, due to the higher number of errors.

Our idea is to restore perfect correctness by letting the sender give a *hint* about the correct key to the receiver. This hint, however, needs to be carefully crafted so that it benefits the correctness of the key exchange without revealing too much information about the password or the key of the sender. For example, simply using a hash of the sender’s key is not possible. This is because during security proofs, the hint may have to be simulated, without knowledge of the key, against an adversary who does know the key and is able to check that whether the hint was computed correctly or not. Instead, the sender hashes the encoding C of s . Based on the uniformity of the error-correcting code, this hint contains sufficient entropy to prevent information leakage on the key. Furthermore, the fact that the hint uniquely determines the correct codeword actually allows us to dispense with the smoothness property completely. Since smoothness, which relied on maximum distance separability, is no longer needed, the space of possible codes which can be used in our construction is much larger, opening the door to codes with even better list-decoding capabilities and fuzzy PAKEs with even higher error tolerance in future.

Altogether, we are able to give a construction with better error correction capability than previous schemes in the literature. More formally, we prove the following in Theorem 2.

Our fuzzy PAKE using a Reed-Solomon code of rank k , block length n , list decoding, a hint as described above and signing under ephemeral public keys to achieve binding, is secure and has an error correction capability of

$$\delta = n - 1 - \sqrt{(k - 1)n}.$$

We can also apply our results to obtain a fixed version of the fuzzy PAKE of [DHP⁺18] et al. by using unique decoding instead of list decoding, offering a trade-off between runtime (unique decoding is generally faster than list decoding) and error correction capability. See Figure 2 for numerical examples of the error correction terms, and Appendix C for some simple timing information for the unique and list decoding algorithms.

Corollary 1. *Our fuzzy PAKE using a Reed-Solomon code of rank k , block length n , **unique** decoding, a hint as described above and signing under ephemeral public keys to achieve binding, is secure and has an error correction capability of $\lfloor \frac{n-k}{2} \rfloor$. The computational overhead over [DHP⁺18] is one signature per message, and one hash by the sender.*

Roadmap. In Section 2 we provide preliminaries on error-correcting codes and list decoding, and on implicit PAKE which is a building block of our and previous ECC-based fuzzy PAKE protocols. Section 3 recaps fuzzy PAKE including its security model in the UC framework, explains the fuzzy PAKE of [DHP⁺18] and presents an attack that allows to test password bits. In Section 4 we give our improved fuzzy PAKE protocol and prove its security.

1.2 Related Work

Prior to the introduction of fuzzy PAKE by Dupont et al. [DHP⁺18], several attempts to base cryptography on low-entropy or noisy shared data have been made. Information reconciliation [BBR88] and fuzzy extractors [DORS08,RW04] let two parties identify common randomness in shared noisy bitstrings. However, the identification comes at the price of leakage, such that these techniques cannot be used when the shared data is of low entropy, such as in the case of passwords. Canetti et al. [CFP⁺16] construct special-purpose schemes for securely comparing the Hamming distance of, e.g., biometric readings. However, the security of their construction again relies on the data having a certain min-entropy. The PAKE-based construction for comparison of biometric readings of Boyen et al. [BDK⁺05] does not protect against offline attacks on the biometric data and hence cannot reach the security level of fuzzy PAKEs.

Our construction improves the parameters of [DHP⁺18] by using list decoding instead of unique decoding, and using a hint to select the correct codeword. In fact, Cramer et al. [CDD⁺15] construct robust secret-sharing schemes using a similar idea, but they ensure that the correct secret is selected using an algebraic manipulation detection (AMD) code. Putting aside the subtle issues in [DHP⁺18] that we address in this paper, our construction is different from using the robust secret-sharing schemes of [CDD⁺15] in the [DHP⁺18] because

incorporating the hint does not constitute an AMD code, and also removes the smoothness requirement from [DHP⁺18]. See Appendix D for a more detailed discussion.

Biometric authentication is a very active area of research. In the following we list several works that aim at providing secure biometric authentication by introducing, e.g., additional trust anchors, or by leveraging the entropy in the biometric scans. BETA [ABM⁺21] generates a token to be used for authenticating a client to a server, by letting three client devices communicate over authenticated channels. In the fuzzy PAKE setting, each entity only needs one device and also a fuzzy PAKE is secure even in unauthenticated channels. Agrawal et al. [ABMR20] present a protocol for authenticating a user with their biometric by letting the user’s phone, an external terminal and a service provider communicate to compute the authentication result. Differences to fuzzy PAKE are that [ABMR20] has three parties instead of two, and their protocol uses cosine similarity instead of Hamming distance to compare the biometrics. Moreover, [ABMR20] focuses on authentication whereas a fuzzy PAKE achieves key exchange. Wang et al. [WHC⁺21] construct authenticated key change for secure messaging in which the users do not need to store their secret keys but instead derive them from their biometrics. [WHC⁺21] uses the biometric as an error term for LWE samples, which places very strong requirements on the biometric distribution. A fuzzy PAKE is stronger; its security cannot rely on the distribution of the authentication data, and hence our protocol is secure even when used with low entropy passwords (or biometrics). Jiang et al. [JLHG22] build key exchange from fuzzy data, which is also the goal of a fuzzy PAKE. However, similar to [WHC⁺21], [JLHG22] relies on high entropy of that fuzzy data. A fuzzy PAKE however does not make any assumptions about the entropy of the authentication data (=passwords). The tools used in [JLHG22] such as secure sketches are known to be inherently insecure when used on low entropy data. Erwig et al. [EHOR20] leverages robust secret sharing, which previous fuzzy PAKEs and our construction also rely on. Contrary to us, [EHOR20] builds an *asymmetric* PAKE, with a client and a server party where the latter is not allowed to store the password in the clear. While this is an extremely relevant setting in practise, their solutions are computationally expensive and the authors note that it seems infeasible to go beyond passwords of 40 bits. The works of Chatterjee et al. [CAA⁺16], [CWP⁺17] propose typo-tolerant password authentication systems. These password authentication systems tolerate mistyping on the user side, i.e., the user can successfully authenticate to some server even when using a password slightly different from the one they registered with. While [CAA⁺16] corrects a predefined set of most common typos (e.g., capitalization of first letter), the TypTop system [CWP⁺17] offers personalized adoption to frequent typos of users, meaning that the system will learn from past typos and apply certain rules whether to accept these errors in the future. The crucial difference to fuzzy PAKE is that both these typo-tolerant systems require the server party to learn the *cleartext* password of the user, as the password needs to be fed into the decision procedure as well as into the evolving “accepted ty-

pos” cache. The main goal of fuzzy PAKE is to prevent such transmission of clear-text passwords among the two parties. The concept of typo tolerance was subsequently carried over to PAKE [PC20]. Typo-tolerant PAKE is a system to log in to a server despite of typos in the password. The protocol in [PC20] requires storage and communication proportional to the number of close passwords, which limits the potential applications of the protocol. The storage and communication cost of our protocol only depends on the length of the password but not on the number of close passwords, which also makes it possible to use our protocol with biometrics instead of passwords.

In [RX23] it is demonstrated that the common notion of PAKE in the UC-framework does not automatically guarantee correctness. More concretely, the ideal functionality $\mathcal{F}_{\text{PAKE}}$ [CHK⁺05] does not guarantee that two honest parties that are not maliciously attacked do output the same key when they run on the same password. Several ways to overcome these definitional issues were proposed in [RX23]. The most simple way is to demand correctness as a separate property from the protocol. This style of listing several properties is usually not desirable in UC-protocols as the additionally demanded properties might not be preserved under composition. However, [RX23, Thm. 3] shows that any PAKE protocol that is correct *and* UC-realizes $\mathcal{F}_{\text{PAKE}}$ can be interpreted as a protocol that is secure in an enhanced UC PAKE formalization that ensures correctness (by modeling the man-in-the-middle adversary as a third protocol party). The observations of [RX23] rely on the way output is produced in $\mathcal{F}_{\text{PAKE}}$ and not on the way how the passwords are checked. Therefore, they also apply to our $\mathcal{F}_{\text{fPAKE}}$ (and $\mathcal{F}_{\text{iPAKE}}$) functionality and we have to assert correctness separately.

2 Preliminaries

Notation. We denote vectors in boldface, e.g., $\mathbf{V} \in \mathbb{F}^n$. Let $Q \subseteq [n] := \{1, \dots, n\}$. Then, we denote with $\mathbf{V}|_Q := (\mathbf{V}_i)_{i \in Q} \in \mathbb{F}^{|Q|}$ the restriction of \mathbf{V} to Q . We write PPT for probabilistic polynomial time.

2.1 Universal Composability (UC)

To formally define and prove security of fuzzy PAKEs, we use the UC framework of Canetti [Can01] as is standard in the PAKE literature. Our goal here is to introduce the basic terminology and notation of the UC framework at a very high level and only to the extent needed for the understanding of later sections. For a more accurate and formal explanation of the UC framework, we refer the reader to the work of Canetti [Can01].

In our case, we always consider protocols between two parties (which we typically denote \mathcal{P}_0 and \mathcal{P}_1) that are executed in the presence of a PPT adversary \mathcal{A} who can corrupt any party at the beginning of the protocol execution (i.e., we consider so-called static corruption). By corruption we mean that the adversary gets a full control over the corrupt parties and learns their internal state. Parties and the adversary \mathcal{A} get their inputs from a special entity, called the *environment*

\mathcal{Z} , which represents everything external to the protocol execution. \mathcal{Z} also receives outputs from both parties and the adversary. The execution of a protocol in presence of an adversary \mathcal{A} is typically referred to as the “real world”.

In the UC framework, security requirements of a protocol are defined via *ideal functionalities*. At a very high level, an ideal functionality \mathcal{F} defines the intended input/output behaviour of parties, the allowed leakage of the protocol and the influence of an adversary on the protocol. \mathcal{F} communicates with the environment through *dummy parties* $\mathcal{P}_0, \mathcal{P}_1$ that simply forward messages between \mathcal{Z} and \mathcal{F} . The functionality additionally communicates with an adversary which is typically called the *simulator* and denoted \mathcal{S} . The execution of \mathcal{F} in the presence of a simulator \mathcal{S} is typically referred to as the “ideal world”. We say that a protocol π UC-realizes (or UC-emulates) an ideal functionality \mathcal{F} if for any PPT adversary \mathcal{A} , there exists a PPT simulator \mathcal{S} such that for any PPT environment \mathcal{Z} the ideal and real worlds are indistinguishable except for negligible probability. In other words, any attack possible on the protocol π can be simulated as an attack on the ideal functionality.

One of the main benefits of the UC framework is that it natively supports protocol composition and hence allows for natural modularization of protocol designs. Concretely, parties in a protocol can communicate with *hybrid* ideal functionalities $\mathcal{H}_1, \mathcal{H}_2, \dots$. In such a case, we say that the protocol is defined in the “ $(\mathcal{H}_1, \mathcal{H}_2, \dots)$ -hybrid world”. The UC composition theorem then allows to securely replace the hybrid ideal functionalities with concrete protocols that UC-realize those functionalities.

2.2 Error Correcting Codes (ECC)

A *linear error-correcting code* \mathcal{C} over a finite field \mathbb{F}_q of order q with rank k and block length n is a linear subspace $\mathcal{C} \subseteq \mathbb{F}_q^n$ of dimension k . We can associate \mathcal{C} with an injective linear *encoding function* $\text{Enc}: \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$, where $\text{Im}(\text{Enc}) = \mathcal{C}$.

- The *Hamming distance* $d(\mathbf{V}, \mathbf{V}')$ between vectors $\mathbf{V}, \mathbf{V}' \in \mathbb{F}_q^n$ is defined by

$$d(\mathbf{V}, \mathbf{V}') := |\{i \in [n]: \mathbf{V}_i \neq \mathbf{V}'_i\}| .$$

- The *minimum distance* $d(\mathcal{C})$ of a linear code \mathcal{C} is defined as

$$d(\mathcal{C}) := \min_{\mathbf{C}, \mathbf{C}' \in \mathcal{C}, \mathbf{C} \neq \mathbf{C}'} d(\mathbf{C}, \mathbf{C}') .$$

A linear code over \mathbb{F}_q with rank k , blocklength n , and minimum distance d is referred to as an $[n, k, d]_q$ code. The “ q ” will usually be omitted.

Error Correction. Let $e \in \mathbb{Z}$ with $0 \leq e < d/2$. A $[n, k, d]$ -code \mathcal{C} is said to be *e -error-correcting* if there exists a function $\text{Dec}: \mathbb{F}_q^n \rightarrow \mathbb{F}_q^k$ such that for all $\mathbf{V} \in \mathbb{F}_q^n$, $\mathbf{M} \in \mathbb{F}_q^k$ with $d(\mathbf{V}, \text{Enc}(\mathbf{M})) \leq e$, we have $\text{Dec}(\mathbf{V}) = \mathbf{M}$. It is well-known that an $[n, k, d]$ -code is $\lfloor \frac{d-1}{2} \rfloor$ -error correcting.

Further, we say that a code is *efficiently* e -error-correcting if Dec can be computed in polynomial time in n . Note that *all* linear codes are efficiently 0-error correcting, as the message associated with an error-free codeword can simply be recovered via Gaussian elimination.

List Decoding. Let $e \in \mathbb{Z}$ with $0 \leq e \leq n$. Let $\ell \in \mathbb{N}$ with $\ell \geq 1$. An $[n, k, d]$ -code \mathcal{C} is said to be (e, ℓ) -*list-decodable* if for all $\mathbf{V} \in \mathbb{F}_q^n$, the list L of codewords $\mathbf{C} \in \mathcal{C}$ with $d(\mathbf{V}, \mathbf{C}) \leq e$ contains at most ℓ codewords.

Further, we say that a code is *efficiently* (e, ℓ) -list-decodable if there is an algorithm LDec which computes the list of codewords in polynomial time in n . This necessarily implies that ℓ is polynomial in n . In this case, we will often simply say that \mathcal{C} is *efficiently* e -list-decodable.

It will be convenient to use shorthand such as “ $\mathbf{C} \in \text{LDec}(\mathbf{V})$ ” when a codeword \mathbf{C} is part of the output list when LDec is run on input \mathbf{V} .

2.3 Randomized Codes

We also consider codes with hiding properties, using terminology from [BCL22]. Let \mathcal{C} be a $[n, k, d]$ code. Let $k_M, k_R \in \mathbb{N}$ with $k_M + k_R = k$. We may consider a *randomized encoding function* $\overline{\text{Enc}}: \mathbb{F}_q^{k_M} \times \mathbb{F}_q^{k_R} \rightarrow \mathbb{F}_q^n$ defined by $\overline{\text{Enc}}(\mathbf{M}, \mathbf{R}) := \text{Enc}(\mathbf{M} \parallel \mathbf{R})$. When considering $\overline{\text{Enc}}$, we refer to \mathcal{C} as a *randomized linear code*.

Uniform Codes. A randomized linear code \mathcal{C} is said to be B -*query uniform* if for any set $Q \subseteq [n]$ with $|Q| \leq B$, the distribution of $\{\overline{\text{Enc}}(\mathbf{M}, \mathbf{R})|_Q : \mathbf{R} \leftarrow \mathbb{F}_q^{k_R}\}$ is uniform over \mathbb{F}_q^Q .

Example: Reed-Solomon Codes. Let $k, n \in \mathbb{N}$. Let \mathbb{F}_q be a field, let s_1, \dots, s_n be distinct points in \mathbb{F}_q , and let $S := \{s_1, \dots, s_n\}$. The *Reed-Solomon* code $\mathcal{RS}[n, k, S]$ is defined as

$$\mathcal{RS}[n, k, S] := \{(p(s_1), \dots, p(s_n)) \mid p \in \mathbb{F}_q[X], \deg(p) \leq k - 1\} ,$$

with associated encoding function

$$\text{Enc}: (M_0, \dots, M_{k-1}) \mapsto (p_M(s_1), \dots, p_M(s_n)) ,$$

where $p_M := \sum_{i=0}^{k-1} M_i X^i$. It is well known that $\mathcal{RS}[n, k, S]$ is a $[n, k, n - k + 1]$ code and that furthermore, when considering a randomised encoding function for some $k_R < k$, the Reed-Solomon code is k_R -query uniform.

We state results on the decodability and list-decodability of Reed-Solomon codes.

Lemma 1 (Berlekamp-Welch algorithm [KR07]). *The code $\mathcal{RS}[n, k, S]$ is $\lfloor \frac{n-k}{2} \rfloor$ -error correcting, and Dec is computable in time $O(n^2)$.*

Lemma 2 (Guruswami-Sudan algorithm [McE03,Gur06,Nie13]). *The code $\mathcal{RS}[n, k, S]$ is (e, ℓ) -list decodable with $e = n - 1 - \sqrt{(k-1)n}$ and $\ell = O(n^2)$. Moreover, the list L can be computed in time $O(n^9)$.*

Note that some references such as [McE03] allow LDec to produce codewords with distance greater than e as part of the list L as they can be easily discarded. We assume that this has already been done and L contains only codewords within Hamming distance e of the input vector.

While both decoding algorithms run in polynomial time, the asymptotic complexity of LDec in Lemma 2 is much higher than that of Dec in Lemma 1. The bound of $O(n^9)$ in Lemma 2 comes from combining Corollary 3.15 and Table 3.1 in [Nie13].

2.4 Implicit-Only PAKE

A PAKE is a cryptographic protocol π that allows two parties to agree on a shared key over an unauthenticated channel assuming that both parties know the same shared password (i.e., a potentially low entropy string), and guarantees that each party gets a freshly sampled random key if the input passwords are different. The absence of an authenticated channel and the low-entropy assumption on the shared password mean that an iPAKE has to take man-in-the-middle and offline dictionary attacks into account.

The main building block of our construction is an *implicit-only PAKE* (or iPAKE for short) put forward by Dupont et al. [DHP⁺18]. An iPAKE is a specific type of PAKE that only achieves *implicit* authentication with respect to the honest parties as well as the adversary. This means that at the end of the protocol execution, the two interacting parties do not know if they derived the same key or not. Moreover, an adversary launching an active attack on the protocol by trying to guess the password does not get any feedback on the correctness of their guesses. Dupont et al. [DHP⁺18] formally defined an iPAKE as an ideal functionality $\mathcal{F}_{\text{iPAKE}}$ which we recall in Figure 12.

Correctness. Roy and Xu [RX23] showed that ideal UC-functionalities for PAKE like $\mathcal{F}_{\text{PAKE}}$ by [CHK⁺05] do not offer correctness. This shortcoming of $\mathcal{F}_{\text{PAKE}}$ also applies to the $\mathcal{F}_{\text{iPAKE}}$ functionality. In a nutshell, all these functionalities offer an adversarial interface to *prevent* successful key exchange that is not subject to any restrictions. Since protocols that emulate a functionality can allow for the same attacks than the functionality, even a PAKE where parties always output random keys could securely realize a PAKE functionality. Further, Roy and Xu [RX23, Thm. 2] show that it is impossible to incorporate correctness into $\mathcal{F}_{\text{PAKE}}$ in the two-party setting. However, they show how to overcome this limitation by defining PAKE as a three-party protocol, where the third party is the man-in-the-middle adversary. This formulation is equivalent to demanding correctness from the PAKE protocol as a separate property [RX23, Thm. 3]. We therefore also demand that an iPAKE protocol has correctness. Let $(K, K') \leftarrow \text{out}_{\mathcal{A}, \pi} \langle \mathcal{P}_0(\text{pw}), \mathcal{P}_1(\text{pw}) \rangle$ denote that \mathcal{P}_0 outputs K and \mathcal{P}_1 outputs K' when

executing π on input $\mathbf{pw} \in \mathcal{D}$, respectively and in the presence of adversary \mathcal{A} . Here \mathcal{D} denotes a finite alphabet (or dictionary). We adapt their correctness definition to the setting of implicit-only PAKE here.

Definition 1. *We say an iPAKE protocol π is ε -correct if for all $\mathbf{pw} \in \mathcal{D}$, for some finite alphabet \mathcal{D} , and for all passive (i.e., honest-but-curious) adversaries \mathcal{A} , the probability that in an execution of π with honest parties $\mathcal{P}_0, \mathcal{P}_1$ the first party \mathcal{P}_0 on input \mathbf{pw} outputs a different key than the second party \mathcal{P}_1 on input \mathbf{pw} in the presence of \mathcal{A} is negligible, i.e.,*

$$\Pr[K \neq K' \mid (K, K') \leftarrow \text{out}_{\mathcal{A}, \pi}(\mathcal{P}_0(\mathbf{pw}), \mathcal{P}_1(\mathbf{pw}))] \leq \varepsilon(\lambda),$$

where the probability is taken over the random coins of $\mathcal{P}_0, \mathcal{P}_1$, and \mathcal{A} . We call ε the correctness error. We just say π is correct if it has negligible correctness error.

2.5 Split Authentication

Another building block in our fPAKE construction is *Split Authentication*, put forward by Barak et al. [BCL⁺05]. It is essentially a protocol that implements something very close to an authenticated channel in an unauthenticated setting. The adversary’s only additional ability is that they can run two separate executions of an authentication protocol (one with each party) without the two communicating parties realizing it. In particular, if both parties are honest, split authentication forces the adversary to decide at the beginning of the communication whether to launch a man-in-the-middle attack and run a separate protocol with each party, or to just observe the messages being exchanged and potentially delay their delivery.

Barak et al. [BCL⁺05] formalized Split Authentication as an ideal functionality \mathcal{F}_{SA} which we recall in Fig. 13.

3 Fuzzy Password-Authenticated Key Exchange

Fuzzy Password-Authenticated Key Exchange (fPAKE for short) is a cryptographic protocol allowing two parties to agree on a shared key in the following setting: the two parties are connected by an unauthenticated channel, but each of them holds a *noisy* version of a shared low-entropy password.

3.1 The $\mathcal{F}_{\text{fPAKE}}$ ideal functionality

Dupont et al. [DHP⁺18] formalized fPAKE as a UC ideal functionality $\mathcal{F}_{\text{fPAKE}}$ which we recall in Fig. 3. Here we describe $\mathcal{F}_{\text{fPAKE}}$ at a high level and refer the reader to the work of Dupont et al. for a more detailed discussion.

Parties $\mathcal{P}_0, \mathcal{P}_1$ initiate a password-authenticated key exchange by sending a `NEWSESSION` message to the ideal functionality including their version of the password. To allow for initiator-responder-style protocols, the ideal functionality

allows parties to specify their role, i.e., whether they are the Sender or the Receiver. Upon receiving a `NEWSESSION` message, the functionality records the received password pw and informs the adversary about the `NEWSESSION` request.

The adversary can then try to guess the recorded password through a `TESTPWD` query. This interface allows an adversary to mount an active attack on the protocol. The recorded pw is marked as **compromised** if adversary’s guess is “close enough” and **interrupted** otherwise. The adversary is informed about the result of their guess. To capture the closeness formally, $\mathcal{F}_{\text{fPAKE}}$ is parametrized by *error tolerance* δ and a metric d , and two passwords pw, pw' are considered close enough for key exchange to succeed if $d(\text{pw}, \text{pw}') \leq \delta$.

The functionality outputs a secret key to a party \mathcal{P}_i after receiving a `NEWKEY` instruction from the adversary. This captures the ability of network attackers to arbitrarily delay the termination of the protocol. $\mathcal{F}_{\text{fPAKE}}$ decides on the secret key being sent in three different ways (for a formal definition of each case, see Fig. 3.):

- (i) The secret key is specified by the adversary if the adversary successfully guessed \mathcal{P}_i ’s recorded password, which can happen either through a `TESTPWD` query, or by corrupting \mathcal{P}_{1-i} and submitting a password on their behalf. This case captures the fact that a protocol participant with a close enough password can bias the final key.
- (ii) If both parties are honest, their recorded passwords are δ -close and $\mathcal{F}_{\text{fPAKE}}$ already sent a secret key to party \mathcal{P}_{1-i} , then \mathcal{P}_i gets the same key as \mathcal{P}_{1-i} . This requirement ensures the core functionality of an fPAKE; namely, that honest parties with close enough passwords output the same key.
- (iii) In any other case, $\mathcal{F}_{\text{fPAKE}}$ sends a freshly sampled random key to ensures pseudorandomness of keys overall, and uniformity of keys output by parties with non-matching passwords.

Dupont et al. [DHP⁺18] additionally defined a modified `TESTPWD` interface that gives more freedom in defining the leakage to the adversary after a `TESTPWD` query. The functionality is now additionally parametrized by a *leakage threshold* $\gamma \geq \delta$ and three leakage functions L_c, L_m, L_f defining the leakage to the adversary depending on the closeness of their guess to the tested password. See Fig. 5 for a graphical explanation of the extended `TESTPWD` interface and Fig. 4 for a formal definition of the extension.

We are particularly interested in a `TESTPWD` interface that leaks positions at which passwords match if the passwords are sufficiently close (following [DHP⁺18], we call this leakage the *mask*). Our proof of security will rely on this leaked information. Formally, let $\mathcal{F}_{\text{fPAKE}}^M$ be the ideal functionality from Fig. 3 except that the `TESTPWD` interface is defined according to Fig. 4 with the following leakage functions (here $\text{pw} = (\text{pw}_1, \dots, \text{pw}_n)$ and $\text{pw}' = (\text{pw}'_1, \dots, \text{pw}'_n)$):

$$\begin{aligned} L_c^M(\text{pw}, \text{pw}') &:= (\{j \text{ s.t. } \text{pw}_j = \text{pw}'_j\}, \text{“correct guess”}), \\ L_m^M(\text{pw}, \text{pw}') &:= (\{j \text{ s.t. } \text{pw}_j \neq \text{pw}'_j\}, \text{“wrong guess”}), \\ L_f^M(\text{pw}, \text{pw}') &:= \text{“wrong guess”}. \end{aligned}$$

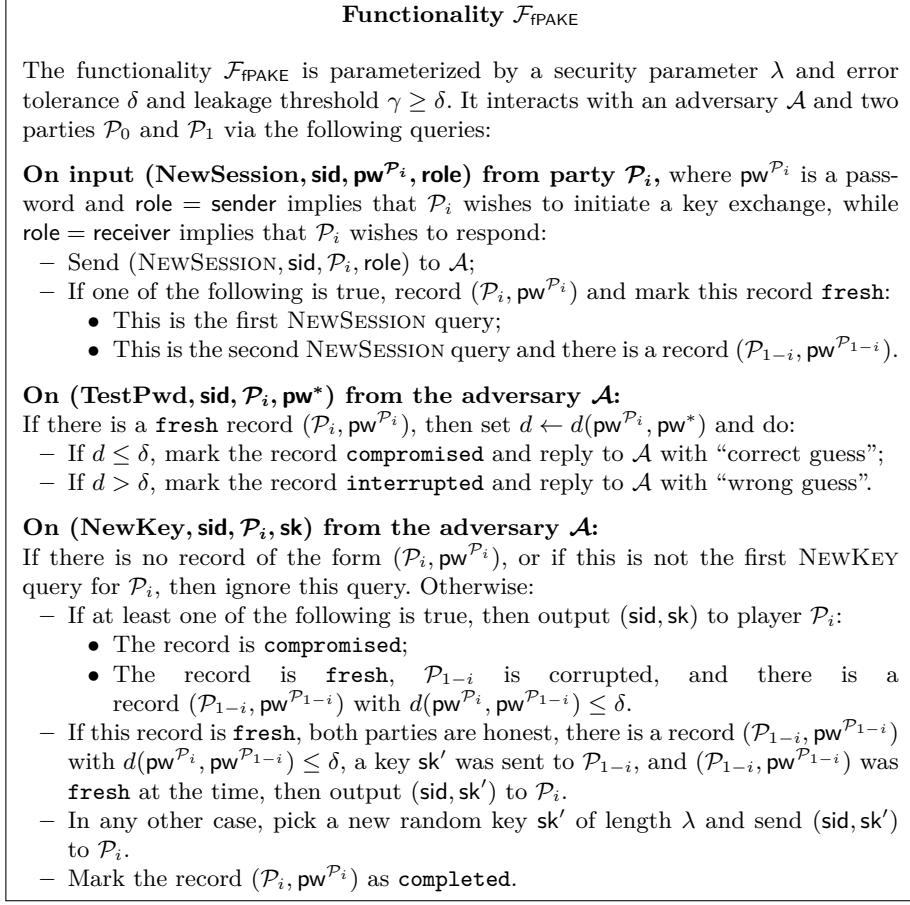


Fig. 3. Ideal functionality $\mathcal{F}_{\text{fPAKE}}$ [DHP⁺18]

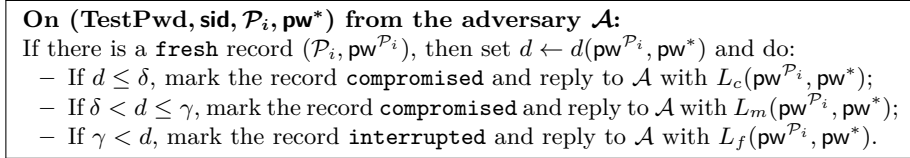


Fig. 4. A modified TESTPWD interface to allow for different leakage [DHP⁺18]

Correctness of fPAKE. For the same reason as for an iPAKE protocol (see discussion in Section 2.4), we also additionally require that a fuzzy PAKE protocol has correctness. Let $(K, K') \leftarrow \text{out}_{\mathcal{A}, \pi}(\mathcal{P}_0(\text{pw}), \mathcal{P}_1(\text{pw}'))$ denote that \mathcal{P}_0 outputs K and \mathcal{P}_1 outputs K' when executing π on input pw and pw' , respectively and in the presence of adversary \mathcal{A} .




Fig. 5. A graphical explanation of the relaxed TESTPWD interface of Figure 4. Here, “KE” stands for successful key exchange (e.g., when 0 to δ errors occur), and “no KE” stands for failure of key exchange, i.e., both parties outputting uniformly random strings.

Definition 2. We say a fuzzy PAKE protocol π is (δ, ε) -correct if for all $\mathbf{pw}, \mathbf{pw}' \in \mathcal{D}$ with $d(\mathbf{pw}, \mathbf{pw}') \leq \delta$, where \mathcal{D} is a finite alphabet, and for all passive (i.e., honest-but-curious) adversaries \mathcal{A} the probability that in an execution of π with honest parties $\mathcal{P}_0, \mathcal{P}_1$ the first party \mathcal{P}_0 on input \mathbf{pw} outputs a different key than the second party \mathcal{P}_1 on input \mathbf{pw}' in the presence of \mathcal{A} is bounded above:

$$\Pr[K \neq K' \mid (K, K') \leftarrow \text{out}_{\mathcal{A}, \pi}(\mathcal{P}_0(\mathbf{pw}), \mathcal{P}_1(\mathbf{pw}'))] \leq \varepsilon(\lambda),$$

where the probability is taken over the random coins of $\mathcal{P}_0, \mathcal{P}_1$, and \mathcal{A} . We call ε the correctness error. We just say π is δ -correct if it has negligible correctness error.

3.2 On the insecurity of previous fuzzy PAKE constructions

Dupont et al. [DHP⁺18] proposed a concrete fPAKE protocol (which they call $\pi_{\text{fPAKE}}^{\text{RSS}}$) that uses an implicit PAKE (iPAKE) and an error-correcting code (ECC)⁷. In this section we explain why, in fact, the $\pi_{\text{fPAKE}}^{\text{RSS}}$ protocol does not realize the $\mathcal{F}_{\text{fPAKE}}$ functionality. In order to do so, we first need to discuss the main ideas of $\pi_{\text{fPAKE}}^{\text{RSS}}$. We refer the reader to the original work [DHP⁺18] for more details.

Description of $\pi_{\text{fPAKE}}^{\text{RSS}}$. Let $\mathbf{pw} = (\mathbf{pw}_1, \dots, \mathbf{pw}_n) \in \mathcal{D} = \mathcal{L}^n$, for some finite alphabet \mathcal{L} , be the input password of the Sender and $\mathbf{pw}' = (\mathbf{pw}'_1, \dots, \mathbf{pw}'_n) \in \mathcal{D} = \mathcal{L}^n$ the input password of the Receiver. We typically talk about passwords being vectors of *password characters*. In the first stage of the protocol, the two parties engage in n iPAKE protocols. In the t -th execution, the Sender’s input is the t -th password character \mathbf{pw}_t and the Receiver’s input is \mathbf{pw}'_t . Parties receive outputs $K_t, K'_t \in \mathbb{F}_q$ respectively. The iPAKE with the correctness property guarantees that if $\mathbf{pw}_t = \mathbf{pw}'_t$ and there is no attack, then $K_t = K'_t$. Otherwise the keys K_t and K'_t are independent. At the end of this stage, the parties hold *character keys* $\mathbf{K} := (K_t)_{t \in [n]}, \mathbf{K}' := (K'_t)_{t \in [n]} \in \mathbb{F}_q^n$ respectively.

⁷ As already discussed in the introduction, Dupont et al. [DHP⁺18] use the terminology of Robust Secret Sharing instead of ECC.

In the second stage, the Sender samples a random field element $U \leftarrow \mathbb{F}_q$ and encodes it using an error correcting code with a hiding property. Let $C \in \mathbb{F}_q^n$ denote the resulting codeword. The Sender now performs a one-time-pad and sends $E := C + K$ to the Receiver. The Receiver computes $C' := E - K'$ and decodes the obtained codeword to U' . The idea is that if K and K' only differ at a few positions, then both codewords C and C' decode to the same value thanks to the error correcting property of the code, and hence both parties output the same key.

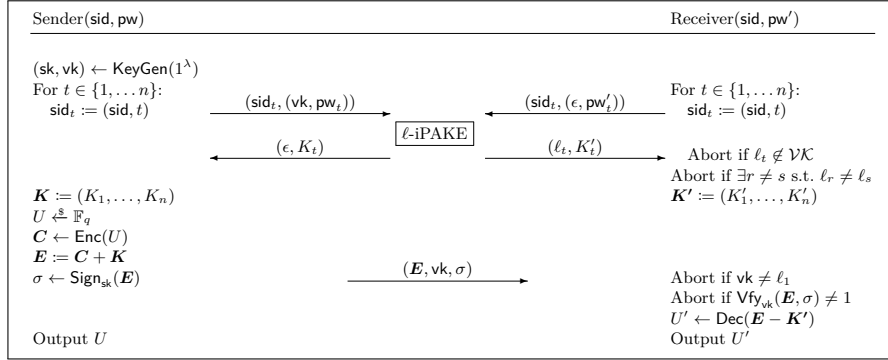


Fig. 6. Description of the $\pi_{\text{iPAKE}}^{\text{RSS}}$ protocol from [DHP⁺18] that uses n instances of an ℓ -iPAKE and a signature scheme ($\text{KeyGen}, \text{Sign}, \text{Vfy}$) whose verification key space is \mathcal{VK} and an error-correcting code (Enc, Dec).

In order to realize the $\mathcal{F}_{\text{iPAKE}}$ functionality, one needs to bind the n iPAKE executions and the final message together, as otherwise, active attacks on individual iPAKE instances allow to derive information about the corresponding password characters as observed already in [DHP⁺18]. To this end, the protocol $\pi_{\text{iPAKE}}^{\text{RSS}}$ actually makes use of a *labelled* iPAKE (or ℓ -iPAKE for short). In a nutshell, ℓ -iPAKE allows each party to additionally input a public label that serves as a public authentication string (i.e., any tampering with the label can be detected efficiently). The Sender then samples a signing key pair (sk, vk) at the beginning of the $\pi_{\text{iPAKE}}^{\text{RSS}}$ protocol and uses vk as their label in all ℓ -iPAKE executions. The Receiver checks that all the output labels of the ℓ -iPAKE executions are the same and aborts if this is not the case. In the last message, the Sender sends their signature $\sigma = \text{Sign}_{sk}(E)$ and vk to the Receiver in addition to E . The Receiver aborts if vk is not equal to the output labels of the ℓ -iPAKE instances or with the signature does not verify. See Fig. 6 for a schematic description of the protocol.

Attack on $\pi_{\text{iPAKE}}^{\text{RSS}}$. It turns out that the way ℓ -iPAKE instances and the last message were tightened together in the $\pi_{\text{iPAKE}}^{\text{RSS}}$ protocol is not sufficient to realize

the $\mathcal{F}_{\text{fPAKE}}$ functionality. Let us assume that both parties enter passwords pw, pw' whose distance is exactly at the threshold δ , i.e., $d(\text{pw}, \text{pw}') = \delta$. This means that the fPAKE would yield matching keys for both parties if there was no active attack on the protocol. The adversary can now make one of the ℓ -iPAKE key exchanges fail by making a wrong guess (without loss of generality, assume the first one). This perfectly emulates a situation where the first password values pw_1 and pw'_1 mismatch. Now, imagine a simulator that ensures indistinguishability of the ideal execution of $\mathcal{F}_{\text{fPAKE}}$ from the real protocol execution $\pi_{\text{fPAKE}}^{\text{RSS}}$ where parties make calls to the ℓ -iPAKE instances. The simulator's task is to figure out whether the described attack actually constitutes a successful DoS attack against the whole fPAKE or not, i.e., whether the passwords entered by both parties were already at the error tolerance δ and if $\text{pw}_1 = \text{pw}'_1$. If so, the result of the attack in the real world is that both parties compute random keys, otherwise they compute the same key. The only leakage about the passwords accessible by the simulator is through the `TESTPWD` interface of $\mathcal{F}_{\text{fPAKE}}$. Calling this interface, however, requires the simulator to provide a password γ -close to pw or pw' , as otherwise parties will already receive randomized keys in the ideal world because of the session getting `interrupted`. The probability that the simulator guesses the password depends on the dictionary size but is not negligibly close to 1. This results in a significant distinguishing advantage for the environment. We highlight that this attack has practical implications on the protocol. For the sake of simplicity, assume that the passwords must match exactly (i.e., $\delta = 0$), and that the password is encoded as a bit string. An attacker can completely retrieve a user's n -bit password from only n executions of the protocol with that user by guessing one bit per execution. Whenever key exchange is successful the bit was guessed correctly. For $\delta > 0$, the attack needs slightly more guesses as the attacker must ensure that the distance from his guess to the user's password is exactly at the δ boundary before the bit-wise guessing will work. We now describe the attack more formally.

Formalizing the attack on $\pi_{\text{fPAKE}}^{\text{RSS}}$ To formalize the attack, we need to show that for every simulator \mathcal{S} there exists an environment \mathcal{Z} that can distinguish the interaction with the hybrid world adversary and the protocol from the interaction with the simulator and the ideal functionality with non-negligible probability. In fact, in Fig. 7, we describe a distinguishing environment \mathcal{Z} that can distinguish the hybrid and ideal worlds with non-negligible probability for all PPT simulators.

Let us first look at the pseudocode in Fig. 7 and consider that \mathcal{Z} interacts with the hybrid world (i.e., the protocol $\pi_{\text{fPAKE}}^{\text{RSS}}$ making calls to $\mathcal{F}_{\ell\text{-iPAKE}}$, and the real world dummy adversary \mathcal{A} that does what \mathcal{Z} tells them to do). One can easily observe that if the environment's choice bit is set to $b = 0$, then the attack results in DoS (i.e., output keys sk and sk' are chosen uniformly at random) and if $b = 1$, the attack does not disturb the key exchange (i.e. $\text{sk} = \text{sk}'$). Hence, if \mathcal{Z} interacts with the hybrid world, they will correctly output HYBRID except with negligible probability (note that $\Pr[\text{sk} = \text{sk}' \mid \text{sk} \xleftarrow{\$} \mathbb{F}_q, \text{sk}' \xleftarrow{\$} \mathbb{F}_q] = \frac{1}{q}$, where $q \approx 2^\lambda$).

Let us now focus on the ideal world, where \mathcal{Z} interacts with the ideal functionality $\mathcal{F}_{\text{fPAKE}}$ and a simulator \mathcal{S} . The task of \mathcal{S} is to make $\mathcal{F}_{\text{fPAKE}}$ produce an output indistinguishable from the hybrid world described above. Let us first summarize all the information that \mathcal{S} gets. From the environment, \mathcal{S} receives $\text{pw}_1^* \in \mathbb{F}_p$ and we can assume that \mathcal{S} knows that this character is different from pw_1 . According to the definition of $\mathcal{F}_{\text{fPAKE}}$, \mathcal{S} gets no information about the passwords $\mathbf{pw}, \mathbf{pw}'$ from $\mathcal{F}_{\text{fPAKE}}$ by default. The only way \mathcal{S} could get information via $\mathcal{F}_{\text{fPAKE}}$ is through the TESTPWD interface, where \mathcal{S} has to guess \mathbf{pw} or \mathbf{pw}' . If \mathcal{S} guesses one of the passwords wrong, this password gets marked as **interrupted**, which means that the simulator loses all power over the keys being output to both parties (i.e. each party gets a randomly sampled key). In order to get useful information about the passwords via the TESTPWD interface without causing random keys, the simulator would have to make a close enough guess of one of the passwords. But since \mathcal{S} knows only that $\text{pw}_1 \neq \text{pw}_1^*$ and the passwords were chosen uniformly at random, the probability that \mathcal{S} succeeds in guessing is not negligibly close to 1 (the exact probability depends, of course, on the size of the dictionary and γ).

Hence, if \mathcal{S} could ensure that the ideal and hybrid worlds were indistinguishable, then \mathcal{S} would be able to guess b correctly with non-negligible probability. Since b was chosen uniformly at random from $\{0, 1\}$ by \mathcal{Z} , this is information-theoretically impossible.

1. Sample a bit $b \in \{0, 1\}$ uniformly at random.
2. Sample $\mathbf{pw} = (\text{pw}_1, \dots, \text{pw}_n) \in \mathbb{F}_p^n$ uniformly at random.
3. Choose a random $\mathbf{pw}' = (\text{pw}'_1, \dots, \text{pw}'_n) \in \mathbb{F}_p^n$ s.t. (1) $d(\mathbf{pw}, \mathbf{pw}') = \delta$ and (2) $\text{pw}'_1 = \text{pw}_1$ if $b = 0$ and $\text{pw}'_1 \neq \text{pw}_1$ if $b = 1$.
4. Choose $\text{pw}_1^* \in \mathbb{F}_p \setminus \{\text{pw}_1\}$.
5. Send (NEWSESSION, sid, \mathbf{pw} , Sender) to \mathcal{P}_0 and (NEWSESSION, sid, \mathbf{pw}' , Receiver) to party \mathcal{P}_1 , where sid chosen arbitrarily.
6. Upon receiving (NEWSESSION, ($\text{sid}, 1$), \mathcal{P}_0 , Sender) and (NEWSESSION, ($\text{sid}, 1$), \mathcal{P}_1 , Receiver), instruct the adversary to send (TESTPWD, ($\text{sid}, 1$), $\mathcal{P}_0, \text{pw}_1^*$).
7. Instruct the adversary to send (NEWKEY, ($\text{sid}, 1$), $\mathcal{P}_0, \text{sk}_1$) and then (NEWKEY, ($\text{sid}, 1$), $\mathcal{P}_1, \text{sk}_1$) for an arbitrary key sk_1 .
8. For every $t = 2, \dots, n$:
 - (a) Upon receiving (NEWSESSION, (sid, t), \mathcal{P}_0 , Sender) and (NEWSESSION, (sid, t), \mathcal{P}_1 , Receiver), instruct the adversary to send (NEWKEY, (sid, t), $\mathcal{P}_0, \text{sk}_t$) and then (NEWKEY, (sid, t), $\mathcal{P}_1, \text{sk}_t$) for an arbitrary key sk_t .
9. Upon receiving (sid, sk) from \mathcal{P}_0 and (sid, sk) from \mathcal{P}_1 distinguish the following two cases:
 - $b = 0$: If $\text{sk} \neq \text{sk}'$, output HYBRID. Otherwise, output IDEAL.
 - $b = 1$: If $\text{sk} = \text{sk}'$, output HYBRID. Otherwise, output IDEAL.

Fig. 7. Pseudocode of an environment \mathcal{Z} distinguishing the interaction with the $\pi_{\text{fPAKE}}^{\text{RSS}}$ protocol in the $\mathcal{F}_{\ell\text{-fPAKE}}$ hybrid world from the interaction with the $\mathcal{F}_{\text{fPAKE}}$ ideal functionality.

3.3 Towards repairing the previous construction

In their work, Dupont et al. mention an alternative way of binding the iPAKE instances and the last message together (see footnote 6 on page 26 of the eprint version of their paper [DHP⁺17]). Instead of using labels and a one-time signature scheme, they suggest letting the two parties sign every message using the split transformation put forward by Barak et al. [BCL⁺05]. A natural first step when trying to fix the $\pi_{\text{iPAKE}}^{\text{RSS}}$ protocol is hence to follow this alternative approach. Unfortunately, it turns out that applying the split transformation is not straightforward.

At a very high level, the split transformation is a generic way to transform a protocol realizing an ideal functionality \mathcal{F} assuming authenticated channels into a protocol that achieves the same without authenticated channels.⁸ Instead, the transformed protocol works in the \mathcal{F}_{SA} -hybrid world, where \mathcal{F}_{SA} is the ideal functionality for split authentication recalled in Section 2.5.

Intuitively, this seems to be exactly what is needed to bind the iPAKE instances and the message of $\pi_{\text{iPAKE}}^{\text{RSS}}$ together. Let $\pi_{\text{iPAKE}}^{\text{plain}}$ denote the protocol which is defined exactly as $\pi_{\text{iPAKE}}^{\text{RSS}}$ except that it uses $\mathcal{F}_{\text{iPAKE}}$ instead of $\mathcal{F}_{\ell\text{-iPAKE}}$ instances and the sender does not sign the last message. Assume that the sender and the receiver have an authenticated channel available which they can use for all the communication happening during all the n iPAKE executions and parties can use the same channel for the last message. Then the $\pi_{\text{iPAKE}}^{\text{plain}}$ protocol would be a secure fuzzy PAKE. Unfortunately, this intuition is misleading as it does not reflect the UC formalization of the $\pi_{\text{iPAKE}}^{\text{RSS}}$ protocol correctly.

The core of the problem is that the result of Barak et al. does not apply to hybrid protocols like $\pi_{\text{iPAKE}}^{\text{plain}}$. In a bit more detail, recall that the protocol $\pi_{\text{iPAKE}}^{\text{plain}}$ is defined in the $\mathcal{F}_{\text{iPAKE}}$ -hybrid world. This, in particular, means that the only message sent between the sender and the receiver in $\pi_{\text{iPAKE}}^{\text{plain}}$ is the last message. The rest of the protocol consists of the communication between a party and $\mathcal{F}_{\text{iPAKE}}$ ideal functionality. Sending (only) the last message over an authenticated channel would not help us to bind the n iPAKE executions together, so such a protocol could not realize $\mathcal{F}_{\text{iPAKE}}$. Hence, applying the split transformation does not have the desired effect of protecting against active attacks on individual password characters, like the one described in Fig. 7, in $\pi_{\text{iPAKE}}^{\text{RSS}}$.

Our approach. To address this issue, we define a protocol that does not work in the $\mathcal{F}_{\text{iPAKE}}$ -hybrid world. Instead, it lets the Sender and the Receiver run the code of a protocol π_{iPAKE} realizing $\mathcal{F}_{\text{iPAKE}}$ directly. In a bit more detail, let us assume that the Sender received $\text{pw} = (\text{pw}_1, \dots, \text{pw}_n)$ as input. For every $t \in [n]$, our protocol instructs the Sender to run the code of the π_{iPAKE} -sender on input pw_t . Whenever the π_{iPAKE} -sender would have sent a message to the π_{iPAKE} -receiver instance, the Sender does so via the channel available for the two parties in our protocol. The Receiver in our protocol is defined analogously. This way,

⁸ This is an oversimplified statement as the transformed protocol does not realize \mathcal{F} but its *split variant*, denoted $s\mathcal{F}$. See the work of Barak et al. [BCL⁺05] for more details.

all communication happening during the iPAKE executions and the last message are sent via the same channel between the Sender and the Receiver. Hence, the split transformation will have the desired effect. We formally prove this intuitive statement in the next section, where we define our final protocol precisely.

Potential alternative approach. Another way to resolve the insufficient binding of the iPAKE instances and the last message together could be to define a new building block, “batched iPAKE”. This new iPAKE primitive would take as input a vector $(\text{pw}_t)_{t \in [n]}$ and output a vector $(K_t)_{t \in [n]}$. An adversary would be allowed to make a TESTPWD query, but would always have to test all password characters $(\text{pw}_t^*)_{t \in [n]}$. Hence, the adversary would be forced to launch a man-in-the-middle attack either against all iPAKE instances, or none of them, which would guarantee the binding of the iPAKEs. In order to bind iPAKEs to the last message, we would require the “batch iPAKE” to be labelled. As in $\pi_{\text{iPAKE}}^{\text{RSS}}$, the sender would use their verification key as a label and then use the corresponding signing key to sign the last message.

4 Our Construction

General idea. The idea behind our construction is very simple: we modify the fuzzy PAKE of Dupont et al. [DHP⁺18] by letting the receiver apply list decoding instead of unique decoding. Since now the receiver ends up with a list of potential key candidates (or codewords, more precisely), we additionally let the sender compute a “hint” on which codeword is the correct one. More formally, the sender adds the hash of the original codeword $h := H(\mathbf{C})$ to the last message. The receiver applies list decoding which results in a list of candidates $\mathbf{C}^1, \dots, \mathbf{C}^\ell$. They then find the one that satisfies $h = H(\mathbf{C}^j)$ and compute $U \leftarrow \text{Dec}(\mathbf{C}^j)$. The value h is a hint on which of the candidates is the correct one. The benefit of our construction is that it can correct more errors in the password: for a Reed-Solomon code of rank k and block length n , unique decoding can correct up to $\lfloor \frac{n-k}{2} \rfloor$ errors (Lemma 1), while list decoding outputs a list containing the original codeword in the presence of $n - 1 - \sqrt{(k-1)n}$ errors. Fig. 2 provides several numerical examples of these terms, demonstrating that for growing password sizes, the list decoding approach yields a significantly better error correction capability than the unique decoding approach in [DHP⁺18].

On the other hand, it is *not* intuitively clear that the extra information passed from the sender to the receiver does not void any of the security guarantees of the protocol. To give an example, consider a different kind of hint $\bar{h} := H(\text{sk})$, where sk is the output key of the sender. Such a hint helps the receiver to choose its output key just as well as the hint $h := H(\mathbf{C})$. However, a fuzzy PAKE protocol sending \bar{h} cannot realize $\mathcal{F}_{\text{iPAKE}}$. The reason is that \bar{h} depends deterministically on secret information known to the sender (namely its output key sk), and hence it leaves no wiggle room for a simulator, who simply has to guess sk in order to simulate correctly. While the hint $H(\mathbf{C})$ does include some randomness contributed by the sender to the encoding algorithm Enc, it is not

straightforward to see that the resulting fuzzy PAKE is simulatable, meaning that it does not involuntarily reveal any information about the sender’s password through the hint.

Hence, in the remainder of this section, we carefully analyze the security of our fuzzy PAKE protocol given in Fig. 8.

Formal protocol description. The protocol is depicted in Fig. 8. To be able to fix previous issues of imperfect binding of messages, and as detailed in Section 3.2, we let parties send messages to each other over functionality \mathcal{F}_{SA} (the functionality was briefly discussed in Section 2.5 and is formally defined in Fig. 13). The protocol parties take a password from a dictionary $\mathcal{D} = \mathcal{L}^n$ as input, where \mathcal{L} is some finite alphabet. We write that the sender chooses $\mathbf{s} \xleftarrow{\$} \mathbb{F}_q^k$ and encodes it using $\text{Enc}(\mathbf{s})$. One can equivalently say that the sender chooses a message $s_0 \xleftarrow{\$} \mathbb{F}_q$ and randomness $\mathbf{r} \xleftarrow{\$} \mathbb{F}_q^{k-1}$ and uses the randomized encoding $\bar{\text{Enc}}(s_0, \mathbf{r})$. To make our protocol description modular, let π_j denote an instance of a protocol that UC-realizes $\mathcal{F}_{\text{IPAKE}}$. We write $X \leftarrow \pi_j(\mathcal{P}_i, \text{sid}, \text{pw}^{\mathcal{P}_i})$ to denote that \mathcal{P}_i outputs X when running π_j on sid and $\text{pw}^{\mathcal{P}_i}$ and $K \leftarrow \pi_j(\mathcal{P}_i, \text{sid}, X)$ to denote that \mathcal{P}_i outputs K as response to input X when running π_j on sid . For the sake of simplicity, we assume that the π_j protocols are one-round protocols. For multi-round protocols, one can send every message again over \mathcal{F}_{SA} to the respective other party, like in the first round. Note that it is important that the receiver only continues the computation with codewords that are within distance of δ to $\mathbf{E} \oplus \mathbf{K}'$. Otherwise, the receiver might output the same key as the sender, even when the passwords were not close enough. We ensure this by letting LDec only output codewords within distance of δ to the input. As explained in Section 2 this can easily be guaranteed by letting LDec discard codewords with distance $> \delta$. A more formal description of the protocol, using the interfaces of $\mathcal{F}_{\text{IPAKE}}$ can be found in Figure 14.

We start by proving the correctness of our protocol.

Theorem 1. *If \mathcal{C} is a randomized linear code which is efficiently δ -list decodable, H_1 is a ε_2 -collision-resistant hash function and π is ε_1 -correct and UC-realizes $\mathcal{F}_{\text{IPAKE}}$, then the protocol in Figure 8 is (δ, ε_0) -correct for*

$$\varepsilon_0 = n\varepsilon_1 + \varepsilon_2.$$

Remark 1 (Implications for previous results). Note that our construction implies (a fixed version of) the original fuzzy PAKE of Dupont et al. [DHP⁺18] when falling back to unique decoding algorithms. Indeed, one can simply set the list size of the list decoding algorithm to 1 and thus obtain unique decoding instead of list decoding as in [DHP⁺18], preserving our security analysis. There are two notable differences in this “fallback” version of our fuzzy PAKE and the one of [DHP⁺18]. First, our version applies a proper binding mechanism by using the split functionality \mathcal{F}_{SA} for sending messages [BCL⁺05], at the cost of one signature per message (see Appendix E for the concrete instantiation). Second, our security proof does not rely on the smoothness property of the underlying

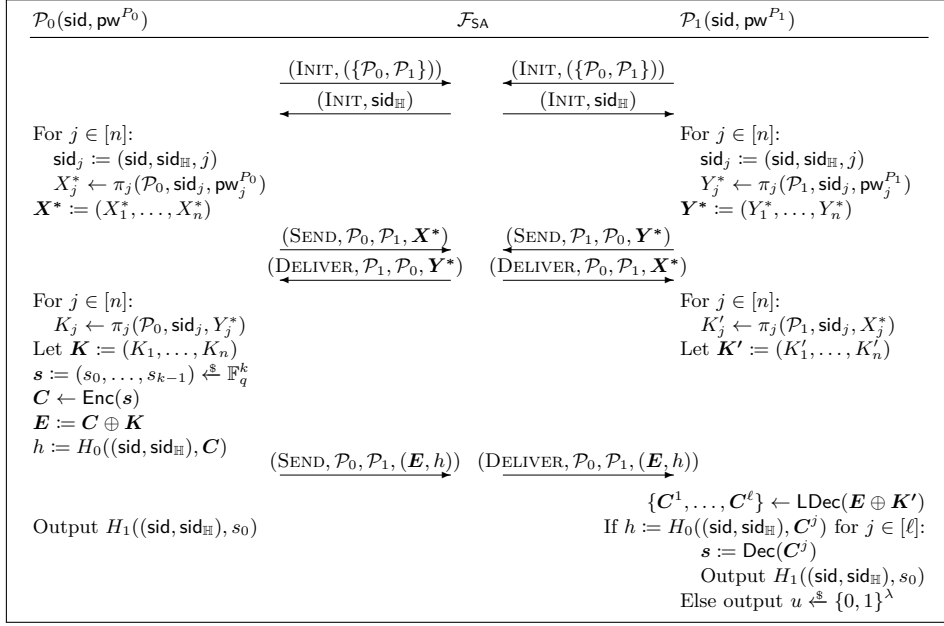


Fig. 8. The protocol in the \mathcal{F}_{SA} -hybrid model. We drop sid from all messages for brevity. We recall \mathcal{F}_{SA} in Figure 13 in the appendix, but for the sake of understanding our protocol it is enough to know that \mathcal{F}_{SA} is used to transmit the protocol messages. With π_j we denote an instances of a protocol that UC-realizes $\mathcal{F}_{\text{iPAKE}}$ (c.f. Figure 12). We write $X \leftarrow \pi_j(\mathcal{P}_i, \text{sid}, \text{pw}_j^{P_i})$ to denote that \mathcal{P}_i outputs X when running π_j on sid with password $\text{pw}_j^{P_i}$ and $K \leftarrow \pi_j(\mathcal{P}_i, \text{sid}, X)$ to denote that \mathcal{P}_i outputs K as response to input X when running π_j on sid .

code. This enables the usage of a variety of codes for which such properties are not well researched, or do not hold at all.

Proof. Assume that two honest parties $\mathcal{P}_0, \mathcal{P}_1$ execute the protocol from Fig. 8 on inputs pw^{P_0} and pw^{P_1} with $d(\text{pw}^{P_0}, \text{pw}^{P_1}) \leq \delta$, respectively. Further, assume that the adversary \mathcal{A} attacking this protocol is passive. Each instance π_j of π is ε_1 -correct by assumption (see Definition 1). We construct a sequence of hybrids $\mathcal{H}_0, \dots, \mathcal{H}_n$ where in the j -th hybrid, we abort the execution of the protocol if the two parties of π_j output different keys even though the respective inputs $\text{pw}_j^{P_0}$ and $\text{pw}_j^{P_1}$ are the same. By the ε_1 -correctness of π , \mathcal{H}_j differs from \mathcal{H}_{j-1} at most by ε_1 for each $j \in \{1, \dots, n\}$. Now in \mathcal{H}_n , as we have $d(\text{pw}^{P_0}, \text{pw}^{P_1}) \leq \delta$, there are at least $n - \delta$ instances of π where \mathcal{P}_0 and \mathcal{P}_1 run on matching inputs $\text{pw}_j^{P_0}$ and $\text{pw}_j^{P_1}$. Thus, we have $K_j = K'_j$ for at least $n - \delta$ keys. So the codeword $\mathbf{C}' = \mathbf{E} \oplus \mathbf{K}'$ that the receiver computes differs from the sender's \mathbf{C} on at most δ entries. As \mathcal{C} is δ -list decodable, the list that is output by LDec contains the same \mathbf{C} that the sender encoded. Now assume that there are two possible words \mathbf{C}, \mathbf{C}^* with $\mathbf{s} = \text{Dec}(\mathbf{C})$ and $\mathbf{s}^* = \text{Dec}(\mathbf{C}^*)$ such that $H_1((\text{sid}, \text{sid}_H), s_0) = H_1((\text{sid}, \text{sid}_H), s_0^*)$ with probability ε_2 . One can construct an adversary against the collision-resistance of H_1 that internally runs $\mathcal{P}_0(\text{pw}^{P_0})$ and $\mathcal{P}_1(\text{pw}^{P_1})$ and outputs s_0, s_0^* as above. We get

$$\Pr[K \neq K' \mid (K, K') \leftarrow \text{out}_{\mathcal{A}, \pi}(\mathcal{P}_0(\text{pw}^{P_0}), \mathcal{P}_1(\text{pw}^{P_1}))] \leq n\varepsilon_1 + \varepsilon_2.$$

Theorem 2. *If \mathcal{C} is a $\mathcal{RS}[n, k, S]$ code over \mathbb{F}_q for $q = 2^{\Omega(\lambda)}$, π UC-realizes $\mathcal{F}_{\text{PAKE}}$ and is correct, and H_0 and H_1 are modeled as random oracles then the protocol in Figure 8 UC-realizes $\mathcal{F}_{\text{PAKE}}^M$ in the \mathcal{F}_{SA} -hybrid model with error tolerance $\delta = n - 1 - \sqrt{(k-1)n}$ and leakage threshold $\gamma = n - k$.*

We discuss the proof of Theorem 2 in Section 5.

4.1 Comparison with properties and proof techniques in [DHP⁺18]

As previously discussed, prior work [DHP⁺18] constructs fuzzy PAKE protocols from linear error-correcting codes, using the fact that the errors in noisy codewords can be uniquely corrected up to the *unique decoding radius*. However, in this work, we wish to offer improved PAKE functionality by using list-decoding algorithms to correct errors beyond the unique decoding radius. The proof strategy employed in [DHP⁺18] does not readily generalise to this setting, and the presence of the hint h becomes crucial in the proof of Theorem 2, as we will explain.

The PAKE construction of [DHP⁺18] builds fuzzy PAKE protocols from robust secret-sharing (RSS) schemes with various special properties, which are in turn constructed from linear error-correcting codes. First, we summarize the RSS properties used in the PAKE construction of [DHP⁺18] and rephrase them as properties of error-correcting codes. Then, we describe the techniques used in the proof of the *reconstruction* and *smoothness* properties, and how this proof strategy leads to challenges when we generalize to list-decoding. Finally, we explain how the inclusion of the hint h allows us to circumvent these issues.

RSS properties in [DHP⁺18]. An RSS scheme consists of a sharing algorithm, which takes a message as input and produces n output shares, and a reconstruction algorithm, which, on input a list of (possibly corrupted) shares, returns a secret.

In this work, in line with our study of PAKE protocols from a decoding perspective, we phrase our algorithms in terms of randomised error-correcting codes as in Section 2.2, instead of RSS schemes. With this view, the randomised encoding algorithm $\overline{\text{Enc}}$ simply corresponds to a share algorithm with message space generalised from \mathbb{F} to \mathbb{F}^{k_M} , and explicit random input $\mathbf{R} \in \mathbb{F}^{k_R}$.

The PAKE construction of [DHP⁺18] uses the following properties of randomised encoding algorithms (with their corresponding RSS terminology from [DHP⁺18] included for clarity):

- B -query uniformity (Section 2.3), which corresponds to strong B -privacy for RSS;
- the existence of an efficient algorithm for e -error correction (Section 2.2), which corresponds to e -robustness for RSS;
- for all $E \subseteq [n]$ with $|E| \leq e$, all $\mathbf{M} \in \mathbb{F}^{k_M}$, the distribution of

$$\left\{ \text{Dec}(\mathbf{V}) : \mathbf{R} \leftarrow \mathbb{F}^{k_R}, \mathbf{V}|_E := \overline{\text{Enc}}(\mathbf{M}, \mathbf{R})|_E, \mathbf{V}|_{\bar{E}} \leftarrow \mathbb{F}^{|\bar{E}|} \right\} \quad (1)$$

is uniform over \mathbb{F}^{k_M} . This corresponds to e -smoothness for RSS;

- e -smoothness on random secrets. This is similar to e -smoothness except that \mathbf{M} is also sampled uniformly at random from \mathbb{F}^{k_M} .

The authors of [DHP⁺18] give a construction of a randomised linear code with each of the above properties from a standard linear error correcting code. The construction was originally given as an RSS.

Lemma 3 ([DHP⁺18, Lemma 5], restated). *Given an $[n+1, k, n-k+2]$ -linear code \mathcal{C}_0 , there is a $[n, k, n-k+1]$ randomised linear code \mathcal{C} with $k_M = 1$, $(k-1)$ -smoothness, $(k-1)$ -query uniformity, $\lfloor \frac{n+k}{2} \rfloor$ -error correction and $(\lfloor \frac{n+k}{2} \rfloor - 1)$ -smoothness on random secrets.*

Arguing smoothness and decoding properties for Lemma 3. The proof of Lemma 3 in [DHP⁺18] starts by defining \mathcal{C} to be the code obtained by puncturing \mathcal{C}_0 in its last coordinate. Then, $(k-1)$ -smoothness is proved with respect to the following decoding algorithm. On input a codeword in \mathcal{C}' , the decoding algorithm selects the subvector of k entries of the codeword and then inverts the submatrix of the generator matrix of \mathcal{C}' corresponding to these entries on the subvector to find the original message. Finally, this message vector is combined with the column of the generator matrix of \mathcal{C} corresponding to its last coordinate.

Intuitively, in the case of Reed-Solomon codes, this is like choosing k codeword entries of a punctured codeword with the last entry missing, and then interpolating to find out what the last entry should be.

With this decoding algorithm, it is relatively easy to prove smoothness, because \mathbf{V} in Equation (1) has at least $n-k+1$ random entries, so at least one

entry out of the k entries is random. Based on the MDS properties of \mathcal{C} , one can show that the “last coordinate” computed from the k selected entries must also be random.

In order for this decoding algorithm to be robust, it must be able to select k error-free positions. Otherwise, the submatrix inversion step will incorporate an error and produce an incorrect message. In the case of Reed-Solomon codes, this could be achieved by using another decoding algorithm (such as the Berlekamp-Welch algorithm) to identify these positions.

Therefore, proving the smoothness property relies heavily on the MDS properties of \mathcal{C} and the existence of efficient decoding algorithms to allow errors in the codeword to be located.

Moving beyond the unique decoding radius. In this work, we intend to construct a PAKE scheme which works for errors beyond the unique decoding radius. In this high-error setting, the natural replacement for a decoding algorithm is a list-decoding algorithm, which produces all possible codewords within a particular radius of the transmitted word. The first challenge for our PAKE algorithm to overcome was to allow the sender and receiver to agree on the same password, even though decoding may no longer produce a unique codeword which can be used to determine the password. We solve this problem by way of the hint $h = H(\mathcal{C})$.

Another challenge is the fact that smoothness, an important property in the security proofs of [DHP⁺18], relies heavily on the special decoding strategy introduced, which must apply a linear map to error-free positions in a codeword in order to ensure robustness. In this high-error setting, it may be impossible to identify a unique set of error-free positions in a codeword. Thus, the decoding strategy above, which involves inverting the generator matrix on an error-free set of codeword entries, may not be possible either. Therefore, it seems challenging to construct a list-decoding algorithm with a generalised smoothness property via a linear decoding algorithm that still returns the correct codeword somewhere in its list, as the linear decoding step may always be applied to strings with errors instead of error-free positions.

However, amazingly, the hint that we introduce lets us circumvent this issue. In order to prove our PAKE protocol secure, we require only that the list decoding algorithm will not produce the same codeword provided by the hint after many random errors are added to it. This follows almost trivially as adding a large number of random errors produces a string which is so far away from the original codeword that the list decoding algorithm cannot produce the original codeword. This relates to \mathbf{G}_8 and Lemma 4 in the security proof, discussed in Section 5.

5 Proof of Security

To prove that the protocol in Figure 8 UC-realizes $\mathcal{F}_{\text{fPAKE}}^M$, we have to show that every PPT environment \mathcal{Z} has at most negligible advantage in distinguishing

a real execution of the protocol from an ideal execution, with $\mathcal{F}_{\text{iPAKE}}^M$ and the simulator that we show in Figs. 9 to 10. We show this using a sequence of hybrid games. In the following, we give intuition for each of the game hops and prove the necessary lemmas. We defer the full proof of Theorem 2 to Appendix A.

On using other iPAKE protocols than EKE: The current simulator in Figures 9 to 10 is written keeping in mind an iPAKE protocol in which each party only sends one message. However, the simulator can be easily adapted to work with any iPAKE protocol as follows: Instead of forwarding only one set of messages $(\{Z_j^*\}_{j=1}^n)$ between \mathcal{Z} and the iPAKE simulators, \mathcal{S} would forward all iPAKE messages between \mathcal{Z} and the corresponding iPAKE simulator.

Proof (Sketch). We depict a UC execution of Fig. 8 in Fig. 14 in Appendix E, which makes all the interfaces of $\mathcal{F}_{\text{iPAKE}}$ explicit. The general proof strategy of showing indistinguishability of this protocol from $\mathcal{F}_{\text{iPAKE}}$ is to start with the real-world execution of the protocol from Fig. 14 and replace parts of the protocol with simulations, step-by-step. The goal is to reach the ideal-world execution where the simulator can work without using the secret passwords of the participating parties, but with help from the ideal functionality $\mathcal{F}_{\text{iPAKE}}^M$. In the following we give intuition for every game-hop we make:

- \mathbf{G}_0 : The proof starts with the real execution of the protocol.
- \mathbf{G}_1 : We move the whole execution into one machine which we will call the simulator. We also add dummy parties and a functionality that merely forwards the input of the parties (i.e. `NEWSESSION` calls) to the simulator. Furthermore, the functionality forwards the output of the simulator to the parties as `NEWKEY` messages. This allows the simulator to execute the protocol on the real inputs $\text{pw}^{\mathcal{P}_i}$ and $\text{pw}^{\mathcal{P}_{1-i}}$. The goal of the remaining steps is to add and adapt interfaces of the ideal functionality until we reach $\mathcal{F}_{\text{iPAKE}}^M$ and to make the simulator independent of the input passwords $\text{pw}^{\mathcal{P}_i}$, $\text{pw}^{\mathcal{P}_{1-i}}$ until we reach the simulator as described in Figs. 9 to 10.
- \mathbf{G}_2 : In this game the simulator aborts on collisions of H_0 .
- \mathbf{G}_3 : In this game we use the hypothesis that the iPAKE protocol π_j UC-realizes $\mathcal{F}_{\text{iPAKE}}$. There is a simulator $\mathcal{S}_{\text{iPAKE}}$ that interacts with $\mathcal{F}_{\text{iPAKE}}$ and simulates a protocol execution of π . Thus, we can replace each instance π_j of π by an instance $\mathcal{S}_{\text{iPAKE}}^j$ of $\mathcal{S}_{\text{iPAKE}}$. The simulator \mathcal{S} must also simulate $\mathcal{F}_{\text{iPAKE}}$ towards $\mathcal{S}_{\text{iPAKE}}$ to make sure that the simulation works. Note that \mathcal{S} still needs the input passwords $\text{pw}^{\mathcal{P}_i}$, $\text{pw}^{\mathcal{P}_{1-i}}$ to internally run $\mathcal{F}_{\text{iPAKE}}$.
- \mathbf{G}_4 : In this game we equip $\mathcal{F}_{\text{iPAKE}}^M$ with the `TESTPWD` interface. Also the `NEWKEY` and `TESTPWD` interfaces mark records. Also the `NEWKEY` interface reacts to the marked records. However, \mathcal{S} does not use the `TESTPWD` interfaces yet.
- \mathbf{G}_5 : $\mathcal{S}_{\text{iPAKE}}$ might send `TESTPWD` queries to $\mathcal{F}_{\text{iPAKE}}$. \mathcal{S} still uses the input passwords to answer these. The games \mathbf{G}_5 to \mathbf{G}_8 will change this. In \mathbf{G}_5 we argue that $\mathcal{S}_{\text{iPAKE}}$ will never send a `TESTPWD` query to $\mathcal{F}_{\text{iPAKE}}$ if both parties are honest and no man-in-the-middle attack is mounted. In this case, \mathcal{S} can ignore `TESTPWD` queries from $\mathcal{S}_{\text{iPAKE}}$.

On (NEWSESSION, $\text{sid}, \mathcal{P}_i, \text{role}$) from $\mathcal{F}_{\text{FPAKE}}$ or (INIT, $(\{\mathcal{P}_0, \mathcal{P}_1\}, \text{sid})$) from a corrupted \mathcal{P}_i to \mathcal{F}_{SA} :

- Create record $\langle \mathcal{P}_i, \text{sid}, \text{role}, \perp \rangle$ and mark it **fresh**.
- Send (INIT, $(\{\mathcal{P}_0, \mathcal{P}_1\}, \text{sid})$) to \mathcal{Z} in the name of \mathcal{F}_{SA} .

On (INIT, $\text{sid}, \mathcal{P}_i, \mathbb{H}, \text{sid}_{\mathbb{H}}$) from \mathcal{Z} in the name of \mathcal{A} to \mathcal{F}_{SA} :

- Perform the same checks as \mathcal{F}_{SA} on $\text{sid}_{\mathbb{H}}$ and \mathbb{H} and ignore the message if one of them fails.
- Record $\langle \mathbb{H}, \text{sid}_{\mathbb{H}} \rangle$.
- If \mathcal{P}_i is corrupted then provide output (INIT, $\text{sid}, \text{sid}_{\mathbb{H}}$) to \mathcal{P}_i . Regardless:
 - If $\mathbb{H} = \{\mathcal{P}_0, \mathcal{P}_1\}$ then set $\text{MitM} := 0$.
 - Else if $\mathbb{H} = \{\mathcal{P}_i\}$ set $\text{MitM} := 1$.
- Retrieve $\langle \mathcal{P}_i, \text{sid}, \text{role}, \perp \rangle$ and replace \perp by MitM .
- If there is not already an instance of $\mathcal{S}_{\text{FPAKE}}$ for session-id $(\text{sid}, \text{sid}_{\mathbb{H}}, j)$ for $j \in [n]$ create one. // I.e., create $\{\mathcal{S}_{\text{FPAKE}}^j\}_{j=1}^n$ if $\text{MitM} = 1$ or if $\text{MitM} = 0$ and $\{\mathcal{S}_{\text{FPAKE}}^j\}_{j=1}^n$ are not yet created for \mathcal{P}_{1-i} .
- For all $j \in [n]$ send (NEWSESSION, $(\text{sid}, \text{sid}_{\mathbb{H}}, j), \mathcal{P}_i, \text{role}$) to $\mathcal{S}_{\text{FPAKE}}^j$.
- Upon receiving output Z_j^* from all n instances of $\mathcal{S}_{\text{FPAKE}}$, send $(\text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, \text{msg} := (Z_1^*, \dots, Z_n^*))$ formatted as coming from \mathcal{F}_{SA} to \mathcal{Z} and store $\langle \text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, \text{msg} \rangle$.

On instruction from \mathcal{Z} to send (DELIVER, $\text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, \text{msg} := \{Z_j^*\}_{j=1}^n$) to \mathcal{F}_{SA} :

- If $\text{MitM} = 0$:
 - If message msg is the one that \mathcal{S} sent, i.e. \mathcal{S} has a record $\langle \text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, \text{msg} \rangle$:
 - * Give instruction to $\mathcal{S}_{\text{FPAKE}}^j$ to send message Z_j^* to \mathcal{P}_{1-i} in the name of \mathcal{P}_i .
 - * Remove one appearance of the record $\langle \text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, \text{msg} \rangle$.
 - Else: do nothing. // message manipulated
- Else // if $\text{MitM} = 1$
 - Give instruction to $\mathcal{S}_{\text{FPAKE}}^j$ to send message Z_j^* to \mathcal{P}_{1-i} in the name of \mathcal{P}_i .

On (TESTPWD, $(\text{sid}, \text{sid}_{\mathbb{H}}, j), \mathcal{P}_i, \text{pw}_j^i$) from $\mathcal{S}_{\text{FPAKE}}^j$ to $\mathcal{F}_{\text{FPAKE}}$
 – store $(\text{test-pwd}, \text{sid}, \mathcal{P}_i, j, \text{pw}_j^i)$. // Remember the query. No need to answer

On instruction from \mathcal{Z} to send (DELIVER, $\text{sid}, \mathcal{P}_{1-i}, \mathcal{P}_i, \text{msg} := (\mathbf{E}', h')$) to \mathcal{F}_{SA} :

- Retrieve record $\langle \mathcal{P}_i, \text{sid}, \text{role}, \text{MitM}, \mathbf{K} \rangle$ marked **waiting**.
- If $\text{MitM} = 0$ and both parties of session sid are honest:
 - Remove one appearance of the record $\langle \text{sid}, \mathcal{P}_{1-i}, \mathcal{P}_i, (\mathbf{E}', h') \rangle$. // Record exists because no MitM attack.
 - If $\text{role} = \text{receiver}$ then send (NEWKEY, $\text{sid}, \mathcal{P}_i, \perp$) to $\mathcal{F}_{\text{FPAKE}}$.
 - Else, ignore this message. // $\text{role} = \text{sender}$ does not expect this message
- Else // $\text{MitM} = 1$ or one of the parties is corrupted
 - If \mathcal{P}_i is honest and $\text{role} = \text{receiver}$ then
 - * Compute $\{C^1, \dots, C^\ell\} \leftarrow \text{LDec}(\mathbf{E}' \oplus \mathbf{K})$.
 - * If $h' = H_0((\text{sid}, \text{sid}_{\mathbb{H}}), C^j)$ for $j \in [\ell]$ then set $\mathbf{s} := \text{Dec}(C^j)$.
 - * If the record $\langle \mathcal{P}_i, \text{sid}, \text{role}, \text{MitM} \rangle$ is marked as **accident**, then choose $\widehat{\mathbf{sk}} \xleftarrow{\$} \{0, 1\}^\lambda$.
 - * Else if the record $\langle \mathcal{P}_i, \text{sid}, \text{role}, \text{MitM} \rangle$ is marked as **compromised** then set $\widehat{\mathbf{sk}} := H_1((\text{sid}, \text{sid}_{\mathbb{H}}), s_0)$.
 - * Else // record interrupted set $\widehat{\mathbf{sk}} := \perp$.
 - Regardless, send (NEWKEY, $\text{sid}, \mathcal{P}_i, \widehat{\mathbf{sk}}$) to $\mathcal{F}_{\text{FPAKE}}$.
 - Else if \mathcal{P}_i is corrupt then output (RECEIVED, $\text{sid}, \mathcal{P}_{1-i}, \mathcal{P}_i, (\mathbf{E}', h')$) to \mathcal{P}_i .
 - Else, ignore this message. // $\text{role} = \text{sender}$ does not expect this message

On a query $((\text{sid}, \text{sid}_{\mathbb{H}}), s_0)$ to H_1 :

- If there is a record $\langle H_1, (\text{sid}, \text{sid}_{\mathbb{H}}), s_0, k \rangle$ return k . Else draw $k \xleftarrow{\$} \{0, 1\}^\lambda$, record $\langle H_1, (\text{sid}, \text{sid}_{\mathbb{H}}), s_0, k \rangle$ and return k .

On a query $((\text{sid}, \text{sid}_{\mathbb{H}}), C)$ to H_0 :

- If there is a record $\langle H_0, (\text{sid}, \text{sid}_{\mathbb{H}}), C, \rho \rangle$ return ρ . Else draw $\rho \xleftarrow{\$} \{0, 1\}^\lambda$. If there is another record $\langle H_0, (\text{sid}, \text{sid}_{\mathbb{H}}), C', \rho' \rangle$ with $C \neq C'$ but $\rho = \rho'$ abort the simulation. Else record $\langle H_0, (\text{sid}, \text{sid}_{\mathbb{H}}), C, \rho \rangle$ and return ρ .

Fig. 9. The first part of the simulator \mathcal{S} for $\mathcal{F}_{\text{FPAKE}}$ and our Fuzzy PAKE with list-decoding protocol. \mathcal{S} simulates \mathcal{F}_{SA} and the random oracles H_0 and H_1 . If we write “retrieve record” and the record does not exist, \mathcal{S} ignores the message.

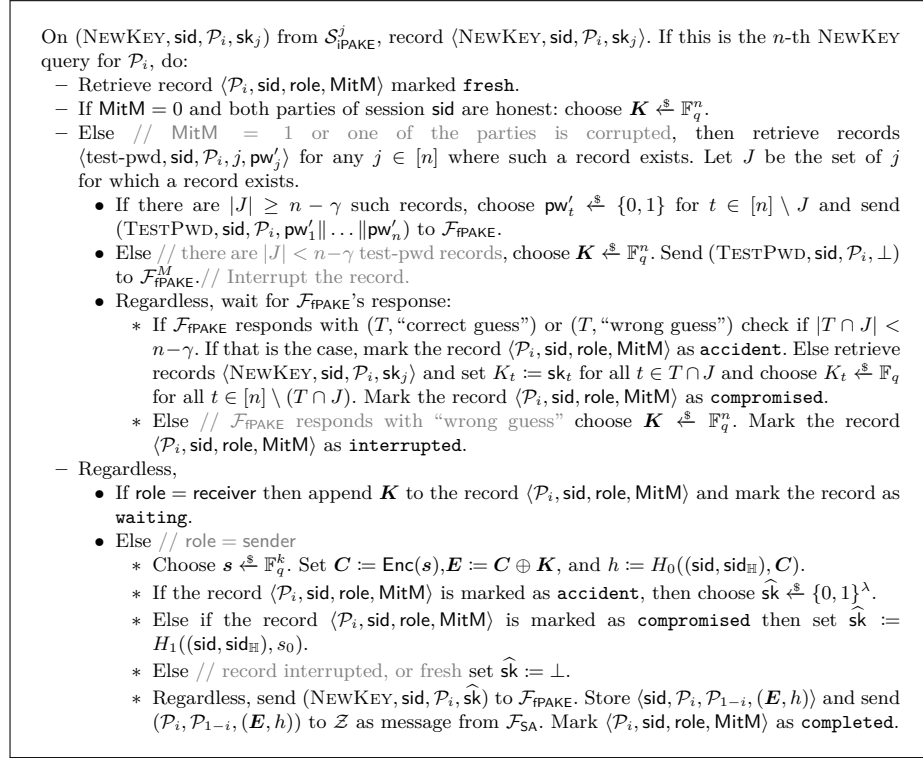


Fig. 10. The second part of the simulator \mathcal{S} for $\mathcal{F}_{\text{iPAKE}}$ and our Fuzzy PAKE with list-decoding protocol. \mathcal{S} simulates \mathcal{F}_{SA} and the random oracles H_0 and H_1 . If we write “retrieve record” and the record does not exist, \mathcal{S} ignores the message.

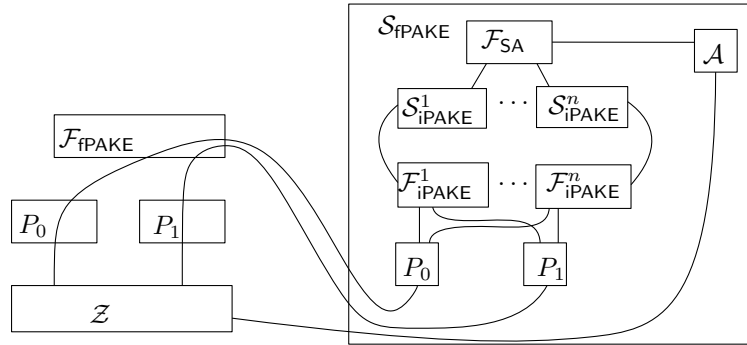


Fig. 11. This figure shows the layout of our simulator $\mathcal{S}_{\text{fPAKE}}$ in \mathbf{G}_3 in a setting without MitM attack.

- **G₆**: In this game we start dealing with attacks on an honest sender. Such an attack will lead to TESTPWD queries from the $\mathcal{S}_{\text{iPAKE}}$. \mathcal{S} can also send a TESTPWD query to $\mathcal{F}_{\text{iPAKE}}^M$. We leverage the fact that communication runs over the split authentication functionality. Because of \mathcal{F}_{SA} , the $\mathcal{F}_{\text{iPAKE}}$ instances of the sender run in a different authentication set than the iPAKE instances of the (malicious) receiver. Thus, $\mathcal{F}_{\text{iPAKE}}$ will never have a successful key exchange. Either the key is adversarially chosen or it is random. This is what mitigates the attack from Section 3.2.
The general idea in this game is that \mathcal{S} collects the TESTPWD queries of the iPAKE simulators, concatenates the passwords and uses this to ask a TESTPWD query to $\mathcal{F}_{\text{iPAKE}}$. If the password was close enough (i.e. within distance γ), then \mathcal{S} receives the position of the correct password characters as leakage from $\mathcal{F}_{\text{iPAKE}}^M$. Our simulator then uses this knowledge to simulate the $\mathcal{F}_{\text{iPAKE}}$ to \mathcal{Z} . If \mathcal{S} 's password guess was too far away, it will not receive any leakage information from $\mathcal{F}_{\text{iPAKE}}^M$. In that case, \mathcal{S} still uses the input passwords $\text{pw}^{\mathcal{P}_i}, \text{pw}^{\mathcal{P}_{1-i}}$ to produce output keys K_j for the iPAKE sessions. The only difference to **G₅** is that in the case of an attack on an honest sender, in **G₆** the sender gets a random key as output from the functionality and not from the simulator anymore. However, by the $(k-1)$ -query uniformity of the code, we can argue that the output was already independent and uniformly random in the previous game.
- **G₇**: In this game we change how \mathcal{S} simulates iPAKE keys of an honest sender when there is a MitM attack or the receiver is corrupt. In the case that the adversary attacked less than $n-\gamma$ iPAKE sessions or guessed less than $n-\gamma$ characters of the password correctly, \mathcal{S} just uses uniformly random iPAKE keys K_j for all $j \in [n]$. So even for the few successfully attacked iPAKE sessions the output key will be random. Again, we can use $(k-1)$ -query uniformity to argue that the few correct entries of \mathcal{C} that an adversary obtains already looked uniformly random before.
- **G₈**: In this game, we deal with attacks on honest receivers. Like in the previous game \mathcal{S} can use split authentication and the TESTPWD interface of $\mathcal{F}_{\text{iPAKE}}^M$ with leakage to deal with $n-\gamma$ or more attacked iPAKE sessions. In the case where less than $n-\gamma$ iPAKE sessions were attacked, \mathcal{S} does not try to decode but instead interrupts the record by sending a TESTPWD query to $\mathcal{F}_{\text{iPAKE}}^M$ with $\text{pw} = \perp$. Then it sends a NEWKEY query with $\text{sk} = \perp$ to $\mathcal{F}_{\text{iPAKE}}^M$, which causes the output of the receiver to be an independent, uniformly random value. One can use Lemma 4 to see that in the case of less than $n-\gamma$ password characters the LDec algorithm will not output anything that matches the hint h and, thus, the receiver's output was an independent uniform random value already even in the previous game.
- **G₉**: This game syntactically changes \mathcal{S} to remove usage of the provided input passwords $\text{pw}^{\mathcal{P}_i}, \text{pw}^{\mathcal{P}_{1-i}}$ in dealing with NEWSESSION and TESTPWD queries. We made sure in the previous games that \mathcal{S} does not need $\text{pw}^{\mathcal{P}_i}$ and $\text{pw}^{\mathcal{P}_{1-i}}$ anymore for those queries.
- **G₁₀**: In this game \mathcal{S} is modified such that honest and undisturbed parties produce output without \mathcal{S} using the provided inputs for the parties. For this,

\mathcal{S} can use the NEWKEY interface of $\mathcal{F}_{\text{iPAKE}}^M$. The functionality $\mathcal{F}_{\text{iPAKE}}^M$ will produce output according to the input passwords. If $d(\text{pw}, \text{pw}') \leq \delta$ then $\mathcal{F}_{\text{iPAKE}}^M$ produces matching passwords. Because of Lemma 6 the real-world protocol does the same. If $d(\text{pw}, \text{pw}') > \delta$ then $\mathcal{F}_{\text{iPAKE}}^M$ produces independent keys. By Lemma 5, the real-world protocol does the same. We also replace the message (\mathbf{E}, h) with a uniformly random message, which is indistinguishable for \mathcal{Z} as the iPAKE keys of honest parties are hidden from \mathcal{Z} .

- \mathbf{G}_{11} : This is the ideal-world experiment with $\mathcal{F}_{\text{iPAKE}}^M$ and the simulator from Figs. 9 to 10.

The following lemma essentially says that given a target vector \mathbf{C} , a vector $\tilde{\mathbf{C}}$ with at least $\gamma + 1 = n - k + 1$ random entries will have $d(\mathbf{C}, \tilde{\mathbf{C}}) \geq \gamma + 1$. This is necessary in our protocol to make sure that a malicious sender with a far away password cannot make the key exchange succeed. In the security proof of our protocol we use this lemma in \mathbf{G}_8 .

Lemma 4. For $q = 2^{\Omega(\lambda)}$, $\forall n, k \in \mathbb{N}$ with $n \geq k$, $\forall A \subseteq [n]$ with $|A| \leq k - 1$, $\forall \mathbf{C}, \tilde{\mathbf{C}} \in \mathbb{F}_q^n$ with $\tilde{\mathbf{C}}|_{\bar{A}} \xleftarrow{\$} \mathbb{F}_q^{n-|A|}$, we have

$$\Pr[d(\mathbf{C}, \tilde{\mathbf{C}}) \geq n - k + 1] \geq 1 - \text{negl}(\lambda).$$

The probability is taken over the randomness used to sample $\tilde{\mathbf{C}}|_{\bar{A}}$.

Proof. Remember that $\gamma = n - k$. Observe that $d(\mathbf{C}, \tilde{\mathbf{C}}) = d(\mathbf{C}|_A, \tilde{\mathbf{C}}|_A) + d(\mathbf{C}|_{\bar{A}}, \tilde{\mathbf{C}}|_{\bar{A}})$. The statement must even hold for $d(\mathbf{C}|_A, \tilde{\mathbf{C}}|_A) = 0$ and A maximally large, i.e., $|A| = k - 1$ and, thus, $|\bar{A}| = n - k + 1 = \gamma + 1$. What remains to be shown is that $d(\mathbf{C}|_{\bar{A}}, \tilde{\mathbf{C}}|_{\bar{A}}) \geq n - k + 1 = \gamma + 1$ with overwhelming probability. For $d(\mathbf{C}|_{\bar{A}}, \tilde{\mathbf{C}}|_{\bar{A}}) \geq \gamma + 1$ to hold, $\mathbf{C}|_{\bar{A}}$ and $\tilde{\mathbf{C}}|_{\bar{A}}$ must be distinct in all $\gamma + 1$ coordinates. The probability that $\mathbf{C}|_{\bar{A}}$ and $\tilde{\mathbf{C}}|_{\bar{A}}$ are distinct in a specific coordinate is $1 - \frac{1}{q}$. Thus, the probability that $\mathbf{C}|_{\bar{A}}$ and $\tilde{\mathbf{C}}|_{\bar{A}}$ are distinct in all $\gamma + 1$ coordinates is $(1 - \frac{1}{q})^{\gamma+1}$. For constant γ and $q = 2^\lambda$ this probability is overwhelming in λ .

The next lemma is similar to Lemma 4. One important difference is that in Lemma 4 \mathbf{C} was an arbitrary vector in \mathbb{F}_q^n , whereas in Lemma 5 \mathbf{C} is a valid codeword. We use this lemma to guarantee that in our protocol in a session with an honest sender and an honest receiver, whose passwords mismatch in more than δ positions, both parties get independent random keys. In particular this lemma guarantees that the receiver will not be able to decode the codeword, which will make them output a random key. This is important in the transition from \mathbf{G}_9 to \mathbf{G}_{10} .

Lemma 5. For $s \xleftarrow{\$} \mathbb{F}_q$, $\mathbf{C} \leftarrow \text{Enc}(s)$, $\tilde{\mathbf{C}} \in \mathbb{F}_q^n$ and $A \subseteq [n]$ such that $|A| \leq k - 1$, $\mathbf{C}|_A = \tilde{\mathbf{C}}|_A$, and $\tilde{\mathbf{C}}|_{\bar{A}} \xleftarrow{\$} \mathbb{F}_q^{n-|A|}$, we have that $\mathbf{C} \in \text{LDec}(\tilde{\mathbf{C}})$ with negligible probability.

Proof. From Lemma 4 we get that $d(\mathbf{C}, \tilde{\mathbf{C}}) \geq n - k + 1$ with overwhelming probability. Because $n - k + 1 > n - 1 - \sqrt{(k - 1)n} = \delta$ and $\text{LDec}(\tilde{\mathbf{C}})$ does not output candidates \mathbf{C}' with $d(\mathbf{C}, \tilde{\mathbf{C}}) > \delta$ by Lemma 2, we get that $\mathbf{C} \in \text{LDec}(\tilde{\mathbf{C}})$ with negligible probability.

Lemma 6 can be seen as the counterpart of Lemma 5, as we use it to guarantee that in our protocol two honest parties that have close enough passwords (i.e. $\leq \delta$) get the same session key. In detail, this lemma guarantees that the honest receiver is able to correctly decode $\mathbf{C}' = \mathbf{E} \oplus \mathbf{K}'$, which then leads to successful key exchange. We use this lemma in the transition from \mathbf{G}_9 to \mathbf{G}_{10} .

Lemma 6. *For $s \xleftarrow{\$} \mathbb{F}_q$, $\mathbf{C} \leftarrow \text{Enc}(s)$, $\tilde{\mathbf{C}} \in \mathbb{F}_q^n$ and $A \subseteq [n]$ such that $|A| \geq n - \delta$, $\mathbf{C}|_A = \tilde{\mathbf{C}}|_A$, and $\tilde{\mathbf{C}}|_{\bar{A}} \xleftarrow{\$} \mathbb{F}_q^{n-|A|}$, we have that $\mathbf{C} \in \text{LDec}(\tilde{\mathbf{C}})$ with overwhelming probability.*

Proof. This follows directly from Lemma 2 with $\delta = e = n - 1 - \lceil \sqrt{(k - 1)n} \rceil$.

References

- ABM⁺21. Shashank Agrawal, Saikrishna Badrinarayanan, Payman Mohassel, Pratyay Mukherjee, and Sikhar Patranabis. BETA: Biometric-enabled threshold authentication. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 290–318. Springer, Heidelberg, May 2021.
- ABMR20. Shashank Agrawal, Saikrishna Badrinarayanan, Pratyay Mukherjee, and Peter Rindal. Game-set-MATCH: Using mobile devices for seamless external-facing biometric matching. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1351–1370. ACM Press, November 2020.
- All22. WiFi Alliance. Wpa3 specification version 3.1. <https://www.wi-fi.org/download.php?file=/sites/default/files/private/WPA32022>.
- BBCW21. Manuel Barbosa, Alexandra Boldyreva, Shan Chen, and Bogdan Warinschi. Provable security analysis of FIDO2. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 125–156, Virtual Event, August 2021. Springer, Heidelberg.
- BBR88. Charles H. Bennett, Gilles Brassard, and Jean-Marc Robert. Privacy amplification by public discussion. volume 17, pages 210–229, 1988.
- BCL⁺05. Boaz Barak, Ran Canetti, Yehuda Lindell, Rafael Pass, and Tal Rabin. Secure computation without authentication. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 361–377. Springer, Heidelberg, August 2005.
- BCL22. Jonathan Bootle, Alessandro Chiesa, and Siqi Liu. Zero-knowledge iops with linear-time prover and polylogarithmic-time verifier. In *Proceedings of the 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques*, EUROCRYPT '22, pages 275–304, 2022.
- BDFK12. Jens Bender, Özgür Dagdelen, Marc Fischlin, and Dennis Kügler. The pace|aa protocol for machine readable travel documents, and its security. In Angelos D. Keromytis, editor, *Financial Cryptography and Data Security - 16th International Conference, FC 2012, Kralendijk, Bonaire, February 27-March 2, 2012, Revised Selected Papers*, volume 7397 of *Lecture Notes in Computer Science*, pages 344–358. Springer, 2012.

- BDK⁺05. Xavier Boyen, Yevgeniy Dodis, Jonathan Katz, Rafail Ostrovsky, and Adam Smith. Secure remote authentication using biometric data. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 147–163. Springer, Heidelberg, May 2005.
- BFK09. Jens Bender, Marc Fischlin, and Dennis Kügler. Security analysis of the PACE key-agreement protocol. In Pierangela Samarati, Moti Yung, Fabio Martinelli, and Claudio A. Ardagna, editors, *Information Security, 12th International Conference, ISC 2009, Pisa, Italy, September 7-9, 2009. Proceedings*, volume 5735 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 2009.
- BFK13. Jens Bender, Marc Fischlin, and Dennis Kügler. The pace|ca protocol for machine readable travel documents. In Roderick Bloem and Peter Lipp, editors, *Trusted Systems - 5th International Conference, INTRUST 2013, Graz, Austria, December 4-5, 2013, Proceedings*, volume 8292 of *Lecture Notes in Computer Science*, pages 17–35. Springer, 2013.
- 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.
- 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.
- 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.
- CAA⁺16. Rahul Chatterjee, Anish Athayle, Devdatta Akhawe, Ari Juels, and Thomas Ristenpart. pASSWORD tYPOS and how to correct them securely. In *2016 IEEE Symposium on Security and Privacy*, pages 799–818. IEEE Computer Society Press, May 2016.
- Can01. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
- CCG⁺07. Hao Chen, Ronald Cramer, Shafi Goldwasser, Robbert de Haan, and Vinod Vaikuntanathan. Secure computation from random error correcting codes. In *Proceedings of the 26th Annual International Conference on the Theory and Applications of Cryptographic Techniques, EUROCRYPT’ 07*, pages 291–310, 2007.
- CDBN15. Ronald Cramer, Ivan Damgård, and Jesper Buus Nielsen. *Secure Multiparty Computation and Secret Sharing*. Cambridge University Press, 2015.
- CDD⁺15. Ronald Cramer, Ivan Bjerre Damgård, Nico Döttling, Serge Fehr, and Gabriele Spini. Linear secret sharing schemes from error correcting codes and universal hash functions. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 313–336. Springer, Heidelberg, April 2015.
- CFP⁺16. Ran Canetti, Benjamin Fuller, Omer Paneth, Leonid Reyzin, and Adam D. Smith. Reusable fuzzy extractors for low-entropy distributions. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 117–146. Springer, Heidelberg, May 2016.

- 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.
- CNPR22. Cas Cremers, Moni Naor, Shahar Paz, and Eyal Ronen. CHIP and CRISP: Protecting all parties against compromise through identity-binding PAKEs. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 668–698. Springer, Heidelberg, August 2022.
- CWP⁺17. Rahul Chatterjee, Joanne Woodage, Yuval Pnueli, Anusha Chowdhury, and Thomas Ristenpart. The TypTop system: Personalized typo-tolerant password checking. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 329–346. ACM Press, October / November 2017.
- DHP⁺17. Pierre-Alain Dupont, Julia Hesse, David Pointcheval, Leonid Reyzin, and Sophia Yakubov. Fuzzy password-authenticated key exchange. *Cryptology ePrint Archive*, Paper 2017/1111, 2017. <https://eprint.iacr.org/2017/1111>.
- DHP⁺18. Pierre-Alain Dupont, Julia Hesse, David Pointcheval, Leonid Reyzin, and Sophia Yakubov. Fuzzy password-authenticated key exchange. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 393–424. Springer, Heidelberg, April / May 2018.
- DORS08. Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.
- EHOR20. Andreas Erwig, Julia Hesse, Maximilian Ortl, and Siavash Riahi. Fuzzy asymmetric password-authenticated key exchange. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 761–784. Springer, Heidelberg, December 2020.
- Gur06. Venkatesan Guruswami. Algorithmic results in list decoding. *Found. Trends Theor. Comput. Sci.*, 2(2), 2006.
- HL19. Björn Haase and Benoît Labrique. AuCPace: Efficient verifier-based PAKE protocol tailored for the IIoT. *IACR TCHES*, 2019(2):1–48, 2019. <https://tches.iacr.org/index.php/TCHES/article/view/7384>.
- HS14. Feng Hao and Siamak Fayyaz Shahandashti. The SPEKE protocol revisited. *IACR Cryptol. ePrint Arch.*, page 585, 2014.
- Hv22. Feng Hao and Paul C. van Oorschot. SoK: Password-authenticated key exchange - theory, practice, standardization and real-world lessons. In Yuji Suga, Kouichi Sakurai, Xuhua Ding, and Kazue Sako, editors, *ASIACCS 22*, pages 697–711. ACM Press, May / June 2022.
- Jab96. David P. Jablon. Strong password-only authenticated key exchange. *Comput. Commun. Rev.*, 26(5):5–26, 1996.
- JLHG22. Mingming Jiang, Shengli Liu, Shuai Han, and Dawu Gu. Fuzzy authenticated key exchange with tight security. In Vijayalakshmi Atluri, Roberto Di Pietro, Christian Damsgaard Jensen, and Weizhi Meng, editors, *ESORICS 2022, Part II*, volume 13555 of *LNCS*, pages 337–360. Springer, Heidelberg, September 2022.
- KR07. Michel Kulhandjian and Atri Rudra. Lecture 27: Berlekamp–welch algorithm, 2007.
- Mac01. Philip MacKenzie. On the security of the SPEKE password-authenticated key exchange protocol. *IACR Cryptol. ePrint Arch.*, page 57, 2001.

- McE03. R J McEliece. The guruswami—sudan decoding algorithm for reed—solomon codes. IPN Progress Report 42-153, 2003.
- Nie13. Johan S. R. Nielsen. List decoding of algebraic codes, 2013.
- PC20. Thitikorn Pongmorrakot and Rahul Chatterjee. tpake: Typo-tolerant password-authenticated key exchange. In Lejla Batina, Stjepan Picek, and Mainack Mondal, editors, *Security, Privacy, and Applied Cryptography Engineering - 10th International Conference, SPACE 2020, Kolkata, India, December 17-21, 2020, Proceedings*, volume 12586 of *Lecture Notes in Computer Science*, pages 3–24. Springer, 2020.
- RW04. Renato Renner and Stefan Wolf. The exact price for unconditionally secure asymmetric cryptography. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 109–125. Springer, Heidelberg, May 2004.
- RX23. Lawrence Roy and Jiayu Xu. A universally composable PAKE with zero communication cost - (and why it shouldn't be considered uc-secure). In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *Public-Key Cryptography - PKC 2023 - 26th IACR International Conference on Practice and Theory of Public-Key Cryptography, Atlanta, GA, USA, May 7-10, 2023, Proceedings, Part I*, volume 13940 of *Lecture Notes in Computer Science*, pages 714–743. Springer, 2023.
- The23. The Sage Developers. Sagemath, the Sage Mathematics Software System (Version 9.5), 2023. <https://www.sagemath.org>.
- Wei16. Mor Weiss. Secure computation and probabilistic checking, 2016.
- WHC⁺21. Mei Wang, Kun He, Jing Chen, Zengpeng Li, Wei Zhao, and Ruiying Du. Biometrics-authenticated key exchange for secure messaging. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2618–2631. ACM Press, November 2021.

Appendix

A Full proof

In the following, we discuss in detail why the game hops presented in Section 5 show that no PPT environment \mathcal{Z} has noticeable advantage in distinguishing a real execution of the protocol from Figure 8 from an ideal execution, with $\mathcal{F}_{\text{fPAKE}}^M$ and the simulator from Figs. 9 to 10.

Proof. The first game, denoted \mathbf{G}_0 , is the real execution. We gradually change the experiment until we reach the ideal execution, denoted \mathbf{G}_{11} . In the following, we introduce the games in between and argue directly why each game is indistinguishable from its predecessor. We write $\Pr[\mathbf{G}_i]$ as a shorthand for the probability that the environment outputs 1 in game \mathbf{G}_i , where the probability is taken over the random coins of the environment and all parties.

Game \mathbf{G}_0 : The real execution.

Game \mathbf{G}_1 : Create functionality and simulator. In this step, we introduce a simulator to the experiment. We move the whole execution of the real protocol to the simulator. Concretely, this means that the simulator runs

all parties of the protocol with their real code. To do this the simulator needs the inputs of the parties (which the simulator won't get any more in later games). We therefore introduce dummy parties and an ideal functionality that merely forwards all messages, i.e., a dummy party forwards its input $(\text{NEWSESSION}, \text{sid}, \text{pw}^{\mathcal{P}_i}, \text{role})$ to the functionality and the functionality forwards $(\text{NEWSESSION}, \text{sid}, \mathcal{P}_i, \text{pw}^{\mathcal{P}_i}, \text{role})$ to the simulator (without changing or suppressing $\text{pw}^{\mathcal{P}_i}$). Finally, we equip the functionality with interfaces $(\text{NEWKEY}, \text{sid}, \mathcal{P}_i, \text{sk})$ that allow the simulator to make any dummy parties output the key sk that the simulator wants.

Note that the internal parties inside of the simulator now execute the protocol exactly as in the real world and also have the same output. We just changed where the parties execute their code. This is merely a syntactical change. Thus, both hybrids are identically distributed, i.e.,

$$\Pr[\mathbf{G}_1] = \Pr[\mathbf{G}_0].$$

Game \mathbf{G}_2 : Abort on collisions of H_0 . In this game, we change the simulator such that it aborts the simulation whenever there are two inputs $\mathcal{C}, \mathcal{C}'$ with $\mathcal{C} \neq \mathcal{C}'$ such that $H_0((\text{sid}, \text{sid}_{\mathbb{H}}), \mathcal{C}) = H_0((\text{sid}, \text{sid}_{\mathbb{H}}), \mathcal{C}')$. As the output for each new hash query is chosen uniformly at random, the probability that two values collide after polynomially many queries is negligible:

$$|\Pr[\mathbf{G}_2] - \Pr[\mathbf{G}_1]| \leq \text{negl}(\lambda).$$

Game \mathbf{G}_3 : Exchanging real iPAKE messages for simulated ones. In this game, we only change the simulator. We introduce a variable MitM , which is set to 1, if there is a man-in-the-middle attack and to 0 otherwise. More precisely, on a message $(\text{INIT}, \text{sid}, \mathcal{P}_i, \mathbb{H}, \text{sid}_{\mathbb{H}})$ from \mathcal{Z} to the simulated \mathcal{F}_{SA} , \mathcal{S} sets $\text{MitM} = 0$, if $\mathbb{H} = \{\mathcal{P}_0, \mathcal{P}_1\}$ and $\text{MitM} = 1$, otherwise. From now on, the simulator runs one internal instance of $\mathcal{F}_{\text{iPAKE}}$, presented in Figure 12, for each distinct session of π . Whenever a party \mathcal{P}_i in our protocol starts a new session of π on a new sid $(\text{sid}, \text{sid}_{\mathbb{H}}, j)$ with $\text{pw}_j^{\mathcal{P}_i}$, \mathcal{S} creates an instance of $\mathcal{F}_{\text{iPAKE}}$ with the same sid and gives it the input $(\text{NEWSESSION}, (\text{sid}, \text{sid}_{\mathbb{H}}, j), \text{pw}_j, \text{role})$ in the name of \mathcal{P}_i . \mathcal{S} also invokes one iPAKE simulator $\mathcal{S}_{\text{iPAKE}}$ for each instance of $\mathcal{F}_{\text{iPAKE}}$, which is guaranteed to exist, because π UC-realizes $\mathcal{F}_{\text{iPAKE}}$. The situation is depicted in Fig. 11. If the environment is not doing a MitM attack on the communication, i.e., $\text{sid}_{\mathbb{H}}$ is the same for both parties, then there will be n instances of $\mathcal{F}_{\text{iPAKE}}$. If the environment is doing a MitM attack, i.e., $\text{sid}_{\mathbb{H}}$ is different for \mathcal{P}_0 and \mathcal{P}_1 , then there will be $2n$ instances of $\mathcal{F}_{\text{iPAKE}}$, because then each party has their own n $\mathcal{F}_{\text{iPAKE}}$ instances that run on different session-ids than the other party's instances.

\mathcal{S} can simply execute the code of $\mathcal{F}_{\text{iPAKE}}$, as in this game \mathcal{S} still knows the parties' passwords and, thus, all inputs to the $\mathcal{F}_{\text{iPAKE}}$ instances. When an instance of $\mathcal{F}_{\text{iPAKE}}$ outputs $((\text{sid}, \text{sid}_{\mathbb{H}}, j), \text{sk})$ to \mathcal{P}_i as response to a NEWKEY query of the corresponding $\mathcal{S}_{\text{iPAKE}}$, \mathcal{S} gives this to \mathcal{P}_i as if coming from π_j . Whenever all n $\mathcal{S}_{\text{iPAKE}}$ belonging to \mathcal{P}_i produce a message Z_j^* for \mathcal{Z} ,

simulating that \mathcal{P}_i sent this message, \mathcal{S} gives (Z_1^*, \dots, Z_n^*) to \mathcal{F}_{SA} as message from \mathcal{P}_i . This is the same as \mathcal{P}_i would have done when all n π_j had produced a message Z_j^* . Finally, whenever \mathcal{F}_{SA} delivers a message (Z_1^*, \dots, Z_n^*) to \mathcal{P}_i (upon instruction from \mathcal{Z}), \mathcal{S} gives the instruction to deliver Z_j^* to $\mathcal{S}_{\text{iPAKE}}^j$ in the name of the environment.

To show that $\mathbf{G}_3 \stackrel{c}{\approx} \mathbf{G}_2$ we use a hybrid argument. In the j -th step of the hybrid argument, we exchange the j -th instance π_j of π with one instance of the ideal execution $\mathcal{F}_{\text{iPAKE}}^j$ and $\mathcal{S}_{\text{iPAKE}}^j$. After at most $2n$ steps (or n steps if there is no man-in-the-middle attack) we have exchanged all instances of π and reached \mathbf{G}_3 .

To see that an individual hybrid step is valid, observe that \mathcal{S} perfectly emulates the internal $\mathcal{F}_{\text{iPAKE}}$ instances. Furthermore, because π UC-realizes $\mathcal{F}_{\text{iPAKE}}$, the outputs of the simulators $\mathcal{S}_{\text{iPAKE}}$ are indistinguishable from the messages of π .

In a bit more detail, if there is an index l , such that an environment \mathcal{Z}_l can distinguish between hybrids l and $l + 1$, then we can use \mathcal{Z}_l to build an environment \mathcal{Z}^* that can distinguish between an execution of π and an execution of the ideal protocol with $\mathcal{F}_{\text{iPAKE}}$. Let \mathcal{S}_l be the simulator in \mathbf{G}_3 at hybrid step l . The environment \mathcal{Z}^* runs \mathcal{Z}_l and \mathcal{S}_l and whenever a party \mathcal{P}_i starts a session of π_{l+1} , \mathcal{Z}^* gives the input to their test-session, which is either the real version of π , or the ideal version with $\mathcal{F}_{\text{iPAKE}}$. If \mathcal{Z}_l guesses that it is in hybrid l , \mathcal{Z}^* outputs that it is in the real world, and if \mathcal{Z}_l guesses that it is in hybrid $l + 1$ then \mathcal{Z}^* outputs that it is in the ideal world. \mathcal{Z}^* has the same advantage as \mathcal{Z}_l .

Thus, we have that

$$|\Pr[\mathbf{G}_3] - \Pr[\mathbf{G}_2]| \leq 2n \mathbf{Dist}_{\mathcal{Z}^*}^{\pi, \{\mathcal{F}_{\text{iPAKE}}, \mathcal{S}_{\text{iPAKE}}\}}(\lambda),$$

where $\mathbf{Dist}_{\mathcal{Z}^*}^{\pi, \{\mathcal{F}_{\text{iPAKE}}, \mathcal{S}_{\text{iPAKE}}\}}(\lambda)$ denotes the distinguishing advantage that a PPT environment \mathcal{Z}^* has in distinguishing a real-world execution of π from an ideal-world execution with $\mathcal{F}_{\text{iPAKE}}$ and $\mathcal{S}_{\text{iPAKE}}$.

Game \mathbf{G}_4 : Functionality makes records of running sessions, marks them, and reacts accordingly. In this game, we only change the ideal functionality. Now, whenever the functionality receives input $(\text{NEWSESSION}, \text{sid}, \text{pw}^{\mathcal{P}_i}, \text{role})$ from the dummy party \mathcal{P}_i it records $\langle \mathcal{P}_i, \text{sid}, \text{role}, \text{pw}^{\mathcal{P}_i} \rangle$ if this was the first NEWSESSION query or if it was the second NEWSESSION query and there already is a record $\langle \mathcal{P}_{1-i}, \text{sid}, \text{pw}^{\mathcal{P}_{1-i}} \rangle$. It marks the record as **fresh**. This means that the ideal functionality now has the full NEWSESSION interface of $\mathcal{F}_{\text{iPAKE}}$ with the addition that it still forwards the passwords to the simulator. Next, we add the TESTPWD interface as in the full $\mathcal{F}_{\text{iPAKE}}^M$ functionality, thus, records can also be marked **interrupted** or **compromised**. Furthermore, we change the functionality's behavior on $(\text{NEWKEY}, \text{sid}, \mathcal{P}_i, \text{sk})$ queries such that it now checks if a record $\langle \mathcal{P}_i, \text{sid}, \text{role}, \text{pw}^{\mathcal{P}_i} \rangle$ exists. If so, it outputs a uniformly random key when the corresponding record is marked as **interrupted**. Also, it outputs the adversarially provided key sk when the

corresponding record is marked as **compromised** or the other party is corrupted as in $\mathcal{F}_{\text{fPAKE}}$. Note, however, that the functionality in case of **fresh** records still uses the key provided by the simulator. This will change in a later game. Regardless of how the record was marked before, the functionality now marks the record as **completed**.

Note that the records and markings are only internal and their only effect on the outside is through the **NEWKEY** interface. In the **NEWKEY** interface the functionality still outputs the key chosen by the simulator, except if the corresponding record is **interrupted**. However, so far our simulator never calls $\mathcal{F}_{\text{fPAKE}}$'s **TESTPWD** interface and, therefore, currently a record can never be **interrupted**. Thus, the changes to the functionality are only syntactical and we have

$$\Pr[\mathbf{G}_4] = \Pr[\mathbf{G}_3].$$

Game \mathbf{G}_5 : Simulator handles TestPwd queries of $\mathcal{F}_{\text{fPAKE}}^M$ for honest and undisturbed session. In this game, we only modify the simulator. In the case of an honest session, i.e. when there is no MitM attack and no party is corrupted, \mathcal{S} does not use the passwords of the parties to simulate **TESTPWD** of $\mathcal{S}_{\text{iPAKE}}$ queries anymore. Whenever \mathcal{S} receives a **TESTPWD** query from $\mathcal{S}_{\text{iPAKE}}$ for a session with two honest parties and where $\text{MitM} = 0$, \mathcal{S} just ignores the query.

This means, for honest sessions, \mathcal{S} can simulate the **NEWSESSION** and the **TESTPWD** interfaces of $\mathcal{F}_{\text{iPAKE}}$ to $\mathcal{S}_{\text{iPAKE}}$ without using the password $\text{pw}^{\mathcal{P}_i}, \text{pw}^{\mathcal{P}_{1-i}}$ of the parties. However, \mathcal{S} still uses $\text{pw}^{\mathcal{P}_i}$ and $\text{pw}^{\mathcal{P}_{1-i}}$ to simulate $\mathcal{F}_{\text{iPAKE}}$'s reaction on **NEWKEY** queries from $\mathcal{S}_{\text{iPAKE}}$.

Below we argue why this does not change the distribution of the game. The only way that ignoring **TESTPWD** queries from the $\mathcal{S}_{\text{iPAKE}}$ can change the distribution of the game is when a $\mathcal{S}_{\text{iPAKE}}$ asks a **TESTPWD** query. However, we only changed the simulator for the case of an honest session and in that case, both parties receive the same $\text{sid}_{\mathbb{H}'} = \text{sid}_{\mathbb{H}}$ from \mathcal{F}_{SA} and \mathcal{F}_{SA} behaves like the normal authenticated channel functionality. Thus, the adversary can either let the session proceed without inserting or changing messages, or it can completely prevent the iPAKE sessions from finishing. Intuitively, in an honest and undisturbed iPAKE session, $\mathcal{S}_{\text{iPAKE}}$ cannot ask a **TESTPWD** query, because $\mathcal{S}_{\text{iPAKE}}$ does not know the password and, thus, this would most likely interrupt the session. Then both parties would get distinct keys, even when they used the same passwords. This should not happen in an honest session. More formally we argue as follows. As we assumed that the iPAKE protocol π is *correct* (see Definition 1) it follows that with overwhelming probability the $\mathcal{S}_{\text{iPAKE}}$ simulators will not send a **TESTPWD** query to \mathcal{S} as both parties are honest and the adversary only eavesdrops on the session. (In the terminology of [RX23], $\mathcal{S}_{\text{iPAKE}}$ is a *perfectly reasonable simulator*. This follows from the fact that π is correct and UC-realizes $\mathcal{F}_{\text{iPAKE}}$, see [RX23, Thm. 1].) Thus, we get

$$|\Pr[\mathbf{G}_5] - \Pr[\mathbf{G}_4]| \leq \text{negl}(\lambda).$$

Game G_6 : Simulator handles attacks on honest senders. Previously, we argued that \mathcal{S} does not need to send a TESTPWD query to $\mathcal{F}_{\text{iPAKE}}^M$ if there is no man-in-the-middle attack and both parties are honest to deal with TESTPWD of $\mathcal{S}_{\text{iPAKE}}$. In this game, we change the simulator's behavior when there is a man-in-the-middle attack against an honest sender or when the receiver is corrupted.

The high-level idea in this game is that \mathcal{S} collects the TESTPWD queries from the iPAKE simulators, concatenates them, and then asks one TESTPWD query to $\mathcal{F}_{\text{iPAKE}}^M$. If the guess was close enough, \mathcal{S} learns the position of the correct password characters. Then \mathcal{S} marks the record of the (internal) $\mathcal{F}_{\text{iPAKE}}$ instance as **compromised** if the corresponding character was correct, or as **interrupted** if that character was wrong. If the guess of \mathcal{S} was too far away, \mathcal{S} still uses the password of the sender to check which of the characters were correct and then sets the $\mathcal{F}_{\text{iPAKE}}$ entry accordingly.

Next, we describe the changes to \mathcal{S} in more detail. When \mathcal{S} receives a (NEWSESSION, sid, \mathcal{P}_i , $\text{pw}^{\mathcal{P}_i}$, role) message from $\mathcal{F}_{\text{iPAKE}}$ then \mathcal{S} records $\langle \mathcal{P}_i, \text{sid}, \text{role} \rangle$. In the following, we only change \mathcal{S} 's behavior if \mathcal{S} recorded $\langle \mathcal{P}_i, \text{sid}, \text{sender} \rangle$, i.e., if \mathcal{P}_i is a sender. On getting a (TESTPWD, sid, \mathcal{P}_i , pw'_j) query from $\mathcal{S}_{\text{iPAKE}}^j$, \mathcal{S} stores the password guess and does not do anything else, because $\mathcal{F}_{\text{iPAKE}}^j$ would also not answer to the query. (Remember that $\mathcal{F}_{\text{iPAKE}}$ models implicit-only PAKE.) When getting a (NEWKEY, sid, \mathcal{P}_i , sk) query from $\mathcal{S}_{\text{iPAKE}}^j$, \mathcal{S} stores sk. As soon as \mathcal{S} receives a NEWKEY query from all n $\mathcal{S}_{\text{iPAKE}}$, \mathcal{S} does the following: Let m be the number of iPAKE TESTPWD queries for \mathcal{P}_i that \mathcal{S} received, so $m \leq n$. We have $m < n$ when there are some simulators $\mathcal{S}_{\text{iPAKE}}^j$ that give input NEWKEY to \mathcal{S} without first giving input TESTPWD to \mathcal{S} . Let δ and γ be as in Figure 4.

- If $m = n$, then \mathcal{S} concatenates all the iPAKE password guesses and asks a TESTPWD query to $\mathcal{F}_{\text{iPAKE}}^M$. Concretely, after receiving (NEWKEY, sid, \mathcal{P}_i , sk_j) from $\mathcal{S}_{\text{iPAKE}}^j$, for all $j \in [n]$, \mathcal{S} retrieves the stored pw'_j and sends (TESTPWD, sid, \mathcal{P}_i , $\text{pw}'_1 \parallel \dots \parallel \text{pw}'_n$) to $\mathcal{F}_{\text{iPAKE}}^M$. If the guess was correct, i.e., at least $n - \delta$ correct password characters, then \mathcal{S} continues the simulation of \mathcal{P}_i with the sk_j from the NEWKEY queries to $\mathcal{F}_{\text{iPAKE}}^j$, where the corresponding password character pw_j was correct. For the incorrect pw_j , \mathcal{S} uses random iPAKE keys, which is the same as $\mathcal{F}_{\text{iPAKE}}$ would do. If the guess was wrong, but there were at least $n - \gamma$ correct password characters, \mathcal{S} learns the position of the correct characters from $\mathcal{F}_{\text{iPAKE}}^M$ through the leakage function. \mathcal{S} continues the simulation as above, i.e. setting the iPAKE keys to sk_j if the j -th password character was correct and to a random value if it was incorrect. If the guess was wrong and \mathcal{S} does not get the position of the correct characters, i.e., less than $n - \gamma$ correct characters, then \mathcal{S} uses the input password of the sender (which in this game it still gets from $\mathcal{F}_{\text{iPAKE}}$) to check the $j \in [n]$ for which the TESTPWD queries by $\mathcal{S}_{\text{iPAKE}}^j$ to $\mathcal{F}_{\text{iPAKE}}$ were correct. For the wrong positions, \mathcal{S} chooses a uniformly random key, and for the correct positions j , the simulator uses the sk_j provided by $\mathcal{S}_{\text{iPAKE}}$.

- If $m < n - \gamma$, then \mathcal{S} uses the input password $\text{pw}^{\mathcal{P}_i}$ of the honest sender as above to check which TESTPWD queries to $\mathcal{F}_{\text{iPAKE}}$ were correct. For the wrong positions \mathcal{S} chooses a uniformly random key and for the correct positions j the simulator uses the sk_j provided by $\mathcal{S}_{\text{iPAKE}}$. Additionally \mathcal{S} sends a TESTPWD query for \mathcal{P}_i with $\text{pw} = \perp$ to $\mathcal{F}_{\text{iPAKE}}^M$ in order to make sure that the record is **interrupted**.
- If $n - \gamma \leq m < n$, \mathcal{S} arbitrarily chooses the missing password characters and asks a TESTPWD query to $\mathcal{F}_{\text{iPAKE}}^M$. If the guess was wrong and \mathcal{S} does not get leakage from $\mathcal{F}_{\text{iPAKE}}^M$, then \mathcal{S} acts as above and uses the input password $\text{pw}^{\mathcal{P}_i}$ of the honest sender to check which TESTPWD queries from $\mathcal{S}_{\text{iPAKE}}$ to $\mathcal{F}_{\text{iPAKE}}$ were correct. For the wrong positions \mathcal{S} chooses a uniformly random key and for the correct positions j the simulator uses the sk_j provided by $\mathcal{S}_{\text{iPAKE}}$.

Now, consider the case that \mathcal{S} gets leakage from $\mathcal{F}_{\text{iPAKE}}^M$ (either because the guess was correct, or at most γ characters were wrong). Then \mathcal{S} needs to check if this was only because of the missing password characters that \mathcal{S} arbitrarily guessed, or because \mathcal{Z} did guess enough correct characters. Because \mathcal{S} learned the position of the correct characters from $\mathcal{F}_{\text{iPAKE}}^M$, \mathcal{S} can check whether \mathcal{Z} did indeed guess enough correct characters. If that is the case, then \mathcal{S} continues the simulations with the sk_j from the NEWKEY queries to $\mathcal{F}_{\text{iPAKE}}$ for the correct positions and with random keys for the incorrect positions. Otherwise (if \mathcal{Z} did not guess enough correct characters), \mathcal{S} uses the sender’s password to determine the position of the correct characters and then sets the iPAKE keys as in a previous case. The only issue is the case where \mathcal{Z} guessed less than $n - \gamma$ correct characters, but \mathcal{S} ’s guesses of the remaining characters were responsible for exceeding $n - \gamma$. In this case, $\mathcal{F}_{\text{iPAKE}}^M$ ’s record is compromised when it should be interrupted. \mathcal{S} recognizes this case, through learning the positions of the correct password characters from the leakage function. Thus, \mathcal{S} can remedy the issue by setting a random key in the NEWKEY query to $\mathcal{F}_{\text{iPAKE}}^M$, thereby imitating a **interrupted** record. Note that \mathcal{S} can do this as the outputs that the internal $\mathcal{F}_{\text{iPAKE}}^j$ functionalities give to parties are not reported to \mathcal{Z} directly but only to \mathcal{S} .

In all of the above cases, if the record is **interrupted** (\mathcal{S} recognizes this by getting “wrong guess” and no leakage as an answer to the TESTPWD query), \mathcal{S} sets $\text{sk} = \perp$ in the NEWKEY query to $\mathcal{F}_{\text{iPAKE}}^M$, as $\mathcal{F}_{\text{iPAKE}}^M$ would ignore sk anyways. If the record is **compromised**, \mathcal{S} still sets sk to the key that the simulated \mathcal{P}_i would output.

We argue in the following why the difference between \mathbf{G}_6 and \mathbf{G}_5 is at most negligible. When there is a man-in-the-middle attack happening or the receiver is corrupted \mathcal{S} has n iPAKE simulators for the \mathcal{P}_i session (and another n iPAKE simulators for the \mathcal{P}_{1-i} session if \mathcal{P}_{1-i} is not corrupted). This is because the adversary conducts a man-in-the-middle attack (\mathcal{S} knows this) and therefore \mathcal{F}_{SA} (played by \mathcal{S}) provides different session identifier $\text{sid}_{\mathbb{H}'} \neq \text{sid}_{\mathbb{H}}$ to \mathcal{P}_i and \mathcal{P}_{i-1} . Therefore \mathcal{P}_i runs its own instance of $\mathcal{F}_{\text{iPAKE}}$ with $\text{sid}_{\mathbb{H}}$ (actually n of them). This instance is in particular independent

of \mathcal{P}_{1-i} 's instance $\mathcal{F}'_{\text{iPAKE}}$ that runs with $\text{sid}_{\mathbb{H}'} \neq \text{sid}_{\mathbb{H}}$. Therefore, $\mathcal{F}_{\text{iPAKE}}$ is not in the case where two honest parties execute the protocol. This means that either the adversary compromises the session and makes \mathcal{P}_i output an adversarially chosen sk' , or the adversary makes \mathcal{P}_i output a random key, either by interrupting the session with a wrong password guess in a TESTPWD to $\mathcal{F}_{\text{iPAKE}}$ query or by sending NEWKEY without a TESTPWD query to $\mathcal{F}_{\text{iPAKE}}$. \mathcal{S} would ignore NEWKEY queries for \mathcal{P}_{1-i} , because $\mathcal{F}_{\text{iPAKE}}$ would also ignore them, since \mathcal{P}_{1-i} never sent a NEWSESSION query (for that specific session-id). \mathcal{S} would also ignore TESTPWD queries from $\mathcal{S}_{\text{iPAKE}}$ for \mathcal{P}_{1-i} for the same reason. Note that this is what prevents the attack from Section 3.2.

We distinguish two cases depending on how many characters of the tested password were correct. Let us first consider the case where \mathcal{Z} guessed less than $n - \gamma = t$ password characters correctly. This includes the case that \mathcal{S} 's guess of the additional characters made the guess appear correct to $\mathcal{F}_{\text{iPAKE}}^M$. In this case, \mathcal{S} still uses the honest sender's password and sets the output of the $\mathcal{F}_{\text{iPAKE}}$ instances the same way as in \mathbf{G}_5 . The only difference to \mathbf{G}_5 is that the record $(\mathcal{P}_i, \text{pw}^{\mathcal{P}_i})$ of $\mathcal{F}_{\text{iPAKE}}^M$ is now interrupted and, thus, the output of the sender is a uniformly random value instead of $H_1((\text{sid}, \text{sid}_{\mathbb{H}}), s_0)$.

Let us explain why this change is not detectable by \mathcal{Z} . Because H_1 is a random oracle, $H_1((\text{sid}, \text{sid}_{\mathbb{H}}), s_0)$ is uniformly random for anybody that did not query H_1 on $((\text{sid}, \text{sid}_{\mathbb{H}}), s_0)$. The chance of guessing s_0 is negligible. The only other values that could contain information about s_0 are the hint $h = H_0((\text{sid}, \text{sid}_{\mathbb{H}}), \mathbf{C})$ and $\mathbf{E} = \mathbf{C} \oplus \mathbf{K}$. Because \mathcal{Z} guessed at most $n - \gamma - 1 = k - 1$ password characters correctly, \mathcal{Z} can recover at most $k - 1$ coordinates of \mathbf{C} . By the $k - 1$ query uniformity of the code, these (at most) $k - 1$ undisturbed coordinates appear uniformly random and, thus, do not reveal any information about s_0 . The hint h is also uniformly random unless one has queried $((\text{sid}, \text{sid}_{\mathbb{H}}), \mathbf{C})$ to H_0 . Let us now bound the probability that \mathcal{Z} queries \mathbf{C} to H_0 . Since $\mathbf{E} = \mathbf{C} \oplus \mathbf{K}$ and \mathbf{K} is uniformly random (and unknown to \mathcal{Z}) in at least $\gamma + 1$ positions, \mathcal{Z} would need to correctly guess the remaining coordinates in $\mathbb{F}_q^{\gamma+1}$. The probability that \mathcal{Z} can do so in polynomial time is negligible.

Let us now consider the case that \mathcal{Z} guessed at least $n - \gamma$ password characters correctly. In this case, \mathcal{S} learned the position of the correct characters via the leakage function and, thus, simulates the $\mathcal{F}_{\text{iPAKE}}$ identically to \mathbf{G}_5 . Therefore, the session key that \mathcal{S} sends to $\mathcal{F}_{\text{iPAKE}}^M$ via the NEWKEY interface is identically distributed to \mathbf{G}_5 . The only difference is that in \mathbf{G}_6 the corresponding record of $\mathcal{F}_{\text{iPAKE}}^M$ is **compromised** instead of **fresh**. However, until now, $\mathcal{F}_{\text{iPAKE}}^M$ outputs \mathcal{S} 's key in both cases. Therefore, in the case that \mathcal{Z} guessed at least $n - \gamma$ password characters correctly, \mathbf{G}_5 and \mathbf{G}_6 are identically distributed. In total we get

$$|\Pr[\mathbf{G}_6] - \Pr[\mathbf{G}_5]| \leq \text{negl}(\lambda).$$

Game \mathbf{G}_7 : Use random iPAKE keys in case of attack on honest sender.

Again we only change the simulator in the case of a MitM attack, or if

the receiver is corrupted. If \mathcal{Z} 's password guess is correct in at most $k - 1 = n - \gamma - 1$ characters, then \mathcal{S} uses random iPAKE keys for all $\mathcal{F}_{\text{iPAKE}}$ instances of the sender. This means that the view of \mathcal{Z} on $\mathbf{C}' := \mathbf{E} \oplus \mathbf{K}'$ is different in \mathbf{G}_6 and \mathbf{G}_7 . Let $t \leq k - 1$ be the number of correct password characters. In \mathbf{G}_6 the receiver still obtained the original code word on t entries, whereas in this game all entries are uniformly random. However, as $t \leq k - 1$, the $(k - 1)$ -query uniformity of the code guarantees that \mathbf{C}' was already uniformly random in \mathbf{G}_6 . Therefore, \mathbf{G}_7 is distributed identically to \mathbf{G}_6 .

Note that in the case of a MitM attack, or if the receiver is corrupted, the simulator does not use the password of the honest sender anymore and instead only uses the information that it received from $\mathcal{F}_{\text{iPAKE}}^M$ as an answer to a TESTPWD query. We get

$$\Pr[\mathbf{G}_7] = \Pr[\mathbf{G}_6].$$

Game \mathbf{G}_8 : Simulator handles attacks on honest receivers. In this game, we change the simulator's behavior when there is a man-in-the-middle attack against an honest receiver happening, or when the sender is corrupted. More precisely, we only change \mathcal{S} 's behavior if \mathcal{S} recorded $\langle \mathcal{P}_i, \text{sid}, \text{sender}, \text{MitM} = 1 \rangle$ or if \mathcal{S} recorded $\langle \mathcal{P}_i, \text{sid}, \text{sender}, \text{MitM} = 0 \rangle$ but \mathcal{P}_{1-i} is corrupted. Like in \mathbf{G}_6 , \mathcal{S} now records TESTPWD queries from $\mathcal{S}_{\text{iPAKE}}^j$ for $j \in [n]$. When it received all n NEWKEY queries from $\mathcal{S}_{\text{iPAKE}}^j$ for \mathcal{P}_i , it distinguishes the following cases:

- If \mathcal{S} recorded $m \geq n - \gamma$ TESTPWD queries with a guess pw_j of $\mathcal{S}_{\text{iPAKE}}^j$ then \mathcal{S} randomly chooses the missing $n - m$ password-entries and submits a TESTPWD query to $\mathcal{F}_{\text{iPAKE}}$. If \mathcal{S} receives leakage from $\mathcal{F}_{\text{iPAKE}}$ as response \mathcal{S} behaves exactly as in \mathbf{G}_6 . This means, \mathcal{S} sets K_j to the value sk_j that was chosen by $\mathcal{S}_{\text{iPAKE}}^j$ for the j where the corresponding guess of pw_j was correct and sets K_j to random values for those $j \in [n]$ where the guess of pw_j was wrong or there was no TESTPWD query of $\mathcal{S}_{\text{iPAKE}}^j$ for that j . In particular, if $\mathcal{F}_{\text{iPAKE}}$ answers that the guess was correct, then \mathcal{S} must check if this was due to its own added password entries or because of the guesses from the $\mathcal{S}_{\text{iPAKE}}^j$ and adjust its NEWKEY query to $\mathcal{F}_{\text{iPAKE}}$ accordingly. However, in the case where \mathcal{S} does only get the answer “wrong guess” and no leakage from $\mathcal{F}_{\text{iPAKE}}$, \mathcal{S} does not use LDec. Instead, \mathcal{S} simply sends a NEWKEY query to $\mathcal{F}_{\text{iPAKE}}^M$ with $\text{sk} = \perp$, as the record is **interrupted** and $\mathcal{F}_{\text{iPAKE}}^M$ would ignore sk anyways.
- If there were less than $n - \gamma$ TESTPWD queries of $\mathcal{S}_{\text{iPAKE}}^j$ recorded, then \mathcal{S} asks a TESTPWD query in order to interrupt the record of $\mathcal{F}_{\text{iPAKE}}^M$. Then \mathcal{S} acts as above when it receives no leakage. This means, it does not decode \mathbf{E} but instead sends a NEWKEY query to $\mathcal{F}_{\text{iPAKE}}^M$ with $\text{sk} = \perp$.

Like in \mathbf{G}_6 , when \mathcal{S} receives leakage from $\mathcal{F}_{\text{iPAKE}}$ the distribution of the experiment does not change as \mathcal{S} uses the same session keys sk_j from $\mathcal{S}_{\text{iPAKE}}^j$ as in \mathbf{G}_6 for the correct $j \in [n]$ and random keys K_j for the other positions. We stress again here that we leverage the fact that communication runs over

\mathcal{F}_{SA} for that argument. As $\text{MitM} = 1$ (or the sender is corrupted) the $\mathcal{F}_{\text{iPAKE}}$ instances of the receiver will never successfully exchange keys, as they run in a different authentication set, i.e., with a different session-id, than the iPAKE of the sender. This is what prevents the attack from Section 3.2.

However, in the case where \mathcal{S} does not receive leakage or less than $n - \gamma$ of the $\mathcal{S}_{\text{iPAKE}}^j$ sent a TESTPWD query, the behavior of \mathcal{S} changed. In \mathbf{G}_7 , as we either have $\text{MitM} = 1$ or a corrupt sender, \mathcal{Z} can instruct the adversary to send a modified message (\mathbf{E}', h') to the honest receiver. But in \mathbf{G}_7 the receiver would at least have $\gamma + 1$ indices $j \in [n]$ where the corresponding iPAKE key K_j is uniformly random. That is because either $\mathcal{S}_{\text{iPAKE}}^j$ did send NEWKEY without sending TESTPWD or because the provided password guess pw_j from the TESTPWD query of $\mathcal{S}_{\text{iPAKE}}^j$ was wrong. But then, the receiver computes $\text{LDec}(\mathbf{E}' \oplus \mathbf{K}')$, where $\mathbf{E}' \oplus \mathbf{K}'$ is uniformly random in at least $\gamma + 1$ positions. Now we need to argue that also in \mathbf{G}_7 the honest receiver would have output an independent, uniformly random key. First consider the case that \mathcal{Z} has not queried the random oracle H_0 on a vector $v \in \mathbb{F}_q^n$ such that $H_0((\text{sid}, \text{sid}_{\mathbb{H}}), v) = h'$. Then the probability that a candidate codeword \mathbf{C}_j output by LDec matches the hint h' is negligible and the honest receiver will output $u \xleftarrow{\$} \{0, 1\}^\lambda$. Now consider the case that \mathcal{Z} did query $v \in \mathbb{F}_q^n$ such that $H_0((\text{sid}, \text{sid}_{\mathbb{H}}), v) = h'$. Then, v is uniquely determined by h' because since \mathbf{G}_2 the simulator aborts on collisions of H_0 . Here we can apply Lemma 4 by setting $\mathbf{C} = v$, $\tilde{\mathbf{C}} = \mathbf{C}$ and A to be the indices of the correct password guesses of the $\mathcal{S}_{\text{iPAKE}}^j$ and get that $d(v, \mathbf{C}') \geq \gamma + 1$, which also means that $d(v, \mathbf{C}) \geq \delta + 1$. Since LDec does not output candidates with a distance greater than δ to \mathbf{C} , LDec will not output v and, thus, none of the candidates will match the hint. Therefore, the honest sender will output $u \xleftarrow{\$} \{0, 1\}^\lambda$.

Thus, in \mathbf{G}_7 in both cases, the honest sender outputs a fresh uniformly random key with overwhelming probability which is the same as in \mathbf{G}_8 . Therefore,

$$|\Pr[\mathbf{G}_8] - \Pr[\mathbf{G}_7]| \leq \text{negl}(\lambda).$$

Game \mathbf{G}_9 : Simulator does not give passwords to $\mathcal{F}_{\text{iPAKE}}$ instances. In this game, we change the simulator. Whenever the simulator receives a message $(\text{NEWSESSION}, \text{sid}, \mathcal{P}_i, \text{pw}^{\mathcal{P}_i}, \text{role})$ from $\mathcal{F}_{\text{iPAKE}}$ for an honest party \mathcal{P}_i it does not use $\text{pw}^{\mathcal{P}_i}$ as input for $\mathcal{F}_{\text{iPAKE}}$. Concretely, instead of internally running $\mathcal{F}_{\text{iPAKE}}^j$ on input $(\text{NEWSESSION}, \text{sid}, \mathcal{P}_i, \text{pw}_j^{\mathcal{P}_i}, \text{role})$ the simulator forwards $(\text{NEWSESSION}, \text{sid}, \mathcal{P}_i, \text{role})$ (without the j -th bit of $\text{pw}^{\mathcal{P}_i}$) to $\mathcal{S}_{\text{iPAKE}}^j$. Further, \mathcal{S} checks if this is the first NEWSESSION query or this is the second NEWSESSION query and if there is a record $\langle \mathcal{F}_{\text{iPAKE}}, \mathcal{P}_{1-i}, (\text{sid}, \text{sid}_{\mathbb{H}}, j) \rangle$. If so, \mathcal{S} creates a record $\langle \mathcal{F}_{\text{iPAKE}}, \mathcal{P}_i, (\text{sid}, \text{sid}_{\mathbb{H}}, j) \rangle$ and marks it **fresh**. On a NEWKEY query, \mathcal{S} still behaves exactly as $\mathcal{F}_{\text{iPAKE}}$ by using $\text{pw}^{\mathcal{P}_i}$ to check if a matching output has to be given to the parties.

For every iPAKE instance, the view on the execution did not change, as \mathcal{S} behaves as $\mathcal{F}_{\text{iPAKE}}$ would.

$$\Pr[\mathbf{G}_9] = \Pr[\mathbf{G}_8].$$

Game \mathbf{G}_{10} : Simulator makes honest and undisturbed parties output a

key. In this game, we change the simulator and the functionality. First, we change $\mathcal{F}_{\text{fPAKE}}^M$'s behavior on receiving a query (NEWKEY, sid, \mathcal{P}_i , sk) from \mathcal{S} . Now, $\mathcal{F}_{\text{fPAKE}}^M$ treats these queries exactly like in Fig. 3. That means concretely, if the record $(\mathcal{P}_i, \text{pw}^{\mathcal{P}_i})$ is **fresh** and both parties are honest and run on input passwords with $d(\text{pw}^{\mathcal{P}_i}, \text{pw}^{\mathcal{P}_{1-i}}) \leq \delta$ then both parties receive the same random key.

We also change the simulator. When the simulator receives a message (NEWSESSION, sid, \mathcal{P}_i , $\text{pw}^{\mathcal{P}_i}$, role) from $\mathcal{F}_{\text{fPAKE}}$ for an honest party \mathcal{P}_i where there is no man-in-the-middle attack \mathcal{S} discards $\text{pw}^{\mathcal{P}_i}$ and stores \perp instead of $\text{pw}^{\mathcal{P}_i}$ in all records of the form $(\mathcal{P}_i, \text{sid}, \text{role}, \text{pw}^{\mathcal{P}_i})$. We also change how \mathcal{S} reacts on receiving all n NEWKEY queries from $\mathcal{S}_{\text{fPAKE}}^j$ for $j \in [n]$.

- If the role of \mathcal{P}_i is role = receiver then \mathcal{S} marks its record $(\mathcal{P}_i, \text{sid}, \text{role}, \text{MitM} = 0)$ as **waiting** but produces no output for \mathcal{P}_i . Instead, on a message (DELIVER, sid, \mathcal{P}_i , \mathcal{P}_{1-i} , (\mathbf{E}, h)) the simulator sends a message (NEWKEY, sid, \mathcal{P}_i , \perp) to $\mathcal{F}_{\text{fPAKE}}$ if (\mathbf{E}, h) is not tampered with.
- If the role of \mathcal{P}_i is role = sender then \mathcal{S} marks its record $(\mathcal{P}_i, \text{sid}, \text{role}, \text{MitM} = 0)$ as **completed** and chooses K_1, \dots, K_n uniformly at random. The simulator uses these keys to produce (\mathbf{E}, h) as an honest party would. \mathcal{S} sends (SEND, sid, \mathcal{P}_i , \mathcal{P}_{1-i} , (\mathbf{E}, h)) to \mathcal{Z} formatted as being sent from \mathcal{P}_i to \mathcal{P}_{1-i} via \mathcal{F}_{SA} . Then, \mathcal{S} sends (NEWKEY, sid, \mathcal{P}_i , \perp) to $\mathcal{F}_{\text{fPAKE}}$.

Note that the changes to $\mathcal{F}_{\text{fPAKE}}^M$ only affect the case where both parties are honest and no man-in-the-middle attack is happening. Because if a man-in-the-middle attack is mounted then \mathcal{S} will send a TESTPWD query to $\mathcal{F}_{\text{fPAKE}}^M$ (see \mathbf{G}_6 and \mathbf{G}_8) either marking $\mathcal{F}_{\text{fPAKE}}^M$'s record $(\mathcal{P}_i, \text{pw}^{\mathcal{P}_i})$ as **interrupted** or as **compromised**, so the record will not be **fresh**. The rest of the NEWKEY interface of $\mathcal{F}_{\text{fPAKE}}^M$ was already like in Fig. 3 since \mathbf{G}_4 .

In the first case role = receiver only the output of the party is changed. This is because the keys K_1, \dots, K_n (output by $\mathcal{F}_{\text{fPAKE}}^j$) were only used internally by \mathcal{P}_i in \mathbf{G}_9 . In this game, \mathcal{P}_i produces output as the ideal functionality $\mathcal{F}_{\text{fPAKE}}$ provides the dummy party with output on receiving NEWKEY from \mathcal{S} . In \mathbf{G}_9 the output was $H_1((\text{sid}, \text{sid}_{\text{H}}), s_0)$. In the second case role = sender, $\mathcal{F}_{\text{fPAKE}}$ provides the dummy party also with output instead of $H_1((\text{sid}, \text{sid}_{\text{H}}), s_0)$. Note that we assumed no man-in-the-middle attack and thus, $\mathcal{F}_{\text{fPAKE}}$ will produce output for \mathcal{P}_i and \mathcal{P}_{1-i} depending on the input passwords $\text{pw}^{\mathcal{P}_i}$ and $\text{pw}^{\mathcal{P}_{1-i}}$. If $d(\text{pw}^{\mathcal{P}_i}, \text{pw}^{\mathcal{P}_{1-i}}) \leq \delta$ then $\mathcal{F}_{\text{fPAKE}}$ provides matching outputs $\text{sk} = \text{sk}'$ to \mathcal{P}_i and \mathcal{P}_{1-i} and if $d(\text{pw}^{\mathcal{P}_i}, \text{pw}^{\mathcal{P}_{1-i}}) > \delta$ then the output sk of \mathcal{P}_i is drawn independently of the output sk' of \mathcal{P}_{1-i} . Clearly, by Lemma 6 both parties output the same random string $\text{sk} = \text{sk}'$ in \mathbf{G}_9 if $d(\text{pw}^{\mathcal{P}_i}, \text{pw}^{\mathcal{P}_{1-i}}) \leq \delta$, with overwhelming probability. Further, in \mathbf{G}_9 if $d(\text{pw}^{\mathcal{P}_i}, \text{pw}^{\mathcal{P}_{1-i}}) > \delta$, we can use Lemma 5 to see that both parties output different keys with overwhelming probability. That is because there are more than δ characters on which the passwords do not match, and so with overwhelming probability, the codeword $\mathbf{C}' = \mathbf{E} \oplus \mathbf{K}'$ computed by the receiver will not decode to \mathbf{C} . It remains to argue that \mathcal{Z} 's view of the protocol transcript does not change. Note that in \mathbf{G}_9 , (\mathbf{E}, h) was computed as $\mathbf{E} := \mathbf{C} \oplus \mathbf{K}$, where $\mathbf{K} =$

(K_1, \dots, K_n) were the outputs of $\mathcal{F}_{\text{iPAKE}}^j$ for $j \in [n]$. In \mathbf{G}_{10} , \mathbf{K} is uniformly random. The output K_j of $\mathcal{F}_{\text{iPAKE}}^j$ is also a uniformly random string when both parties are honest and no man-in-the-middle attack happened. The only difference is that both parties, \mathcal{P}_i and \mathcal{P}_{1-i} receive the same K_j as output for that session. But as we argued above, the receiver only uses K_j internally, so \mathcal{Z} does not see whether both parties output the same K_j for an iPAKE session. We get

$$|\Pr[\mathbf{G}_{10}] - \Pr[\mathbf{G}_9]| \leq \text{negl}(\lambda).$$

Game \mathbf{G}_{11} : Simulator does not get inputs of the parties anymore. In this game, we change the functionality. Instead of sending $(\text{NEWSESSION}, \text{sid}, \mathcal{P}_i, \text{pw}^{\mathcal{P}_i}, \text{role})$ to \mathcal{S} the functionality now keeps the input $\text{pw}^{\mathcal{P}_i}$ secret and sends only $(\text{NEWSESSION}, \text{sid}, \mathcal{P}_i, \text{role})$ to \mathcal{S} . Since \mathbf{G}_{10} , \mathcal{S} no longer needs the input $\text{pw}^{\mathcal{P}_i}$ to simulate the n instances of $\mathcal{F}_{\text{iPAKE}}$ for two honest parties with no man-in-the-middle attack. Further, we argued in \mathbf{G}_6 and \mathbf{G}_8 that in a session with a corrupted party or where a man-in-the-middle attack is mounted, \mathcal{S} does not need $\text{pw}^{\mathcal{P}_i}$ either, as it can either use random keys K_1, \dots, K_n or it learns which iPAKE sessions were successful through a TESTPWD query to $\mathcal{F}_{\text{fPAKE}}$. Overall, the simulator does not use the provided input $\text{pw}^{\mathcal{P}_i}$ anymore. Hence, we get

$$\Pr[\mathbf{G}_{11}] = \Pr[\mathbf{G}_{10}].$$

Note that \mathbf{G}_{11} is the same as the ideal execution of $\mathcal{F}_{\text{fPAKE}}$ with the simulator from Figs. 9 to 10. This concludes the proof.

B Additional preliminaries

We recap the implicit PAKE functionality of Dupont et al. [DHP⁺18] in Fig. 12 as well as the split authentication functionality of Barak et al. [BCL⁺05] in Fig. 13.

C Example Parameter Set with Timings

In this section, we compare the timings of list decoding and unique decoding algorithms for Reed-Solomon codes for a PAKE parameter setting related to conventional passwords.

Consider the setting of conventional alphanumeric passwords. There are 62 case-sensitive alphanumeric characters (A-Z, a-z and 0-9). For a uniformly random alphanumeric password to have 128 bits of entropy, it should be $\left\lceil \frac{128}{\log_2 62} \right\rceil = 22$ characters long.

Suppose that we would like to use the fuzzy PAKE construction of [DHP⁺18] with a tolerance of 3 password errors. Setting $n = 22$ and $k = 16$ achieves this error tolerance. In this case, our fuzzy PAKE construction also tolerates 3 errors.

The functionality $\mathcal{F}_{\text{iPAKE}}$ is parameterized by a security parameter λ . It interacts with an adversary \mathcal{A} and the (dummy) parties \mathcal{P}_0 and \mathcal{P}_1 via the following queries:

- **Upon receiving a query (NewSession, sid, pw ^{\mathcal{P}_i} , role) from party \mathcal{P}_i :**
 - Send (NEWSESSION, sid, \mathcal{P}_i , role) to \mathcal{A} ;
 - If one of the following is true, record $(\mathcal{P}_i, \text{pw}^{\mathcal{P}_i})$ and mark this record **fresh**:
 - * This is the first NEWSESSION query
 - * This is the second NEWSESSION query and there is a record $(\mathcal{P}_{1-i}, \text{pw}^{\mathcal{P}_{1-i}})$
- **Upon receiving a query (TestPwd, sid, \mathcal{P}_i , pw*) from \mathcal{A} :**
 If there is a **fresh** record $(\mathcal{P}_i, \text{pw}^{\mathcal{P}_i})$, then:
 - If $\text{pw}^{\mathcal{P}_i} = \text{pw}^*$, mark the record **compromised**;
 - If $\text{pw}^{\mathcal{P}_i} \neq \text{pw}^*$, mark the record **interrupted**.
- **Upon receiving a query (NewKey, sid, \mathcal{P}_i , sk) from \mathcal{A} , where $|\text{sk}| = \lambda$:**
 If there is no record of the form $(\mathcal{P}_i, \text{pw}^{\mathcal{P}_i})$, or if this is not the first NEWKEY query for \mathcal{P}_i , then ignore this query. Otherwise:
 - If at least one of the following is true, then output (sid, sk) to player \mathcal{P}_i :
 - * The record is **compromised**
 - * \mathcal{P}_i is **corrupted**
 - * The record is **fresh**, \mathcal{P}_{1-i} is **corrupted**, and there is a record $(\mathcal{P}_{1-i}, \text{pw}^{\mathcal{P}_{1-i}})$ with $\text{pw}^{\mathcal{P}_i} = \text{pw}^{\mathcal{P}_{1-i}}$
 - If this record is **fresh**, both parties are honest, there is a record $(\mathcal{P}_{1-i}, \text{pw}^{\mathcal{P}_{1-i}})$ with $\text{pw}^{\mathcal{P}_i} = \text{pw}^{\mathcal{P}_{1-i}}$, a key sk' was sent to \mathcal{P}_{1-i} , and $(\mathcal{P}_{1-i}, \text{pw}^{\mathcal{P}_{1-i}})$ was **fresh** at the time, then output (sid, sk') to \mathcal{P}_i ;
 - In any other case, pick a new random key sk' of length λ and send (sid, sk') to \mathcal{P}_i .
 - Mark the record $(\mathcal{P}_i, \text{pw}^{\mathcal{P}_i})$ as **completed**.

Fig. 12. Functionality $\mathcal{F}_{\text{iPAKE}}$ [DHP⁺18].

Functionality $\overline{\mathcal{F}}_{\text{SA}}$
<p><i>Initialization:</i></p> <p>Upon activation with input (INIT, sid), where sid = (\mathbb{P}, sid') and \mathbb{P} is a set of PIDs, record \mathbb{P} and forward it to the adversary.</p> <p>On (INIT, sid, \mathcal{P}, \mathbb{H}, sid$_{\mathbb{H}}$) from the adversary, verify that:</p> <ul style="list-style-type: none"> – $\mathcal{P} \in \mathbb{P}$ and $\mathcal{P} \in \mathbb{H}$, and – for all recorded sets \mathbb{H}', either <ul style="list-style-type: none"> • $\mathbb{H} \cap \mathbb{H}'$ contains only corrupt parties and sid$_{\mathbb{H}} \neq$ sid$_{\mathbb{H}'}$, or • $\mathbb{H}' = \mathbb{H}$ and sid$_{\mathbb{H}} =$ sid$_{\mathbb{H}'}$. <p>If so, then output (INIT, sid, sid$_{\mathbb{H}}$) to \mathcal{P} and record (\mathbb{H}, sid$_{\mathbb{H}}$) if not yet recorded. Else, do nothing.</p> <p><i>Message authentication:</i></p> <p>On input (SEND, sid, \mathcal{P}, \mathcal{P}', msg) from \mathcal{P}, where $\mathcal{P}' \in \mathbb{P}$, send ($\mathcal{P}$, \mathcal{P}', msg) to \mathcal{A} and add (\mathcal{P}, \mathcal{P}', msg) to an (initially empty) list \mathbb{W} of waiting messages. Note that the same triple can appear multiple times in the list. Else do nothing.</p> <p>On (DELIVER, sid, \mathcal{P}, \mathcal{P}', msg) from \mathcal{A}, proceed as follows:</p> <ul style="list-style-type: none"> – If \mathcal{P}' did not previously receive an (INIT, sid, \cdot) output then do nothing. – Else, if \mathcal{P} is in the authentication set \mathbb{H} of \mathcal{P}', and \mathcal{P} is honest, then do: <ul style="list-style-type: none"> • If there is a triple (\mathcal{P}, \mathcal{P}', msg) \in \mathbb{W}, then remove one appearance of the triple from \mathbb{W} and output (RECEIVED, sid, \mathcal{P}, \mathcal{P}', msg) to \mathcal{P}'. • Otherwise do nothing. – Else (i.e., \mathcal{P}' received (INIT, sid, \cdot), and either \mathcal{P} is corrupted or $\mathcal{P} \notin \mathbb{H}$), output (RECEIVED, sid, \mathcal{P}, \mathcal{P}', msg) to \mathcal{P}', regardless of \mathbb{W}.

Fig. 13. Ideal functionality $\overline{\mathcal{F}}_{\text{SA}}$ for split authentication [BCL⁺05].

Consider the Reed-Solomon code of blocklength n and rank k over \mathbb{F}_q , where $q = 2^{256} - 2^{32} - 2^9 - 2^8 - 2^7 - 2^6 - 2^4 - 1$, the prime used to define the base field for elliptic curve `secp256k1`. The time taken to decode a random codeword with 3 errors using Gao’s unique decoding algorithm is roughly 120ms, while the time taken to run the Guruswami-Sudan list-decoding algorithm is roughly 1.8s.

These timings were obtained in SageMath [The23] using a laptop with a 2.40 GHz Intel i9 processor. Minor modifications of Sage’s Gao and Guruswami-Sudan implementations were required to allow them to work with simple generator matrices rather than Generalized Reed-Solomon Code objects, as defining codes via their generator matrices was much faster for large q .

D On the similarities to AMD codes

In Section 4, we explained how our protocol can narrow down the list of “candidate decodings” produced by list-decoding to a single codeword by using the hint h . Cramer et al. [CDD⁺15] had a similar idea. To construct a robust secret

sharing scheme, they use a list-decodable code. But before encoding the message with this code, they apply an *Algebraic Manipulation Detection* (AMD) code to the message. Intuitively, these codes allow the manipulation of an encoded message to be detected, and abort decoding if the codeword was manipulated. Now, one could suspect that the hint $h = H_0(\mathbf{C})$ serves as AMD code in our construction. But we argue in the following that hashing the codeword and appending the hash is not an AMD code. Roughly speaking, if we used our construction without split authentication it would be trivial for an adversary to tamper with the codeword and the hint, such that a receiver decodes the wrong message. We elaborate on this more formally.

An AMD code is defined as follows:

Definition 3 (Algebraic Manipulation Detection Code). *Let q be a prime-power, $n > k$ be integers and $\rho > 0$. A (k, n, ρ) -AMD \mathcal{C} consists of a probabilistic encoding algorithm $\text{Enc} : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^n$ and a deterministic decoding algorithm $\text{Dec} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^k \cup \{\perp\}$, such that the following holds for every $\mathbf{x} \in \mathbb{F}_q^k$:*

- Correctness: $\text{Dec}(\text{Enc}(\mathbf{x})) = \mathbf{x}$ with probability 1.
- Manipulation detection: for every $\Delta \in \mathbb{F}_q^n$ and for $\mathbf{C} \leftarrow \text{Enc}(\mathbf{x})$ it holds that $\text{Dec}(\mathbf{C} \oplus \Delta) \in \{\perp, \mathbf{x}\}$ except with probability at most ρ .

We can try to formalize the hint mechanism as above by setting $\text{Enc}(\mathbf{x}) := \mathbf{x} \| H(\mathbf{x})$. The decoding parses a codeword as $\mathbf{C} := \mathbf{x} \| h$ and outputs \mathbf{x} if $h = H(\mathbf{x})$ and else an error symbol \perp . But it turns out that this is *not* an AMD code. The reason is that manipulation detection must hold for all $\Delta \in \mathbb{F}_q^n$. Let $\mathbf{x}' \in \mathbb{F}_q^k$ with $\mathbf{x}' \neq \mathbf{x}$. Then manipulation detection must hold in particular for the vector

$$\Delta = (\mathbf{x} \oplus \mathbf{x}', H(\mathbf{x}) \oplus H(\mathbf{x}')).$$

But the decoding algorithm would now output \mathbf{x}' and not \perp or \mathbf{x} on input $\mathbf{C} \oplus \Delta$.

Now, why does this weakness not allow for an attack on our fuzzy PAKE protocol? The reason is that we send messages over a split-authenticated channel. If both parties are authenticated, i.e., in the same authentication set \mathbb{H} then an adversary can not tamper with the message by adding Δ as above to the codeword. Otherwise, if there is a MitM attack happening then we distinguish two cases. (1) If the adversary (playing as corrupted sender) guesses at least $n - \delta$ password characters of the receiver correctly then the receiver gets iPAKE keys of the adversary's choice. That means, the adversary can make the client output a shared key of its choice anyway, by changing the iPAKE keys accordingly. So tampering with h does not give the adversary additional power. (2) If the adversary guessed less than $n - \delta$ password characters correctly then there is already so much randomness added to the codeword by the iPAKE keys that the list decoding algorithm will output some specific codeword only with negligible probability.

E UC Execution and Instantiation of our fuzzy PAKE

We give a UC execution of our fuzzy PAKE protocol of Figure 8 using the interfaces of $\mathcal{F}_{\text{fPAKE}}$ in Fig. 14. We can instantiate the building blocks of our

The parties \mathcal{P}_0 and \mathcal{P} communicate with the hybrid functionalities \mathcal{F}_{SA} and \mathcal{F}_{RO} .

On (NEWSESSION, $\text{sid}, \mathcal{P}_i, \text{role}$) from $\mathcal{F}_{\text{fPAKE}}$ or (INIT, $(\{\mathcal{P}_i, \mathcal{P}_{1-i}\}, \text{sid})$) from \mathcal{Z} in the name of \mathcal{P}_i to \mathcal{F}_{SA}

- Record $(\text{sid}, \text{role})$
- Send (INIT, $(\{\mathcal{P}_i, \mathcal{P}_{1-i}\}, \text{sid})$) to \mathcal{F}_{SA} .

On (INIT, sid, sid_H) from \mathcal{F}_{SA}

- For $j \in [n]$:
 - Run $X_j^* \leftarrow \pi_j(\mathcal{P}_i, (\text{sid}, \text{sid}_H, j), \text{pw}_j)$
- Send (SEND, $\text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, (X_1^*, \dots, X_n^*)$) to \mathcal{F}_{SA}

On (DELIVER, $\text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, (Y_1^*, \dots, Y_n^*)$) from \mathcal{F}_{SA} and if \mathcal{P}_i recorded $(\text{sid}, \text{sender})$

- For $j \in [n]$:
 - Run $K_j \leftarrow \pi_j(\mathcal{P}_i, (\text{sid}, \text{sid}_H, j), Y_j^*)$
- Let $\mathbf{K} := (K_1, \dots, K_n)$
- $s = (s_0, \dots, s_{k-1}) \xleftarrow{\$} \mathbb{F}_q^k$
- $\mathbf{C} \leftarrow \text{Enc}(s)$
- $\mathbf{E} := \mathbf{C} \oplus \mathbf{K}$
- $h := H_0((\text{sid}, \text{sid}_H), \mathbf{C})$
- Send (SEND, $\text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, (\mathbf{E}, h)$) to \mathcal{F}_{SA}
- Send $(\text{sid}, \text{sid}_H, s_0)$ to \mathcal{F}_{RO} . On response ρ from \mathcal{F}_{RO} output ρ .

On (DELIVER, $\text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, (Y_1^*, \dots, Y_n^*)$) from \mathcal{F}_{SA} and if \mathcal{P}_i recorded $(\text{sid}, \text{receiver})$

- For $j \in [n]$:
 - Run $K_j \leftarrow \pi_j(\mathcal{P}_i, (\text{sid}, \text{sid}_H, j), Y_j^*)$
- Let $\mathbf{K} := (K_1, \dots, K_n)$
- Append \mathbf{K} to $(\text{sid}, \text{receiver})$

On (NEWSESSION, $\text{sid}, \mathcal{P}_i, \text{role}$) from $\mathcal{F}_{\text{fPAKE}}$ or (INIT, $(\{\mathcal{P}_i, \mathcal{P}_{1-i}\}, \text{sid})$) from \mathcal{Z} in the name of \mathcal{P}_i to \mathcal{F}_{SA}

- Send (INIT, $(\{\mathcal{P}_i, \mathcal{P}_{1-i}\}, \text{sid})$) to \mathcal{F}_{SA} .

On (DELIVER, $\text{sid}, \mathcal{P}_i, \mathcal{P}_{1-i}, \text{msg} := (\mathbf{E}', h')$) from \mathcal{F}_{SA} and if \mathcal{P}_i recorded $(\text{sid}, \text{receiver})$:

- Run $\{\mathbf{C}^1, \dots, \mathbf{C}^\ell\} \leftarrow \text{LDec}(\mathbf{E}' \oplus \mathbf{K}')$
- If $h = H_0((\text{sid}, \text{sid}_H), \mathbf{C}^j)$ for a $j \in [\ell]$
 - $s' \leftarrow \text{Dec}(\mathbf{C}^j)$
 - Send $(\text{sid}, \text{sid}_H, s_0)$ to \mathcal{F}_{RO} . On response ρ from \mathcal{F}_{RO} output ρ .
- else output $u \xleftarrow{\$} \{0, 1\}^\lambda$

Fig. 14. A UC Execution of our Protocol.

protocol with EKE2 [BPR00] as iPAKE and a signature-based transformation to instantiate \mathcal{F}_{SA} by Barak et al. [BCL⁺05]. The fully instantiated protocol is depicted in Fig. 15.

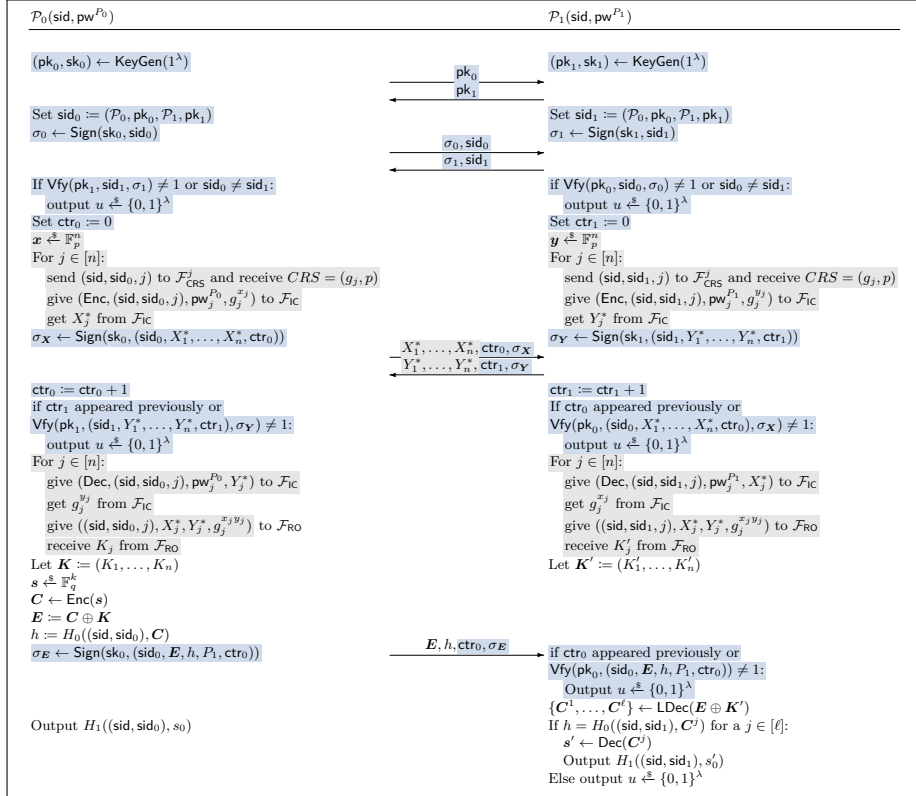


Fig. 15. The protocol in the $\mathcal{F}_{\text{CRS}}, \mathcal{F}_{\text{RO}}, \mathcal{F}_{\text{IC}}$ -hybrid model without any black-box PAKE functionality but with explicit EKE and with signing like in the split transform. Assume that a party aborts the execution if any check fails. Parts in grey are for the n EKE2 key exchanges and parts in blue are for the split-transform-like authentication.