Simpler Session-Key Generation from Short Random Passwords*

Minh-Huyen Nguyen † Salil Vadhan ‡

July 30, 2004

Abstract

Goldreich and Lindell (CRYPTO '01) recently presented the first protocol for password-authenticated key exchange in the standard model (with no common reference string or set-up assumptions other than the shared password). However, their protocol uses several heavy tools and has a complicated analysis.

We present a simplification of the Goldreich–Lindell (GL) protocol and analysis for the special case when the dictionary is of the form $\mathcal{D} = \{0,1\}^d$, i.e. the password is a short random string (like an ATM PIN number). Our protocol can be converted into one for arbitrary dictionaries using a common reference string of logarithmic length. The security bound achieved by our protocol is somewhat worse than the GL protocol. Roughly speaking, our protocol guarantees that the adversary can "break" the scheme with probability at most $O(\mathrm{poly}(n)/|\mathcal{D}|)^{\Omega(1)}$, whereas the GL protocol guarantees a bound of $O(1/|\mathcal{D}|)$.

We also present an alternative, more natural definition of security than the "augmented definition" of Goldreich and Lindell, and prove that the two definitions are equivalent.

---

*An extended abstract of this paper appeared in the *First Theory of Cryptography Conference (TCC '04)* [23].

†Harvard University, Cambridge, MA. E-mail: `mnguyen@eecs.harvard.edu`. Supported by NSF grant CCR-0205423 and ONR grant N00014-04-1-0478.

‡Harvard University, Cambridge, MA. E-mail: `salil@eecs.harvard.edu`. Supported by NSF grant CCR-0205423, a Sloan Research Fellowship, and ONR grant N00014-04-1-0478. Part of this work done while at the Radcliffe Institute for Advanced Study.

Contents

1 Introduction **1**

2 Definition of Security **3**
2.1 The Initial Definition . . . 3
2.2 Security with respect to the Session-Key Challenge . . . 5
2.3 Security with respect to the Environment . . . 6

3 An Overview of the Protocol **9**
3.1 Secure Polynomial Evaluation . . . 10
3.2 Motivation for our Protocol . . . 10
3.3 Description of the Protocol . . . 11

4 Security Theorems **13**

5 Overview of the Proof **15**

6 Proof of Security against Passive Adversaries **18**

7 Key-Match Property for the First Scheduling **19**

8 Key-Match Property for the warm-up case **21**

9 Key-Match Property for the Second Scheduling **22**
9.1 Mental Experiment Protocol . . . 22
9.2 Key-match Property for the Second Scheduling . . . 22

10 Adapting the GL Techniques to our Protocol **25**
10.1 Simulation of the (A, C') Execution . . . 26
10.2 Simulation of the (C, B) Execution . . . 27
10.3 Security Theorem . . . 29

References **31**

A Secure Two-Party Computation **33**

B Almost Pairwise Independence **33**

1 Introduction

What is the minimal amount of information that two parties must share in order to perform nontrivial cryptography? This fundamental question is at the heart of many of the major distinctions we draw in cryptography. Classical private-key cryptography assumes that the legitimate parties share a long random key. Public-key cryptography mitigates this by allowing the sharing of information to be done through public keys that need not be hidden from the adversary. However, in both cases, the amount of information shared by the legitimate parties (e.g. as measured by mutual information) needs to be quite large. Indeed, the traditional view is that security comes from the adversary's inability to exhaustively search the keyspace.

Thus it is very natural to ask: *can we do nontrivial cryptography using "low-entropy" keys?* That is, using a keyspace that is feasible to exhaustively search. In addition to being a natural theoretical question, it has clear relevance to the many "real-life" situations where we need security but only have a low-entropy key (e.g. an ATM PIN number, or human-chosen password on a website).

Public-key cryptography provides an initial positive answer to this question: key-exchange protocols, as in [10], do not require *any* prior shared information. However, this holds only for passive adversaries, and it is well known that without any prior shared information between the legitimate parties, an active adversary can always succeed through a person-in-the-middle attack. Thus, it remains an interesting question to achieve security against active adversaries using a low-entropy shared key. This has led researchers to consider the problem of *password-authenticated key exchange,* which we describe next.

Password-Authenticated Key Exchange. The password-authenticated key exchange problem was first suggested by Bellovin and Merritt [4]. We assume that two parties, Alice and Bob, share a password w chosen uniformly at random from a dictionary $\mathcal{D} \subseteq \{0,1\}^n$. This dictionary can be very small, e.g. $|\mathcal{D}| = \mathrm{poly}(n)$, and in particular it may be feasible for an adversary to exhaustively search it. Our aim is to construct a protocol enabling Alice and Bob to generate a "random" session key $K \in \{0,1\}^n$, which they can subsequently use for standard private-key cryptography. We consider an active adversary that completely controls the communication channel between Alice and Bob. The adversary can intercept, modify, drop, and delay messages, and in particular can attempt to impersonate either party through a person-in-the-middle attack.

Our goal is that, even after the adversary mounts such an attack, Alice and Bob will generate a session key that is indistinguishable from uniform even given the adversary's view. However, our ability to achieve this goal is limited by two unpreventable attacks. First, since the adversary can block all communication, it can prevent one or both of the parties from completing the protocol and obtaining a session key. Second, the adversary can guess a random password $\tilde{w} \leftarrow \mathcal{D}$ and attempt to impersonate one of the parties. With probability $1/|\mathcal{D}|$, the guess equals the real password (i.e., $\tilde{w} = w$), and the adversary will succeed in the impersonation and therefore learn the session key. Thus, we revise our goal to effectively limit the adversary to these two attacks. Various formalizations for this problem have been developed through several works [3, 18, 26, 2, 7, 15]. We follow the definitional framework of Goldreich and Lindell [15], which is described in more detail in Sec. 2.

In addition to addressing what can be done with a minimal amount of shared information, the study of this problem is useful as another testbed for developing our understanding of *concurrency* in cryptographic protocols. The concurrency implicitly arises from the person-in-the-middle attack, which we can view as two simultaneous executions of the protocol, one between Alice and the adversary and the other between Bob and the adversary.

The first protocols for the password-authenticated key exchange problem were proposed in the security literature, based on informal definitions and heuristic arguments (e.g. [5, 28]). The first rigorous proofs of security were given in the random oracle model [2, 7]. Only recently were rigorous solutions without random oracles given, in independent works by Goldreich and Lindell [15] and Katz, Ostrovsky, and Yung [19]. One of the main differences between these two protocols is that the KOY protocol (and the subsequent protocols of [20, 13]) is in the "public parameters model," requiring a string to be generated and published by a trusted third party, whereas the GL protocol requires no set-up assumption other than the shared password. Thus, even though the KOY protocol has a number of practical and theoretical advantages over the GL protocol (which we will not enumerate here), the GL protocol is more relevant to our initial question about the minimal amount of shared information needed for nontrivial cryptography.

The Goldreich–Lindell Protocol. As mentioned above, the Goldreich–Lindell protocol [15] is remarkable in that the only set-up assumption it requires is that the two parties share a password chosen at random from an arbitrary dictionary. Their protocol can be based on general complexity assumptions (the existence of trapdoor permutations), can be implemented in a constant number of rounds (under stronger assumptions), and achieves a nearly optimal security bound (the adversary has probability only $O(1/|\mathcal{D}|)$ of "breaking" the scheme).

Despite giving such a strong result, the Goldreich–Lindell protocol does not leave us with a complete understanding of the password-authenticated key exchange problem. First, the protocol makes use of several "heavy" tools: secure two-party polynomial evaluation (building on [22], who observed that this yields a protocol for password-authenticated key exchange against passive adversaries), nonmalleable commitments (as suggested in [6]), and the specific concurrent zero-knowledge proof of Richardson and Kilian [25]. It is unclear whether all of these tools are really essential for solving the key exchange problem. Second, the proof of the protocol's security is extremely complicated. Goldreich and Lindell do introduce nice techniques for analyzing concurrent executions (arising from the person-in-the-middle attack) of two-party protocols whose security is only guaranteed in the stand-alone setting (e.g. the polynomial evaluation). But these techniques are applied in an intricate manner that seems inextricably tied to the presence of the nonmalleable commitment and zero-knowledge proof. Finally, finding an efficient instantiation of the Goldreich–Lindell protocol would require finding efficient instantiations of all three of the heavy tools mentioned above, which seems difficult. In particular, the Richardson-Kilian zero-knowledge proof is used to prove an **NP** statement that asserts the consistency of a transcript of the nonmalleable commitment, a standard commitment, and the output of an iterated one-way permutation. For such an **NP** statement, it seems difficult to avoid using a generic zero-knowledge proof system for **NP**, which are almost always inefficient due to the use of Cook's theorem.

Our Protocol. Our main result is a simplification of the Goldreich–Lindell protocol and analysis for the special case when the dictionary is of the form $\mathcal{D} = \{0,1\}^d$, i.e. the password is a short random string (like an ATM PIN number)[1]. This special case still retains many of the key features of the problem: the person-in-the-middle attack and the resulting concurrency issues are still present, and the adversary can still exhaustively search the dictionary (since we allow the password length d to be as small as $O(\log n)$, where n is the security parameter). Moreover, our protocol can be converted into one for arbitrary dictionaries in the common reference string model (using the common reference string as the seed of a randomness extractor [24]). For dictionaries $\mathcal{D} \subset \{0,1\}^n$, the common reference string is a uniform string of only logarithmic length (specifically, $O(\log n + \log|\mathcal{D}|)$), and thus retains the spirit of minimizing the amount of shared information between the legitimate parties. In contrast, the previous protocols in the public parameters model [19, 20, 13] require a public string of length $\mathrm{poly}(n)$ with special number-theoretic structure.

The main way in which we simplify the GL protocol is that we remove the nonmalleable commitments and the Richardson–Kilian zero-knowledge proof. Instead, our protocol combines secure polynomial evaluation with a combinatorial tool (almost pairwise independent hashing), in addition to using "lightweight" cryptographic primitives also used in [15] (one-way permutations, one-time MACs, standard commitments). Our analysis is also similarly simpler. While it has the same overall structure as the analysis in [15] and utilizes their techniques for applying the stand-alone properties of the polynomial evaluation in the concurrent setting, it avoids dealing with the nonmalleable commitments and the zero-knowledge proof (which is the most complex part of the GL analysis).

Removing the nonmalleable commitments and the RK zero-knowledge proof has two additional implications. First, finding an efficient implementation of our protocol only requires finding an efficient protocol for secure polynomial evaluation (in fact, only for linear polynomials).[2] Since this is a highly structured special case of secure two-party computation, it does not seem beyond reach to find an efficient protocol. Indeed, Naor and Pinkas [22] have already given an efficient polynomial evaluation protocol for passive adversaries. Second, our protocol can be implemented in a constant number of rounds assuming only the existence of trapdoor permutations, whereas implementing the Goldreich–Lindell protocol in a constant number of rounds

[1] More generally, the password can be chosen uniformly from any dictionary of size 2^d whose elements we can efficiently enumerate because the enumeration provides a bijection with $\{0,1\}^d$.

[2] Actually, we require a slightly augmented form of polynomial evaluation, in which one of the parties commits to its input beforehand and the protocol ensures consistency with this committed input.

requires additional assumptions, such as the existence of claw-free permutations (for [25]) and some sort of exponential hardness assumption (to use [1]).

We note that the security bound achieved by our protocols is somewhat worse than in previous works. Roughly speaking, our protocol guarantees that the adversary can "break" the scheme with probability at most $O\left(\frac{\mathrm{poly}(n)}{|\mathcal{D}|}\right)^{\Omega(1)}$, whereas previous works guarantee a bound of $O(1/|\mathcal{D}|)$.

An additional result in our paper involves the definition of security in [15]. As pointed out by Rackoff (cf., [2]), it is important that a key exchange protocol provide security even if the party who completes the protocol first starts using the generated key in some application before the second party completes the protocol. In order to address this issue, Goldreich and Lindell [15] augmented their definition with a "session-key challenge", in which the adversary is given either the generated key or a uniform string with probability $1/2$ upon the first party's completion of the protocol. We present an arguably more natural definition that directly models the use of the generated key in an arbitrary application, and prove its equivalence to the augmented definition of Goldreich and Lindell [15]. (This result is analogous to the result of Shoup [26] for non-password-based key exchange protocols.)

2 Definition of Security

We adopt the notation of Goldreich and Lindell and refer the reader to [15] for more details.

- C denotes the probabilistic polynomial time adversary through which the honest parties A and B communicate. We model this communication by giving C oracle access to a single copy of A and a single copy of B. Here the oracles A and B have memory and represent honest parties executing the session-key generation protocol. We denote by $C^{A(x),B(y)}(\sigma)$ an execution of C with auxiliary input σ when it communicates with A and B, with respective inputs x and y. The output of the channel C from this execution is denoted by $\mathrm{output}\left(C^{A(x),B(y)}(\sigma)\right)$.
- The security parameter is denoted by n. The password dictionary is denoted by $\mathcal{D} \subseteq \{0,1\}^n$ and we write $\epsilon = \frac{1}{|\mathcal{D}|}$.

We denote by U_n the uniform distribution over strings of length n, by $\mathrm{neg}(n)$ a negligible function and write $x \xleftarrow{\text{R}} S$ when x is chosen uniformly from the set S.

For a function $\gamma : \mathbb{N} \to [0,1]$, we say that the probability ensembles $\{X_n\}$ and $\{Y_n\}$ are $(1-\gamma)$-*indistinguishable* (denoted by $\{X_n\} \stackrel{\gamma}{\equiv} \{Y_n\}$) if for every nonuniform PPT distinguisher D and all n,

$$|\Pr[D(X_n) = 1] - \Pr[D(Y_n) = 1]| < \gamma(n) + \mathrm{neg}(n)$$

We say that $\{X_n\}$ and $\{Y_n\}$ are *computationally indistinguishable*, which we denote by $X_n \stackrel{\text{c}}{\equiv} Y_n$, if they are 1-indistinguishable. We say that $\{X_n\}$ is $(1-\gamma)$ *pseudorandom* if it is $(1-\gamma)$ indistinguishable from U_n.

We will now formalize the problem of session-key generation using human passwords. We first follow the presentation of the problem as in [15] and then contrast it with our definition.

2.1 The Initial Definition

The definition in [15] follows the standard paradigm for secure computation: define an ideal functionality (using a trusted third party) and require that every adversary attacking the real protocol can be simulated by an ideal adversary attacking the ideal functionality. Note that in the real protocol, the active adversary C can prevent one or both of the parties A and B from having an output. Thus, in the ideal model, we will allow $\mathrm{C_{ideal}}$ to specify two input bits, dec_C^A and dec_C^B, which determine whether A and B obtain a session key or not.

Ideal model Let A, B be the honest parties and let $\mathrm{C_{ideal}}$ be any PPT ideal adversary with auxiliary input σ.

1. A and B receive $w \stackrel{\mathrm{R}}{\leftarrow} \mathcal{D}$.
2. A and B both send w to the trusted party.
3. $\mathrm{C}_{\mathrm{ideal}}$ sends $(\mathrm{dec}_C^A, \mathrm{dec}_C^B)$ to the trusted party.
4. The trusted party chooses $K \stackrel{\mathrm{R}}{\leftarrow} \{0,1\}^n$. For each party $i \in \{A, B\}$, the trusted party sends K if $\mathrm{dec}_C^i = 1$ and sends $\perp$ if $\mathrm{dec}_C^i = 0$.

The ideal distribution is defined by:

$$\mathrm{IDEAL}_{\mathrm{C}_{\mathrm{ideal}}}(\mathcal{D}, \sigma) = (w, \mathrm{output}(A), \mathrm{output}(B), \mathrm{output}(\mathrm{C}_{\mathrm{ideal}}(\sigma)))$$

We note that this description of the ideal model differs slightly from the original definition in [15] since we allow B to finish first and A to reject in the ideal model (this is to take into account protocols in which no party is guaranteed to terminate with a session key). However, as described in Section 3.3, our protocol will guarantee that A always accepts. Moreover, we will show that any real adversary can be simulated by an ideal adversary who always chooses A to conclude first and accept.

Real model Let A, B be the honest parties and let C be any PPT real adversary with auxiliary input σ.

At some initialization stage, A and B receive $w \stackrel{\mathrm{R}}{\leftarrow} \mathcal{D}$. The real protocol is executed by A and B communicating via C. We will augment C's view of the protocol with A and B's decision bits, denoted by dec_A and dec_B, where $\mathrm{dec}_A = \texttt{reject}$ if $\mathrm{output}(A) = \perp$, and $\mathrm{dec}_A = \texttt{accept}$ otherwise (dec_B is defined similarly). (Indeed, in typical applications, the decisions of A and B will be learned by the real adversary C: if A obtains a session key, then it will use it afterwards; otherwise, A will stop communication or try to re-initiate an execution of the protocol.) C's augmented view is denoted by $\mathrm{output}(C^{A(w),B(w)}(\sigma))$.

The real distribution is defined by:

$$\mathrm{REAL}_C(\mathcal{D}, \sigma) = (w, \mathrm{output}(A), \mathrm{output}(B), \mathrm{output}(C^{A(w),B(w)}(\sigma)))$$

One might want to say that a protocol for password-based session-key generation is secure if the above ideal and real distributions are computationally indistinguishable. Unfortunately, as pointed in [15], an active adversary can guess the password and successfully impersonate one of the parties with probability $\frac{1}{|\mathcal{D}|}$. This implies that the real and ideal distributions are always distinguishable with probability at least $\frac{1}{|\mathcal{D}|}$. Thus we will only require that the distributions be distinguishable with probability at most $O(\gamma)$ where the goal is to make γ as close to $\frac{1}{|\mathcal{D}|}$ as possible. In the case of a passive adversary, we require that the real and ideal distributions be computationally indistinguishable (for all subsequent definitions, this requirement will be implicit).

Definition 2.1 (Initial definition) *A protocol for password-based authenticated session-key generation is* $(1-\gamma)$-secure *for the dictionary $\mathcal{D} \subseteq \{0,1\}^n$ (where γ is a function of the dictionary size $|\mathcal{D}|$ and n) if:*

1. *For every real* passive *adversary, there exists an ideal adversary* $\mathrm{C}_{\mathrm{ideal}}$ *which always sends (1,1) to the trusted party such that for every auxiliary input $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$,*

$$\{\mathrm{IDEAL}_{\mathrm{C}_{\mathrm{ideal}}}(\mathcal{D}, \sigma)\}_\sigma \stackrel{\mathrm{c}}{\equiv} \{\mathrm{REAL}_C(\mathcal{D}, \sigma)\}_\sigma$$

2. *For every real adversary C, there exists an ideal adversary* $\mathrm{C}_{\mathrm{ideal}}$ *such that for every auxiliary input $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$,*

$$\{\mathrm{IDEAL}_{\mathrm{C}_{\mathrm{ideal}}}(\mathcal{D}, \sigma)\}_\sigma \stackrel{\mathrm{O}(\gamma)}{\equiv} \{\mathrm{REAL}_C(\mathcal{D}, \sigma)\}_\sigma$$

By the discussion above, the best we can hope for is $\gamma = \frac{1}{|\mathcal{D}|}$. Note that in [15], their definition and protocol refer to any dictionary $\mathcal{D} \subseteq \{0,1\}^n$ and $\gamma = \frac{1}{|\mathcal{D}|}$. In contrast, our protocol will be $(1-\gamma)$-secure for dictionaries of the form $\mathcal{D} = \{0,1\}^d$ and $\gamma = \left(\frac{\mathrm{poly}(n)}{|\mathcal{D}|}\right)^{\Omega(1)}$.

2.2 Security with respect to the Session-Key Challenge

The above definition is actually not completely satisfying because of a subtle point raised by Rackoff: the adversary controls the scheduling of the interactions (A, C) and (C, B) so the honest parties do not necessarily end at the same time. A might use its session key K_A before the interaction (C, B) is completed: A's use of K_A leaks information which C might use in its interaction with B to learn K_A, K_B or the password w.

In [15], Goldreich and Lindell augment the above definition with a *session-key challenge* to address this issue. Suppose that A completes the protocol first and outputs a session key K, then the adversary is given a session key challenge K_β, which is the session key K with probability $1/2$ (i.e. $\beta = 1$) or a truly random string K_0 with probability $1/2$ (i.e. $\beta = 0$). The adversary C will be given the session-key challenge in both the ideal and real models, as soon as the first honest party outputs a session key K. We call the resulting definition *security with respect to the session-key challenge.*

Ideal model Let A, B be the honest parties and let $\mathrm{C}_{\mathrm{ideal}}$ be any PPT ideal adversary with auxiliary input σ.

1. A and B receive $w \stackrel{\mathrm{R}}{\leftarrow} \mathcal{D}$.
2. A and B both send w to the trusted party
3. $\mathrm{C}_{\mathrm{ideal}}$ decides which party $i \in \{A, B\}$ concludes first and whether it is a successful execution or not, i.e. $\mathrm{C}_{\mathrm{ideal}}$ sends a bit dec_C^i to the trusted party.
4. The trusted party chooses $K \stackrel{\mathrm{R}}{\leftarrow} \{0,1\}^n$. If $\mathrm{dec}_C^i = 1$, the trusted party sends K to party i; otherwise it sends $\perp$.
5. Session-key challenge: if party i received $\perp$, then the trusted party gives $\perp$ to $\mathrm{C}_{\mathrm{ideal}}$. Otherwise, the trusted party chooses $\beta \stackrel{\mathrm{R}}{\leftarrow} \{0,1\}$ and gives $\mathrm{C}_{\mathrm{ideal}}$ the string K_β where $K_1 = K$ and $K_0 \stackrel{\mathrm{R}}{\leftarrow} \{0,1\}^n$.
6. $\mathrm{C}_{\mathrm{ideal}}$ decides whether the second party's execution is successful or not, i.e. $\mathrm{C}_{\mathrm{ideal}}$ sends dec_C^j for $j \neq i$ to the trusted party.
7. If $\mathrm{dec}_C^j = 1$, the trusted party sends K to party j. Otherwise, it sends $\perp$.

The augmented ideal distribution is defined by:

$$\mathrm{IDEAL} - \mathrm{SK}_{\mathrm{C}_{\mathrm{ideal}}}(\mathcal{D}, \sigma) = (w, \mathrm{output}(A), \mathrm{output}(B), \mathrm{output}(\mathrm{C}_{\mathrm{ideal}}(\sigma, K_\beta)), \beta)$$

Real model C has oracle access to a single copy of $A(w)$ and a single copy of $B(w)$. The adversary C controls which party (A or B) concludes first. If the first party concludes with $\perp$, then C is given $\perp$. If the first party concluding outputs locally a session key K, then a bit $\beta \stackrel{\mathrm{R}}{\leftarrow} \{0,1\}$ is chosen and C is given the session-key challenge K_β where $K_1 = K$ and $K_0 \stackrel{\mathrm{R}}{\leftarrow} \{0,1\}^n$. C completes its interaction with the other party.

The augmented real distribution is defined by:

$$\mathrm{REAL} - \mathrm{SK}_C(\mathcal{D}, \sigma) = (w, \mathrm{output}(A), \mathrm{output}(B), \mathrm{output}(C^{A(w), B(w)}(\sigma, K_\beta)), \beta)$$

Definition 2.2 (Security with respect to the session-key challenge [15]) *A protocol for password-based authenticated session-key generation is* $(1-\gamma)$-secure with respect to the session-key challenge *for the dictionary* $\mathcal{D} \subseteq \{0,1\}^n$ *if for every real adversary* C*, there exists* $\mathrm{C}_{\mathrm{ideal}}$ *such that for every auxiliary input* $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$,

$$\{\mathrm{IDEAL} - \mathrm{SK}_{\mathrm{C}_{\mathrm{ideal}}}(\mathcal{D}, \sigma)\}_\sigma \stackrel{\mathrm{O}(\gamma)}{\equiv} \{\mathrm{REAL} - \mathrm{SK}_C(\mathcal{D}, \sigma)\}_\sigma$$

Goldreich and Lindell give some intuition as to why the session-key challenge solves the flaw mentioned earlier. First, note that the ideal adversary cannot distinguish between the case $\beta = 0$ and the case $\beta = 1$ since in the ideal model, both K_0 and K are truly uniform strings. Consider the real adversary who has been given the session-key challenge: if C has been given K_0, then the session-key challenge does not help C in attacking the protocol, since C could generate K_0 on its own. Suppose that instead C has been given K and that C can somehow use it to attack the protocol (this corresponds to the situation where A uses the session-key K; $C(K)$ can simulate A's use of the key), then it would mean that C can tell if it is in the case $\beta = 0$ or $\beta = 1$, which is not possible if the protocol is secure with respect to the session-key challenge.

2.3 Security with respect to the Environment

Our intuitive notion of security is that *no matter how A uses its session-key K* before the execution (C, B) is completed, the ideal and real distributions should be $(1 - O(\gamma))$-indistinguishable. It is not immediate that the session-key challenge captures this. Thus we propose an alternative augmentation to Definition 2.1 that corresponds more directly to this goal.

We model the different ways the party A could use its session-key K by considering an arbitrary probabilistic polynomial time machine Z which is given the key K (as soon as A outputs a session-key K) and interacts with the adversary in both the ideal and real models. This is similar to the "application" queries in Shoup's model for (non-password-based) secure key-exchange [26], which was later extended to password protocols in [7]. Z can also be thought of in terms of "environment" as in the definition of universal composability by Canetti [8]: Z models an arbitrary environment (or application) in which the key generated by the session-key generation protocol is used (note that this is not as general as the definition of Canetti since the environment Z is only given the session-key and not the password w).

Examples of environments follow:

1. $Z(K) = \perp$: A does not use its session-key.
2. $Z(K) = K$: A publicly outputs its session-key.
3. $Z(K) = \begin{cases} K & \text{with probability } 1/2, \\ U_n & \text{with probability } 1/2. \end{cases}$

   This corresponds to the session-key challenge.
4. $Z(K) = \mathrm{Enc}_K(0^n)$: A uses its session-key for secure private-key encryption.
5. C sends a query m_1, $Z(K)$ answers with $\mathrm{Enc}_K(m_1)$, C sends a query m_2, $Z(K)$ answers with $\mathrm{Enc}_K(m_2)$ and so on. This corresponds to an interactive environment Z which models a chosen plaintext attack.

We call the definition obtained by adding (in both the ideal and real models) the environment Z *security with respect to the environment*. Informally, a real protocol is secure with respect to the environment if every adversary attacking the real protocol and interacting with an arbitrary environment can be simulated, with probability $1 - O(\gamma)$, by an ideal adversary attacking the ideal functionality and interacting with the same environment in the ideal model. (More precisely, for every real adversary, there should be a single ideal adversary that simulates it well for *every* environment.)

Ideal model Let A and B be the honest parties, $\mathrm{C_{ideal}}$ any PPT ideal adversary with auxiliary input σ and Z any PPT with auxiliary input τ.

1. A and B receive $w \stackrel{R}{\leftarrow} \mathcal{D}$
2. A and B both send w to the trusted party.
3. $\mathrm{C_{ideal}}$ decides which party $i \in \{A, B\}$ concludes first and whether it is a successful execution or not, i.e. $\mathrm{C_{ideal}}$ sends dec_C^i to the trusted party.
4. The trusted party chooses $K \stackrel{R}{\leftarrow} \{0,1\}^n$. If $\mathrm{dec}_C^i = 1$, it sets $L_1 = K$; otherwise, $L_1 = \perp$. The trusted party sends L_1 to party i and Z.
5. $\mathrm{C_{ideal}}$ interacts with $Z(L_1, \tau)$.
6. $\mathrm{C_{ideal}}$ decides whether the second party's execution is successful or not, i.e. $\mathrm{C_{ideal}}$ sends dec_C^j for $j \neq i$ to the trusted party.
7. If $\mathrm{dec}_C^j = 1$, the trusted party sets $L_2 = K$; otherwise, $L_2 = \perp$. It sends L_2 to party j.

The ideal distribution is defined by:

$$\mathrm{IDEAL}_{Z,\tau,\mathrm{C_{ideal}}}(\mathcal{D}, \sigma) = (w, \mathrm{output}(A), \mathrm{output}(B), \mathrm{output}(Z(L_1, \tau)), \mathrm{output}({\mathrm{C_{ideal}}}^{Z(L_1,\tau)}(\sigma)))$$

Real model C has oracle access to a single copy copy of $A(w)$ and a single copy of $B(w)$. The adversary C controls which party (A or B) concludes first. Let $L_1 \in \{0,1\}^n \cup \perp$ be the output of the first concluding party. C interacts with $Z(L_1, \tau)$ and completes its interaction with the other party,

The real distribution is defined by:

$$\mathrm{REAL}_{Z,\tau,C}(\mathcal{D}, \sigma) = (w, \mathrm{output}(A), \mathrm{output}(B), \mathrm{output}(Z(L_1, \tau)), \mathrm{output}(C^{A(w),B(w),Z(L_1,\tau)}(\sigma)))$$

Definition 2.3 (Security with respect to the environment) *A protocol for password-based authenticated session-key generation is* $(1-\gamma)$-secure with respect to the environment *for the dictionary* $\mathcal{D} \subseteq \{0,1\}^n$ *if for every PPT* C*, there exists* $\mathrm{C}_{\mathrm{ideal}}$ *such that for every auxiliary input* $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$ *and every PPT* Z *with every auxiliary input* $\tau \in \{0,1\}^{\mathrm{poly}(n)}$*,*

$$\{\mathrm{IDEAL}_{Z,\tau,\mathrm{C}_{\mathrm{ideal}}}(\mathcal{D}, \sigma)\}_\sigma \overset{\mathrm{O}(\gamma)}{\equiv} \{\mathrm{REAL}_{Z,\tau,C}(\mathcal{D}, \sigma)\}_\sigma$$

Note that security with respect to the environment implies security with respect to the session-key challenge since it suffices to consider the PPT $Z(K)$ which generates $\beta \overset{\mathrm{R}}{\leftarrow} \{0,1\}$ and outputs the key K if $\beta = 1$ or a truly random string K_0 if $\beta = 0$. We show that the two definitions are actually equivalent:

Theorem 2.4 *A protocol* (A, B) *is* $(1-\gamma)$*-secure with respect to the session-key challenge iff it is* $(1-\gamma)$*-secure with respect to the environment.*

This is similar to a result of Shoup [26] showing the equivalence of his definition and the Bellare-Rogaway [3] definition for non-password-based key exchange. The "application" queries in Shoup's definition are analogous to our environment Z, and the "test" queries in [3] are analogous to the session-key challenge. Though both of these definitions have been extended to password-authenticated key exchange [7, 2], it is not immediate that Shoup's equivalence result extends directly to our setting. For example, the definitions of [3, 2] are not simulation-based and do not directly require that the password remain pseudorandom, whereas here we are relating two simulation-based definitions that do ensure the password's secrecy.

Given Theorem 2.4, the relationship between security with respect to the environment and security with respect to the session-key challenge is analogous to the relationship between semantic security and indistinguishability for encryption schemes [17, 21]. Though both are equivalent, the former captures our intuitive notion of security better, but the latter is typically easier to establish for a given protocol (as it involves only taking into account a specific environment Z). For conciseness of notation in the proof, we omit "output" in the distributions.

Proof: Let (A, B) be a protocol that is secure with respect to the session-key challenge. To prove the theorem, it suffices to prove that for every PPT C, there exists a PPT $\mathrm{C}_{\mathrm{ideal}}$ such that for every Z and every auxiliary input τ:

$$\{w, A, B, Z(M_1, \tau), C^{A(w),B(w),Z(M_1,\tau)}(\sigma)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{w, A, B, Z(L_1, \tau), {\mathrm{C}_{\mathrm{ideal}}}^{Z(L_1,\tau)}(\sigma)\}$$

where M_1 is the output of the first concluding party in the real execution $C^{A(w),B(w)}$ and L_1 is the output of the first concluding party in the ideal execution.

We denote by M_2 the output of the second concluding party in the real execution $C^{A(w),B(w)}$ and by L_2 the output of the second concluding party in the ideal execution. Hence, we want to prove that for every PPT C, there exists a PPT $\mathrm{C}_{\mathrm{ideal}}$ such that for every Z and every auxiliary input τ:

$$\{w, M_1, M_2, Z(M_1, \tau), C^{A(w),B(w),Z(M_1,\tau)}(\sigma)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{w, L_1, L_2, Z(L_1, \tau), {\mathrm{C}_{\mathrm{ideal}}}^{Z(L_1,\tau)}(\sigma)\}$$

where we also require that if $\mathrm{C}_{\mathrm{ideal}}$ sets $i \in \{A, B\}$ as the first concluding party, then i concludes first in the simulated view $\mathrm{C}_{\mathrm{ideal}}$ outputs.

We will introduce

- In the ideal model of Definition 2.2, the trusted party gives the adversary $\mathrm{C}_{\text{ideal}}$ the session-key challenge K_β

$$K_\beta = \begin{cases} L_1 & \text{if } \beta = 1 \text{ or } L_1 = \perp \\ U_n & \text{if } \beta = 0 \text{ and } L_1 \neq \perp \end{cases}$$

- In the real model of Definition 2.2, the trusted party gives the adversary C the session-key challenge K_β

$$K_\beta = \begin{cases} M_1 & \text{if } \beta = 1 \text{ or } M_1 = \perp \\ U_n & \text{if } \beta = 0 \text{ and } M_1 \neq \perp \end{cases}$$

We fix the real adversary C (against security with respect to the environment) and define the real adversary C' (against security with respect to the session-key challenge) that, on auxiliary input (σ, τ) and receiving K from the first concluding party, simulates $Z(K, \tau)$ on its own. Hence $C'^{A(w),B(w)}(\sigma, \tau, K) \equiv \{K, \tau, C^{A(w),B(w),Z(K,\tau)}(\sigma)\}$.

By security with respect to the session-key challenge, there exists an ideal adversary $\mathrm{C}_{\text{ideal}}{}'$ such that

$$\begin{aligned} \mathrm{IDEAL} - \mathrm{SK}_{\mathrm{C}'_{\text{ideal}}} &\overset{\mathrm{O}(\gamma)}{\equiv} \mathrm{REAL} - \mathrm{SK}_{C'} \\ \Rightarrow \{w, L_1, L_2, \mathrm{C}'_{\text{ideal}}(\sigma, \tau, K_\beta), \beta\} &\overset{\mathrm{O}(\gamma)}{\equiv} \{w, M_1, M_2, C'^{A(w),B(w)}(\sigma, \tau, K_\beta), \beta\} \end{aligned}$$

Moreover, as β has only two possible values, we know that:

$$\{w, L_1, L_2, \mathrm{C}'_{\text{ideal}}(\sigma, \tau, L_1)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{w, M_1, M_2, C'^{A(w),B(w)}(\sigma, \tau, M_1)\} \tag{1}$$

$$\{w, L_1, L_2, \mathrm{C}'_{\text{ideal}}(\sigma, \tau, K_0)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{w, M_1, M_2, C'^{A(w),B(w)}(\sigma, \tau, K_0)\} \tag{2}$$

We will first prove that the real outputs of the honest parties are indistinguishable from ideal outputs, even when the environment Z is present. This is formalized by the following claim:

Claim 2.5 *For every Z, every τ and every σ,*

$$\{w, M_2, M_1, \tau, C^{A(w),B(w),Z(M_1,\tau)}(\sigma)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{w, L_2, L_1, \tau, C^{A(w),B(w),Z(L_1,\tau)}(\sigma)\}$$

where M_i is the output of the ith concluding party in the real execution $C^{A(w),B(w)}$ and L_1, L_2 are defined as follows:

- *if the first party in the real execution $C^{A(w),B(w)}$ accepts, then $L_1 = U_n$. Otherwise, $L_1 = \perp$.*
- *if the second party in the real execution $C^{A(w),B(w)}$ accepts and $L_1 \neq \perp$, then $L_2 = L_1$. If the second party accepts and $L_1 = \perp$, then $L_2 = U_n$. If the second party rejects, then $L_2 = \perp$.*

Proof of claim:

By definition of C' and Equation (1), we know that

$$\begin{aligned} \{w, M_2, M_1, \tau, C^{A(w),B(w),Z(M_1,\tau)}(\sigma)\} &\equiv \{w, M_2, C'^{A(w),B(w)}(\sigma, \tau, M_1)\} \\ &\overset{\mathrm{O}(\gamma)}{\equiv} \{w, L_2, \mathrm{C}'_{\text{ideal}}(\sigma, \tau, L_1)\} \end{aligned}$$

Again, by definition of C', we have

$$\{w, L_2, L_1, \tau, C^{A(w),B(w),Z(L_1,\tau)}(\sigma)\} \equiv \{w, L_2, C'^{A(w),B(w)}(\sigma, \tau, L_1)\}$$

Note that in both the real and ideal models, the string K_0 is distributed identically to L_1, hence by Equation 2, we have

$$\begin{aligned} \{w, C'^{A(w),B(w)}(\sigma, \tau, L_1)\} &\overset{\mathrm{O}(\gamma)}{\equiv} \{w, \mathrm{C}'_{\text{ideal}}(\sigma, \tau, L_1)\} \\ \Rightarrow \{w, L_2, C'^{A(w),B(w)}(\sigma, \tau, L_1)\} &\overset{\mathrm{O}(\gamma)}{\equiv} \{w, L_2, \mathrm{C}'_{\text{ideal}}(\sigma, \tau, L_1)\} \end{aligned}$$

□

We will now prove that if the real outputs of the honest parties are replaced by ideal outputs, then the protocol leaks no information about the password w.

Claim 2.6 *For every Z, every τ and every σ,*

$$\{w, L_2, L_1, \tau, C^{A(w),B(w),Z(L_1,\tau)}(\sigma)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{\tilde{w}, L_2, L_1, \tau, C^{A(w),B(w),Z(L_1,\tau)}(\sigma)\}$$

where $\tilde{w} \overset{R}{\leftarrow} \mathcal{D}$ and L_1, L_2 are defined as follows:

- *if the first party in the real execution $C^{A(w),B(w)}$ accepts, then $L_1 = U_n$. Otherwise, $L_1 = \perp$.*
- *if the second party in the real execution $C^{A(w),B(w)}$ accepts and $L_1 \neq \perp$, then $L_2 = L_1$. If the second party accepts and $L_1 = \perp$, then $L_2 = U_n$. If the second party rejects, then $L_2 = \perp$.*

Proof of claim: We define the real adversary C'' (against security with respect to the session-key challenge) that, on auxiliary input (σ, τ) and receiving L_1 from the first concluding party, simulates $Z(L_1, \tau)$ on its own such that $C''^{A(w),B(w)}(\sigma, \tau, L_1) \equiv \{L_2, L_1, \tau, C^{A(w),B(w),Z(L_1,\tau)}\}$, where L_2 is computed according to the above rule. Since L_1 is distributed identically to K_0, by security with respect to the session-key challenge (for the case where $\beta = 0$) there exists $\mathrm{C}''_{\mathrm{ideal}}$ such that

$$\{w, \mathrm{C}''_{\mathrm{ideal}}(\sigma, \tau, L_1)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{w, C''^{A(w),B(w)}(\sigma, \tau, L_1)\} \tag{3}$$

which in turn implies (by non-uniform indistinguishability or samplability of $\mathcal{D}$)

$$\{\tilde{w}, \mathrm{C}''_{\mathrm{ideal}}(\sigma, \tau, L_1)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{\tilde{w}, C''^{A(w),B(w)}(\sigma, \tau, L_1)\} \tag{4}$$

where $\tilde{w} \overset{R}{\leftarrow} \mathcal{D}$. Note that in the ideal model, the adversary $\mathrm{C}''_{\mathrm{ideal}}(\sigma, \tau, L_1)$ learns nothing about the password w since L_1 is independent of the password w. Hence we have

$$\{\tilde{w}, \mathrm{C}''_{\mathrm{ideal}}(\sigma, \tau, L_1)\} \equiv \{w, \mathrm{C}''_{\mathrm{ideal}}(\sigma, \tau, L_1)\} \tag{5}$$

From Equations (3), (4), (5) and transitivity of indistinguishability, we conclude that

$$\{w, C''^{A(w),B(w)}(\sigma, \tau, L_1)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{\tilde{w}, C''^{A(w),B(w)}(\sigma, \tau, L_1)\}$$

□

Note that the distributions $\{\tilde{w}, L_2, L_1, \tau, C^{A(w),B(w),Z(L_1,\tau)}(\sigma)\}$ and $\{w, L_2, L_1, \tau, C^{A(\tilde{w}),B(\tilde{w}),Z(L_1,\tau)}(\sigma)\}$, where $\tilde{w} \overset{R}{\leftarrow} \mathcal{D}$, are equivalent. Combining Claims 2.5 and 2.6, we obtain

$$\{w, M_1, M_2, \tau, C^{A(w),B(w),Z(M_1,\tau)}(\sigma)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{w, L_1, L_2, \tau, C^{A(\tilde{w}),B(\tilde{w}),Z(L_1,\tau)}(\sigma)\}$$

We now describe the ideal adversary ${\mathrm{C}_{\mathrm{ideal}}}^{Z(L_1,\tau)}(\sigma)$: $\mathrm{C}_{\mathrm{ideal}}$ generates a random password $\tilde{w} \overset{R}{\leftarrow} \mathcal{D}$ and simulates the honest parties A and B in the interaction $(A(\tilde{w}), B(\tilde{w}))$. $\mathrm{C}_{\mathrm{ideal}}$ interacts with $Z(L_1)$ as soon as the first party in the simulated execution concludes. $\mathrm{C}_{\mathrm{ideal}}$ can therefore simulate $C^{A(\tilde{w}),B(\tilde{w}),Z(L_1,\tau)}(\sigma)$.

Hence, for any PPT C, there exists $\mathrm{C}_{\mathrm{ideal}}$ such that for every Z and every τ:

$$\{w, M_1, M_2, Z(M_1, \tau), C^{A(w),B(w),Z(M_1,\tau)}(\sigma)\} \overset{\mathrm{O}(\gamma)}{\equiv} \{w, L_1, L_2, Z(L_1, \tau), {\mathrm{C}_{\mathrm{ideal}}}^{Z(L_1,\tau)}(\sigma)\}$$

■

3 An Overview of the Protocol

Before presenting our protocol, we introduce the polynomial evaluation functionality, which is an important tool for the rest of the paper. In [22], it is observed that a secure protocol for polynomial evaluation immediately yields a protocol for session-key generation which is secure against *passive* adversaries. In [15], Goldreich and Lindell work from the intuition (from [6]) that by augmenting a secure protocol for polynomial evaluation with additional mechanisms, one can obtain a protocol for session-key generation which is secure against *active* adversaries. Our protocol also comes from this intuition but the additional tools we are using are different.

3.1 Secure Polynomial Evaluation

In a secure polynomial evaluation, a party A knows a polynomial Q over some field $\mathbb{F}$ and a party B wishes to learn the value $Q(x)$ for some element $x \in \mathbb{F}$ such that A learns nothing about x and B learns nothing else about the polynomial Q but the value $Q(x)$. More specifically, for our problem, we will assume that $\mathbb{F} = \mathrm{GF}(2^n) \approx \{0,1\}^n$, Q is a non-constant linear polynomial over $\mathbb{F}$, and x is a string in $\{0,1\}^n$.

Definition 3.1 (Polynomial evaluation) *The polynomial evaluation functionality is defined as:*

Inputs *The input of A is a non-constant linear polynomial Q over* $\mathrm{GF}(2^n)$. *The input of B is a value* $x \in \mathrm{GF}(2^n)$.

Outputs *B receives $Q(x)$. A receives nothing.*

As observed in [22], a secure protocol for polynomial evaluation yields immediately a protocol for session-key generation which is secure against passive adversaries as follows: A chooses a random linear non-constant polynomial Q, and A and B engage in a secure polynomial evaluation protocol, where A inputs Q and B inputs w, so that B obtains $Q(w)$. Since A has both Q and w, A can also obtain $Q(w)$, and the session key is set to be $K = Q(w)$.

This protocol is secure against passive adversaries because the key K is a random string (since Q is a random polynomial), and it can be shown that an eavesdropper learns nothing about w or $Q(w)$ (due to the security of the polynomial evaluation).

However, the protocol is not secure against active adversaries. For example, an active adversary C can input a fixed polynomial Q_C in its interaction with B, say the identity polynomial id, and a fixed password w_C in its interaction with A. A outputs the session key $Q_A(w)$ and B outputs the session key $Q_C(w) = w$. With probability $1 - 2^{-n}$, the two session keys are different, whereas the definition of security requires them to be equal with probability $1 - O(\gamma)$.

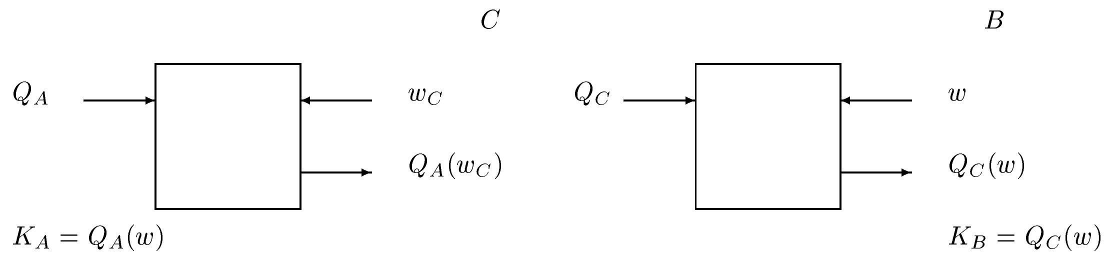

Figure 1: Protocol which is insecure against active adversaries

3.2 Motivation for our Protocol

The main deficiency of the secure polynomial evaluation protocol against active adversaries is that it does not guarantee that A and B output the same random session key. Somehow, the parties have to check that they computed the same random session key before starting to use it. It can be shown that A's session key $K_A = Q_A(w)$ is pseudorandom to the adversary, so A can start using it without leaking information. However, B cannot use its key $K_B = Q_C(w)$ because it might belong to a set of polynomial size (for example, if $Q_C = id$, then $Q_C(w) \in \mathcal{D}$ where the dictionary is by definition a small set). Hence Goldreich and Lindell added a validation phase in which A sends a message to B so that B can check if it computed the same session key, say A sends $f^n(K_A)$ where f is a one-way permutation. Since f^n is a 1-1 map, this uniquely defines K_A (the session-key used now consists of hardcore bits of $f^i(K_A)$, for $i = 0, \cdots, n-1$) : B will compute $f^n(K_B)$ and compare it with the value it received.

But it is still not clear that this candidate protocol is secure. Recall that the security of the polynomial evaluation protocol applies only in the stand-alone setting and guarantees nothing in the concurrent setting. In particular, it might be that C inputs a polynomial Q_C in the polynomial evaluation between C and B such that the polynomials Q_A and Q_C are related in some manner, say for any $w \in \mathcal{D}$, it is easy to compute the correct validation message $f^{2n}(Q_C(w))$ given the value of $f^{2n}(Q_A(w))$; yet B's key does not equal A's key.

To prevent this from happening, Goldreich and Lindell force the polynomial Q input in the polynomial evaluation phase to be consistent with the message sent in the validation phase (which is supposedly $f^{2n}(Q(w))$). The parties have to commit to their inputs at the beginning and then prove in a zero-knowledge manner that the messages sent in the validation phase are consistent with these commitments. Because of the person-in-the-middle attack and the concurrency issues mentioned earlier, Goldreich and Lindell cannot use standard commitment schemes and standard zero-knowledge proofs but rather they use non-malleable commitments and the specific zero-knowledge proof of Richardson and Kilian.

Our approach is to allow C to input a polynomial Q_C related to Q_A, but to prevent C from being able to compute a correct validation message with respect to B's session-key, even given A's validation message. Suppose that the parties have access to a family of pairwise independent hash functions $\mathcal{H}$. In the validation phase, we require A to send $h(f^{2n}(K_A)) = h(f^{2n}(Q_A(w)))$ for some function $h \stackrel{R}{\leftarrow} \mathcal{H}$. Then, even if $K_A = Q_A(w)$ and $K_B = Q_C(w)$ are related (but distinct), the values $h(f^{2n}(K_A))$ and $h(f^{2n}(K_B))$ will be independent and C cannot do much better than randomly guess the value of $h(f^{2n}(K_B))$.

One difficulty arises at this point: the parties have to agree on a common random hash function $h \stackrel{R}{\leftarrow} \mathcal{H}$. But the honest parties A and B only share the randomness coming from the password w so this password w has to be enough to agree on a random hash function. To make this possible, we assume that the password is of the form (w, w') where w and w' are chosen independently of one another: w is chosen at random from an arbitrary dictionary $\mathcal{D} \subseteq \{0,1\}^n$ and w' is uniformly distributed in $\mathcal{D}' = \{0,1\}^{d'}$. (For example, these can be obtained by splitting a single random password from $\{0,1\}^{d''}$ into two parts.) The first part of the password, w, will be used in the polynomial evaluation protocol whereas the second part of the password, w', will be used as the index of a hash function. Indeed, if we assume that $\mathcal{D}' = \{0,1\}^{d'}$, there exists a family of almost pairwise independent hash functions $\mathcal{H} = \{h : \{0,1\}^n \to \{0,1\}^m\}$, where each hash function is indexed by a password $w' \in \mathcal{D}'$ and $m = \Omega(d')$.

We formalize these ideas in the protocol described below.

3.3 Description of the Protocol

Like in [15], we will need a secure protocol for an augmented version of polynomial evaluation.

Definition 3.2 (Augmented polynomial evaluation) *The augmented polynomial evaluation functionality is defined as:*

Earlier phase *A sends a commitment $c_A = \mathrm{Commit}(Q_A, r_A)$ to a linear non-constant polynomial Q_A for a randomly chosen r_A. B receives a commitment c_B. We assume that the commitment scheme used is perfectly binding and computationally hiding.*

Inputs *The input of A is a linear non-constant polynomial Q_A, a commitment c_A to Q_A and a corresponding decommitment r_A. The input of B is a value $x \in \mathrm{GF}(2^n)$ and a commitment c_B.*

Outputs
- *In the case of correct inputs, i.e. $c_A = c_B$ and $c_A = \mathrm{Commit}(Q_A, r_A)$, B receives $Q_A(x)$ and A receives nothing.*
- *In the case of incorrect inputs, i.e. $c_A \neq c_B$ or $c_A \neq \mathrm{Commit}(Q_A, r_A)$, B receives a special failure symbol $\perp$ and A receives nothing.*

The other cryptographic tools we will need are:

Commitment scheme : Let Commit be a perfectly binding, computationally hiding string commitment.

Seed-committed pseudorandom generator : similarly to [15], we will use the seed-committed pseudorandom generator $G(s) = (b(s)b(f(s)) \cdots b(f^{n+\ell-1}(s))f^{n+\ell}(s))$ where f is a one-way permutation with hardcore bit b.

One-time MAC with pseudorandomness property : Let MAC be a message authentication code for message space $\{0,1\}^{p(n)}$ (for a polynomial $p(n)$ to be specified later) using keys of length $\ell = \ell(n)$ that is secure against one query attack, i.e. a PPT A which queries the tagging algorithm MAC_K on *at most one* message of its choice cannot produce a valid forgery on a different message. Additionally, we will require the following pseudorandomness property:

- Let K be a uniform key of length ℓ
- The adversary queries the tagging algorithm MAC_K on the message m of its choice
- The adversary selects $m' \neq m$. We require that the value $\mathrm{MAC}_K(m')$ be pseudorandom with respect to the adversary's view.

Two examples of such a MAC are:

- $\mathrm{MAC}_s(m) = f_s(m)$ where $\{f_s\}_{s\in\{0,1\}^\ell}$ is a pseudorandom function family
- $\mathrm{MAC}_{a,b}(m) = am + b$ where $\ell(n) = 2p(n)$ and $a, b \in \mathrm{GF}(2^{\ell/2})$.

Almost pairwise independent hash functions The family of functions $\mathcal{H} = \{h_{w'} : \{0,1\}^n \to \{0,1\}^m\}_{w'\in\{0,1\}^{d'}}$ is said to be *pairwise δ-dependent* or *almost pairwise independent* if:

1. (uniformity) $\forall x \in \{0,1\}^n$, when we choose $w' \xleftarrow{R} \{0,1\}^{d'}$, $h_{w'}(x)$ is uniform over $\{0,1\}^m$.
2. (pairwise independence) $\forall x_1 \neq x_2 \in \{0,1\}^n, \forall y_1, y_2 \in \{0,1\}^m$, when we choose $w' \xleftarrow{R} \{0,1\}^{d'}$,

$$\Pr_{w'\in\{0,1\}^{d'}}[h_{w'}(x_1) = y_1 \wedge h_{w'}(x_2) = y_2] = \frac{1+\delta}{2^{2m}}$$

We also require that for a fixed $w' \in \{0,1\}^{d'}$, the function $h_{w'}$ is regular, i.e. it is 2^{n-m} to 1. In other words, $h_{w'}(U_n) \equiv U_m$. Throughout this paper, we write $\mu \stackrel{\text{def}}{=} \frac{1+\delta}{2^m}$.

Lemma 3.3 *For the fixed dictionary $\mathcal{D}' = \{0,1\}^{d'} \subseteq \{0,1\}^n$, there exists a family of almost pairwise independent hash functions $\mathcal{H} = \{h_{w'} : \{0,1\}^n \to \{0,1\}^m\}_{w'\in\mathcal{D}'}$ for $\mu = O\left(\frac{n}{|\mathcal{D}'|^{1/3}}\right)$.*

The formal description of the protocol follows.

Protocol 3.4 1. **Inputs** The parties A and B have a joint password (w, w'), where w is chosen at random from an arbitrary dictionary $\mathcal{D} \subseteq \{0,1\}^n$ and w' is uniformly distributed in $\mathcal{D}' = \{0,1\}^{d'} \subseteq \{0,1\}^n$. w and w' are chosen independently.

2. **Commitment:** A chooses a random non-constant linear polynomial Q_A over $\mathrm{GF}(2^n)$ and coin tosses r_A and sends $c_A = \mathrm{Commit}(Q_A, r_A)$. B receives some commitment c_B.

3. **Augmented polynomial evaluation**

   (a) A and B engage in a polynomial evaluation protocol: A inputs the polynomial Q_A, the commitment c_A and the coin tosses r_A it used for the commitment; B inputs the commitment c_B it received and the password w seen as an element of $GF(2^n)$.

   (b) The output of B is denoted Π_B, which is supposed to be equal to $Q_A(w)$.

   (c) A internally computes $\Pi_A = Q_A(w)$.

4. **Validation**

   (a) A sends the string $y_A = h_{w'}(f^{n+\ell}(\Pi_A))$.

(b) Let t_A be the session transcript so far as seen by A. A computes $k_1(\Pi_A) = b(\Pi_A)\cdots b(f^{\ell-1}(\Pi_A))$ and sends the string $z_A = \mathrm{MAC}_{k_1(\Pi_A)}(t_A)$.

5. **Decision**

(a) A always accepts and outputs $k_2(\Pi_A) = b(f^{\ell}(\Pi_A))\cdots b(f^{\ell+n-1}(\Pi_A))$

(b) B accepts (this event is denoted by $dec_B = \texttt{accept}$) if the strings y_B and z_B it received satisfy the following conditions :

- $y_B = h_{w'}(f^{n+\ell}(\Pi_B))$
- $\mathrm{Ver}_{k_1(\Pi_B)}(t_B, z_B) = \texttt{accept}$, where t_B is the session transcript so far as seen by B and $k_1(\Pi_B)$ is defined analogously to $k_1(\Pi_A)$.

If $\Pi_B = \perp$, then B will immediately reject.
If B accepts, it outputs $k_2(\Pi_B) = b(f^{\ell}(\Pi_B))\cdots b(f^{\ell+n-1}(\Pi_B))$.

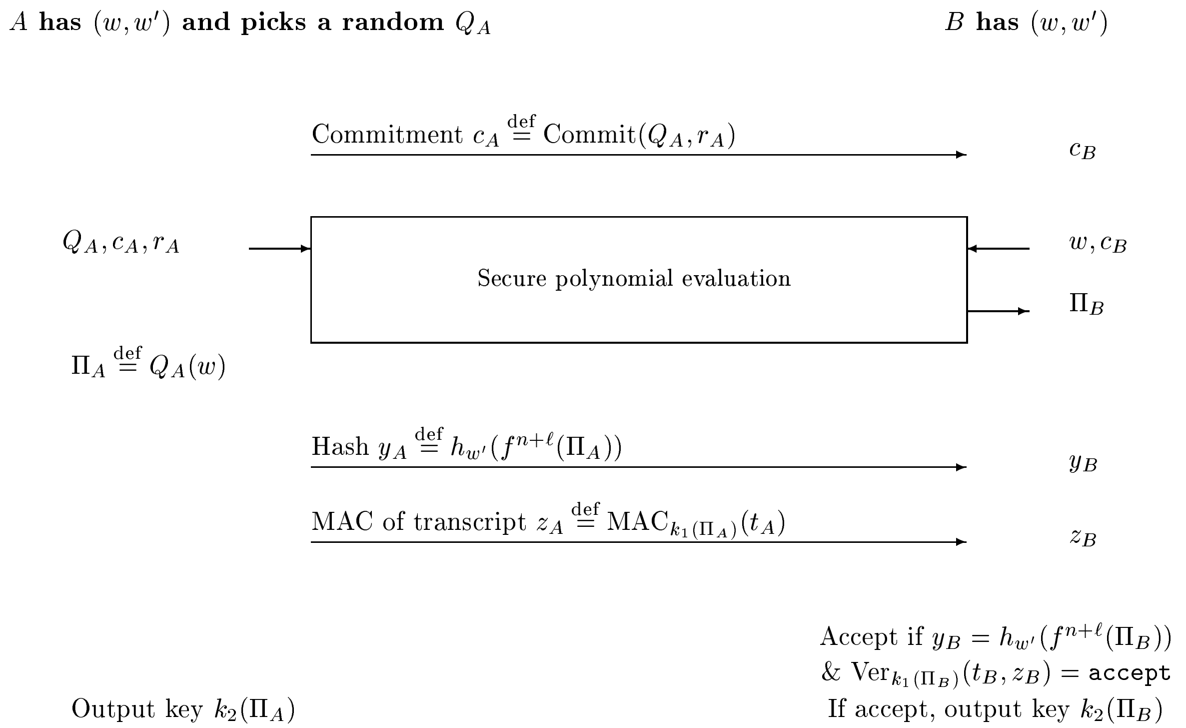

Figure 2: Overview of our protocol

4 Security Theorems

We begin by stating our protocol's security against passive adversaries.

Theorem 4.1 *Protocol 3.4 is secure for the dictionary* $\mathcal{D} \times \mathcal{D}' = \mathcal{D} \times \{0,1\}^{d'}$ *against passive adversaries. More formally, for every passive PPT real adversary C, there exists an ideal adversary* $\mathrm{C}_{\mathrm{ideal}}$ *which always sends* $(\mathrm{dec}_C^A, \mathrm{dec}_C^B) = (1,1)$ *to the trusted party such that for every auxiliary input* $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$*:*

$$\{\mathrm{IDEAL}_{\mathrm{C}_{\mathrm{ideal}}}(\mathcal{D} \times \mathcal{D}', \sigma)\}_\sigma \stackrel{\mathrm{c}}{\equiv} \{\mathrm{REAL}_C(\mathcal{D} \times \mathcal{D}', \sigma)\}_\sigma$$

Next we state the basic security theorem against active adversaries, in the plain model with a dictionary of the form $\mathcal{D} \times \{0,1\}^{d'}$.

Theorem 4.2 *Protocol 3.4 is $(1-\gamma)$-secure with respect to the session-key challenge for the dictionary $\mathcal{D} \times \mathcal{D}' = \mathcal{D} \times \{0,1\}^{d'}$, for $\gamma = \max\left\{\frac{1}{|\mathcal{D}|}, \left(\frac{\mathrm{poly}(n)}{|\mathcal{D}'|}\right)^{\Omega(1)}\right\}$. More precisely, $\gamma = \max\left\{\frac{1}{|\mathcal{D}|}, O\left(\frac{n^3}{|\mathcal{D}'|}\right)^{1/6}\right\}$.*

Finally, we show how the shared dictionary of the form $\mathcal{D} \times \{0,1\}^d$ required in Theorem 4.2 can be realized from several other types of dictionaries $\mathcal{D}''$, achieving security bounds of the form $(\mathrm{poly}(n)/|\mathcal{D}''|)^{\Omega(1)}$ in all cases.

Single Random Password. We can split a single random password from a dictionary $\mathcal{D}'' = \{0,1\}^{d''}$ into two parts, one of length d and one of length $d' = d'' - d$. Optimizing, we set $d = (d'' - 3\log n)/7$, and obtain a security bound of

$$\gamma = \max\left\{\frac{1}{2^d}, O\left(\frac{n^3}{2^{d'}}\right)^{1/6}\right\} = O\left(\frac{n^3}{|\mathcal{D}''|}\right)^{1/7}.$$

Arbitrary Password with Common Random String. We can convert a password from an arbitrary dictionary $\mathcal{D}''' \subseteq \{0,1\}^n$ into a single random password (as in the previous construction) in the common random string model, using randomness extractors, which we define now.

A random variable X is a *k-source* if for all x, $\Pr[X = x] \leq 2^{-k}$. (In other words, X has *min-entropy* at least k.) Note that the uniform distribution on $\mathcal{D}'''$ is a k-source for $k = \log \mathcal{D}'''$.

Definition 4.3 ([24]) *A function* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^\ell \to \{0,1\}^m$ *is a* (strong) (k,ϵ)-extractor *if for every k-source X on $\{0,1\}^n$, the random variable $(U_\ell, \mathrm{Ext}(X, U_\ell))$ is ϵ-close to (U_ℓ, U_m).*

That is, using a random *seed* of length ℓ, the function Ext extracts m almost-uniform bits from the k-source X. We call Ext *explicit* if it is computable in polynomial time (in n and ℓ).

We will use the following construction of "low min-entropy" extractors.

Lemma 4.4 ([27]) *For every n, $k \leq n$, and $\epsilon > 0$, there exists an explicit (k,ϵ)-extractor* $\mathrm{Ext} : \{0,1\}^n \times \{0,1\}^\ell \to \{0,1\}^m$ *with $\ell = O(\log n + m) + 2\log(1/\epsilon)$ and $m = k - 2\log(1/\epsilon) - O(1)$.*

To use extractors with our protocol, we view the common random string as the seed for the extractor, and apply the extractor to convert the password from the arbitrary dictionary $\mathcal{D}''' \subseteq \{0,1\}^n$ into $d'' = m$ almost-uniform bits, which we use in place of the "single random password" in the previous construction. We pay an additive loss of ϵ (the error of the extractor) in the security bound, and also lose because the extractor cannot extract *all* of the min-entropy in the source (i.e. d'' will be smaller than $\log|\mathcal{D}'''|$). Optimizing with the extractor of Lemma 4.4, we set $k = \log|\mathcal{D}'''|$ and $\epsilon = (n^3/|\mathcal{D}'''|)^{1/9}$, and obtain $d'' = m = k - 2\log(1/\epsilon) - O(1)$, i.e. $|\mathcal{D}''| = 2^{d''} = \Omega(\epsilon^2 \cdot 2^k) = \Omega(n^{2/3} \cdot |\mathcal{D}'''|^{7/9})$. Then we have:

$$\gamma = O\left(\frac{n^3}{|\mathcal{D}''|}\right)^{1/7} = O\left(\frac{n^3}{n^{2/3} \cdot |\mathcal{D}'''|^{7/9}}\right)^{1/7} = O\left(\frac{n^3}{|\mathcal{D}'''|}\right)^{1/9},$$

for a final security bound of

$$\gamma + \epsilon = O\left(\frac{n^3}{|\mathcal{D}'''|}\right)^{1/9}.$$

The length of our common random string is $\ell = O(\log n + k) = O(\log n + \log|\mathcal{D}'''|)$. Note that this is only logarithmic in the security parameter n, whereas the protocols of [19, 13] require common reference strings of length polynomially related to n (and, moreover, their reference strings are not merely uniformly distributed random strings, but are supposed to be generated according to more complex distributions). On the other hand, using our protocol requires knowing (or assuming) a lower bound on the size of the dictionary (and this lower bound is what determines the security). The protocols of [15, 19, 13] do not require such a lower bound.

Two Independent Passwords. If the parties share *two* independent passwords w_1, w_2 coming from arbitrary dictionaries $\mathcal{D}_1, \mathcal{D}_2 \subseteq \{0,1\}^r$, then they can apply a (seedless) extractor for 2 independent weak random sources [9] to convert these into a single random password. Even better is to use the following variant of 2-source extractors:

Definition 4.5 ([11]) *A function* $\mathrm{Ble} : \{0,1\}^r \times \{0,1\}^r \to \{0,1\}^m$ *is a* (strong) (k_1, k_2, ϵ)-blender *if for every k_1-source X_1 and independent k_2-source X_2 on $\{0,1\}^r$, the random variable $(X_1, \mathrm{Ble}(X_1, X_2))$ is ϵ-close to (X_1, U_m).*

Thus, if the parties share two independent passwords w_1, w_2 coming from arbitrary dictionaries, a strong blender can be used to convert w_2 into an almost-uniform string $w' = \mathrm{Ble}(w_1, w_2)$ that is essentially independent of the other password, and thus (w_1, w') can be used in our original construction. Nonconstructively, strong (k_1, k_2, ϵ)-blenders are known to exist with $m = k_2 - 2\log(1/\epsilon) - O(1)$, provided that $k_1 > \log r + 2\log(1/\epsilon) + O(1)$. If there were explicit constructions matching these parameters, we would obtain a protocol with security bound of

$$\gamma = O\left(\max\left\{\left(\frac{n}{|\mathcal{D}_1|}\right)^{1/2}, \left(\frac{n^3}{|\mathcal{D}_2|}\right)^{1/8}\right\}\right).$$

Unfortunately, explicit constructions of blenders (or 2-source extractors) are only known in cases when either k_1 or k_2 are at least $r/2$. (See [12] and the references therein for the current state-of-the-art.) Thus we would not obtain a protocol that could work for arbitrary dictionaries $\mathcal{D}_1, \mathcal{D}_2 \subseteq \{0,1\}^n$ of size $\mathrm{poly}(n)$. However, these constructions do allow us to obtain a protocol for arbitrary dictionaries $\mathcal{D}_1, \mathcal{D}_2 \subseteq \{0,1\}^r$ of size, say, $2^{.51r}$, for $r \leq n$ and even $r = O(\log n)$.

5 Overview of the Proof

Notations

- Without loss of generality, we will assume that the real adversary's output equals its view of the execution (since the output is efficiently computable from the view). We will also often omit the auxiliary input σ of the adversary.
- Recall that we denote by $C^{A(w,w'),B(w,w')}$ an execution of C when it communicates with A and B, with common input w. We denote by $C^{A(Q_A,w,w'),B(w,w')}$ the execution of C with A and B where Q_A specifies the random non-constant linear polynomial to be used by A.
- A channel C is *reliable* in a given protocol execution if C runs the (A, C) and (C, B) executions in a synchronized manner and does not modify any message sent by A or B. If C was reliable in the given execution, we denote this event by $\mathrm{reliable}_C = \texttt{true}$; otherwise, we write $\mathrm{reliable}_C = \texttt{false}$.

Although the security theorems are stated in terms of Definition 2.2, the proofs will concern mainly Definition 2.1 and the modifications necessary to take into account the session-key challenge are given.

Similarly to [15], the proof of Theorem 4.2 is in four steps:

1. **Key-Match property:** we show that if $\Pi_A \neq \Pi_B$, then B will reject with probability $1 - O(\gamma)$.
2. **Simulation of the (C, B) interaction:** we show that if the key-match property holds, then the interaction (C, B) can be simulated by an adversary C' interacting only with A, even if the interaction (A, C) is concurrent.
3. **Simulation of the (A, C') interaction:** we show that the interaction (A, C') as a stand-alone can be simulated.
4. Combining the above steps, we obtain a proof of security against active adversaries. The real adversary's view of the concurrent interactions (A, C) and (C, B) can be simulated by a PPT C'' which is non-interactive and can therefore be simulated by an ideal adversary with no input.

Like in [15], the main part of the proof of Theorem 4.2 is the *key-match property*. Once the key-match property is established, we can easily adapt the proofs in [15] to our specific protocol to build an ideal adversary which simulates the real adversary's view.

Theorem 5.1 (Key-Match property) *For every PPT real adversary C and all sufficiently large values of n*

$$\Pr\left[\mathrm{dec}_B = \texttt{accept} \wedge \Pi_A \neq \Pi_B\right] < \sqrt{\mu} + \mu + \epsilon + \mathrm{neg}(n)$$

where $\epsilon = \frac{1}{|\mathcal{D}|}$ *and* $\mu = \frac{1+\delta}{2^m} = O\left(\frac{n}{|\mathcal{D}'|^{1/3}\log|\mathcal{D}'|}\right)$.

The main part of our proof that is new (and simpler than [15]) is the key-match property. As noted in the introduction, the adversary C has total control over the scheduling of the two interactions (A, C) and (C, B). Hence the key-match property will be proved for every possible scheduling case, including those for which these interactions are concurrent. Nevertheless, the key-match property will be established by tools of secure two-party computation, which *a priori* only guarantee security in the stand-alone setting.

For each scheduling, we want to bound from above the probability $\Pr\left[\mathrm{dec}_B = \texttt{accept} \wedge \Pi_A \neq \Pi_B\right]$. Recall that B accepts iff two conditions are satisfied: the string y_B received must equal $h_{w'}(f^{n+\ell}(\Pi_B))$ and the MAC z_B received must be a valid MAC, i.e. $\mathrm{Ver}_{k_1(\Pi_B)}(t_B, z_B) = \texttt{accept}$. Hence, to obtain an upper bound we can omit the verification of the MAC by B and only consider the probability that C succeeds in sending the value $h_{w'}(f^{n+\ell}(\Pi_B))$ when $\Pi_A \neq \Pi_B$. (Like in [15], the MAC is only used to reduce the simulation of active adversaries to the simulation of passive adversaries plus the key-match property.) For convenience, we will decompose the adversary into two algorithms.

- The first algorithm is denoted by C. C is the channel through which A and B communicate. For a given execution, we denote by $C^{A(Q_A,w,w'),B(w,w')}$ the view of C when it communicates with A and B with respective inputs (Q_A, w, w') and (w, w') *until just before C sends a string y_B to B.*
- The second algorithm is denoted by $\mathrm{C_{hash}}$. $\mathrm{C_{hash}}$ takes as an input the above view $C^{A(Q_A,w,w'),B(w,w')}$ and tries to compute the hash value $h_{w'}(f^{n+\ell}(\Pi_B))$.

Hence to establish the key-match property, for each scheduling, we will bound from above the probability

$$\Pr\left[\mathrm{C_{hash}}(C^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B\right]$$

Note that since B always rejects if $\Pi_B = \bot$, we can adopt the convention that

$$\Pr\left[\mathrm{C_{hash}}(C^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_B = \bot\right] = 0$$

We consider two scheduling cases (see Figures 3 and 4):

Scheduling 1 : C sends the commitment c_B to B *after* A sends the hash value y_A.

The intuition for this case is that we have two sequential executions (A, C) and (C, B). Using the security of the polynomial evaluation (A, C), we show that even if C receives y_A, the hash index w' is $(1-\epsilon)$ pseudorandom with respect to the adversary's view. Hence, by the uniformity property of the hash functions, C cannot do much better than randomly guess the value of $h_{w'}(f^{n+\ell}(\Pi_B))$.

Scheduling 2 : C sends the commitment c_B to B *before* A sends the hash value y_A.

The almost pairwise independence property means that for fixed values $x_1 \neq x_2 \in \{0,1\}^n$, if the index w' is chosen at random and independently of x_1 and x_2, then given the value $h_{w'}(x_1)$, one cannot do much better than randomly guess the value $h_{w'}(x_2)$. Before y_A is sent, the hash index w' is random (since it has not been used by A). Thus, if we show that the values Π_A and Π_B can be computed before y_A is sent, then w' is independent of $x_1 = f^{n+\ell}(\Pi_A)$ and $x_2 = f^{n+\ell}(\Pi_B)$ and the adversary cannot guess $h_{w'}(x_2)$ even given $y_A = h_w(x_1)$. To show that Π_A and Π_B can be computed before y_A is sent, we used an *ideal* polynomial evaluation (C, B) to extract an opening of the adversary's commitment c_B. (The adversary must input such an opening in the ideal evaluation, else B will reject).

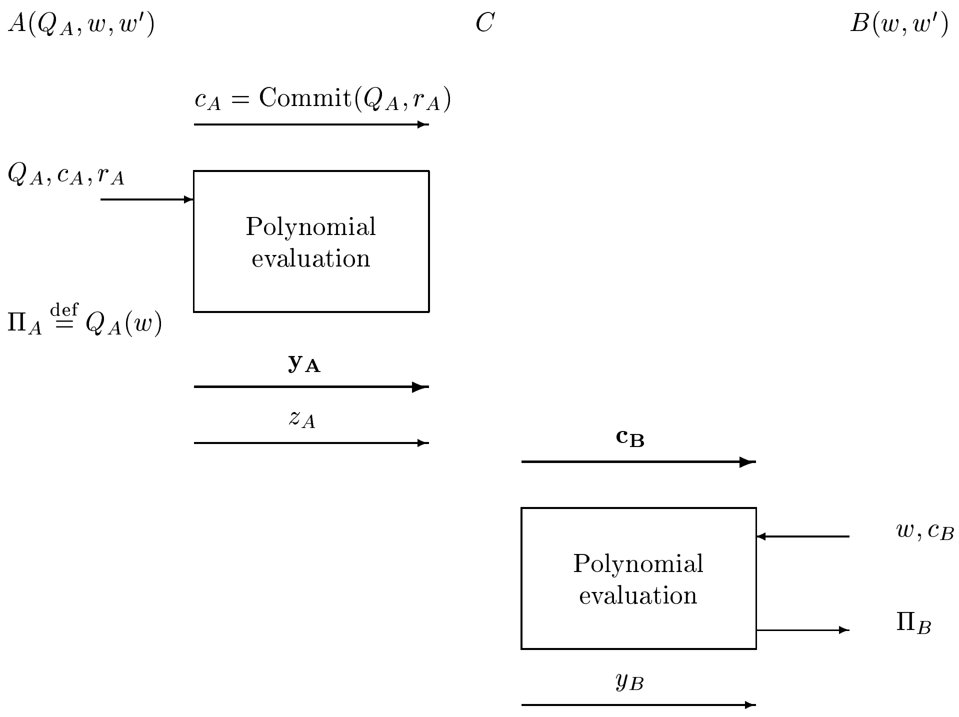

Figure 3: First scheduling

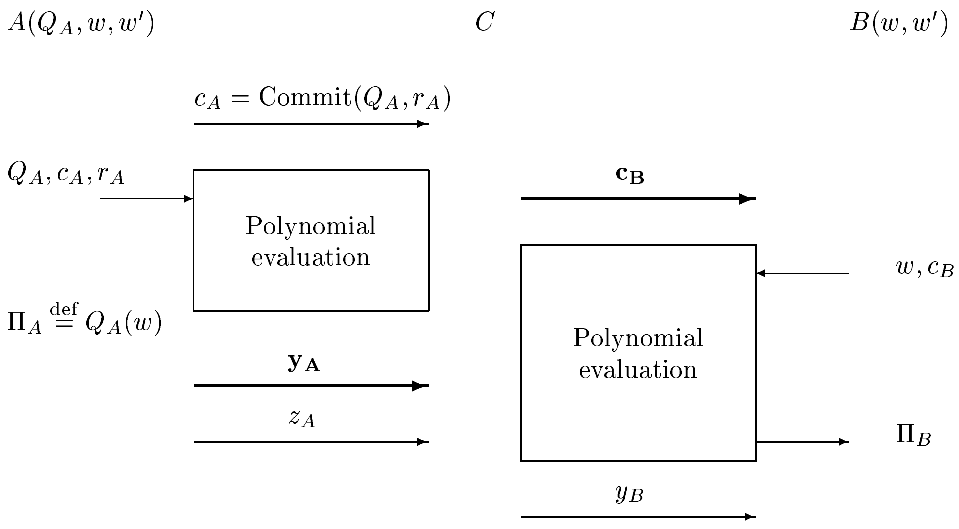

Figure 4: Second scheduling

6 Proof of Security against Passive Adversaries

Theorem 6.1 *Protocol 3.4 is secure for the dictionary $\mathcal{D} \times \mathcal{D}' = \mathcal{D} \times \{0,1\}^{d'}$ against passive adversaries. More formally, for every passive PPT real adversary C, there exists an ideal adversary $\mathrm{C}_{\mathrm{ideal}}$ which always sends $(\mathrm{dec}_C^A, \mathrm{dec}_C^B) = (1,1)$ to the trusted party such that for every auxiliary input $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$:*

$$\{\mathrm{IDEAL}_{\mathrm{C}_{\mathrm{ideal}}}(\mathcal{D} \times \mathcal{D}', \sigma)\}_\sigma \stackrel{c}{\equiv} \{\mathrm{REAL}_C(\mathcal{D} \times \mathcal{D}', \sigma)\}_\sigma$$

Recall that a passive adversary just eavesdrops on the interaction between the honest parties so in this case, the parties A and B output the same session-key ($\mathrm{output}(A) = \mathrm{output}(B)$) and both accept. In the ideal model, the session-key $\mathrm{K}_{\mathrm{ideal}}$ is distributed according to U_n.

Thus, to prove Theorem 6.1, it suffices to prove the following proposition:

Proposition 6.2 *For every passive PPT real adversary C, there exists an ideal adversary $\mathrm{C}_{\mathrm{ideal}}$ such that*

$$\{w, w', \mathrm{output}(A), \mathrm{output}(C^{A(w,w'),B(w,w')})\} \stackrel{c}{\equiv} \{w, w', U_n, \mathrm{output}(\mathrm{C}_{\mathrm{ideal}})\}$$

Proof: The view of the real adversary consists of a transcript of the execution of the protocol by A and B. We can think of this transcript as the concatenation of:

- The commitment to Q_A and the transcript of the augmented polynomial evaluation. We denote these by $T(Q_A, w)$.
- The hash value $y_A \stackrel{\mathrm{def}}{=} h_{w'}(f^{n+\ell}(\Pi_A))$ where $\Pi_A \stackrel{\mathrm{def}}{=} Q_A(w)$.
- The MAC-key $k_1(\Pi_A)$ (it suffices to include the MAC-key rather than the MAC itself, since the latter is easily computable from the MAC-key and the transcript so far).

Claim 6.3

$$\{w, Q_A, T(Q_A, w)\} \stackrel{c}{\equiv} \{w, Q_A, T(\tilde{Q}_A, \tilde{w})\}$$

where Q_A and $\tilde{Q}_A$ are random non-constant linear polynomials and $w, \tilde{w}$ are taken uniformly at random (and independently) from $\mathcal{D}$.

Proof Sketch: The claim follows from the security of the augmented polynomial evaluation.

The commitment scheme we consider is computationally hiding hence a commitment to Q_A is indistinguishable from a commitment to $\tilde{Q}_A$. Combining this with the security of the augmented polynomial evaluation and the connectedness of non-constant linear polynomials (for every Q_A and $\tilde{Q}_A$, there exists $\hat{Q}_A$ and values x_1 and x_2 such that $Q_A(x_1) = \hat{Q}_A(x_1)$ and $\tilde{Q}_A(x_2) = \hat{Q}_A(x_2)$), we know [3] that $\forall w, Q_A, \tilde{w}, \tilde{Q}_A$,

$$T(Q_A, w) \stackrel{c}{\equiv} T(\tilde{Q}_A, \tilde{w})$$

□

Claim 6.3 implies that

$$\begin{aligned} \{w, Q_A(w), T(Q_A, w)\} &\stackrel{c}{\equiv} \{w, Q_A(w), T(\tilde{Q}_A, \tilde{w})\} \\ &\equiv \{w, U_n, T(\tilde{Q}_A, \tilde{w})\} \end{aligned} \tag{6}$$

where Equation (6) comes from the fact that for a random Q_A, $\Pi_A = Q_A(w)$ is uniformly distributed in $\{0,1\}^n$ and Q_A is independent of $T(\tilde{Q}_A, \tilde{w})$.

Note that w' is independent from the variables in Equation (6) hence we have:

$$\{w, w', Q_A(w), T(Q_A, w)\} \stackrel{c}{\equiv} \{w, w', U_n, T(\tilde{Q}_A, \tilde{w})\}$$

[3] See Claim 5.2 in [15]

We can then apply the deterministic polytime function $G(.) = (f^{n+\ell}(.), k_1(.), k_2(.))$ to the third component of each distribution to obtain:

$$\{w, w', k_2(\Pi_A), f^{n+\ell}(\Pi_A), k_1(\Pi_A), T(Q_A, w)\} \stackrel{c}{\equiv} \{w, w', k_2(U_n), f^{n+\ell}(U_n), k_1(U_n), T(\tilde{Q}_A, \tilde{w})\}$$

Since $G(s) = (f^{n+\ell}(s), k_1(s), k_2(s))$ is a PRG, we have:

$$\begin{aligned}
\{w, w', k_2(\Pi_A), f^{n+\ell}(\Pi_A), k_1(\Pi_A), T(Q_A, w)\} &\stackrel{c}{\equiv} \{w, w', U_n^1, U_n^2, U_\ell, T(\tilde{Q}_A, \tilde{w})\} \\
\Rightarrow \{w, w', k_2(\Pi_A), h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A), T(Q_A, w)\} &\stackrel{c}{\equiv} \{w, w', U_n^1, h_{w'}(U_n^2), U_\ell, T(\tilde{Q}_A, \tilde{w})\}
\end{aligned}$$

For a fixed $w' \in \mathcal{D}'$, $h_{w'}$ is a regular map, so

$$\{w, w', k_2(\Pi_A), h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A), T(Q_A, w)\} \stackrel{c}{\equiv} \{w, w', U_n^1, U_m, U_\ell, T(\tilde{Q}_A, \tilde{w})\}$$

The ideal adversary $\mathrm{C}_{\text{ideal}}$ will do the following:

1. Generate a random password $\tilde{w} \in \mathcal{D}$ and a random non-constant linear polynomial $\tilde{Q}_A$
2. Simulate the honest parties in the augmented polynomial evaluation to produce the transcript $T(\tilde{Q}_A, \tilde{w})$
3. Generate random strings U_m and U_ℓ.
4. Output $(U_m, U_\ell, T(\tilde{Q}_A, \tilde{w}))$

■

7 Key-Match Property for the First Scheduling

Scheduling 1 is defined as "C sends the commitment c_B to B *after* A sends y_A". Without loss of generality we can assume that C sends the commitment c_B to B after A sends z_A (since obtaining z_A can only help C).

The intuition for this case is that we have two sequential executions (A, C) and (C, B). Using the security of the polynomial evaluation (A, C), we show that even if C receives y_A, the hash index w' is $(1-\epsilon)$ pseudorandom with respect to the adversary's view. Hence, by the uniformity property of the hash functions, C cannot do much better than randomly guess the value of $h_{w'}(f^{n+\ell}(\Pi_B))$.

Proposition 7.1 *For every PPT real adversary C and all sufficiently large values of n*

$$\Pr[\mathrm{dec}_B = \texttt{accept} \wedge \Pi_A \neq \Pi_B \wedge \mathrm{Sch1}] < \epsilon + \mu + \mathrm{neg}(n)$$

where $\epsilon = \frac{1}{\mathcal{D}}$ and $\mu = O\left(\frac{n}{|\mathcal{D}'|^{1/3} \log |\mathcal{D}'|}\right)$. Sch1 *denotes the event that the execution follows the first scheduling.*

Proof: From the discussion in Section 5, recall that:

$$\Pr[\mathrm{dec}_B = \texttt{accept} \wedge \Pi_A \neq \Pi_B \wedge \mathrm{Sch1}] \leq \Pr\left[\mathrm{C}_{\text{hash}}(C^{A(Q_A, w, w'), B(w, w')}) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B \wedge \mathrm{Sch1}\right]$$

We decompose the adversary into two algorithms:

- C_1 refers to the adversary until just before the commitment c_B is sent. Let (τ, y_A, z_A) denote the view of the adversary C_1 when interacting with $A(Q_A, w, w')$.
- C_2 refers to the adversary once the (A, C) interaction is over, i.e. C_2 will be given as inputs (τ, y_A, z_A). Since C_2 and B are executing the secure (augmented) polynomial evaluation in the stand-alone setting, we know that there exists an ideal adversary $\mathrm{C}_{2,\text{ideal}}$ such that for every τ, y_A, z_A,

$$\{\Pi_{B,\text{ideal}}, {\mathrm{C}_{2,\text{ideal}}}^{B(w, c_B)}(\tau, y_A, z_A)\} \stackrel{c}{\equiv} \{\Pi_B, C_2^{B(w, c_B)}(\tau, y_A, z_A)\}$$

where $\Pi_{B,\text{ideal}} \stackrel{\text{def}}{=} \mathrm{output}(B^{\mathrm{C}_{2,\text{ideal}}(\tau, y_A, z_A)}(w, c_B))$ and $\Pi_B \stackrel{\text{def}}{=} \mathrm{output}(B^{C_2(\tau, y_A, z_A)}(w, c_B))$.

$$
\begin{aligned}
\Pr & \left[\mathrm{C_{hash}}(C^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B \wedge \mathrm{Sch1} \right] \\
& \leq \Pr \left[\mathrm{C_{hash}}(C_2^{B(w,c_B)}(\tau, y_A, z_A)) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \mathrm{Sch1} \right] \\
& \leq \Pr \left[\mathrm{C_{hash}}(\mathrm{C_{2,ideal}}^{B(w,c_B)}(\tau, y_A, z_A)) = h_{w'}(f^{n+\ell}(\Pi_{B,\mathrm{ideal}})) \wedge \mathrm{Sch1} \right] + \mathrm{neg}(n) \qquad (7) \\
& \leq \Pr \left[\mathrm{C_{hash}}'(\tau, h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A)) = h_{w'}(f^{n+\ell}(Q_C(w))) \right] + \mathrm{neg}(n)
\end{aligned}
$$

where $\mathrm{C_{hash}}'$ simulates $\mathrm{C_{2,ideal}}$'s view of the ideal polynomial evaluation with B and Q_C is $\mathrm{C_{2,ideal}}$'s input (wlog we will assume that we are in the correct input case in the augmented polynomial evaluation (C, B) since by convention we define $\mathrm{C_{hash}}(\sigma) \neq h_{w'}(f^{n+\ell}(\Pi_B))$ if $\Pi_B = \perp$). Equation (7) comes from the fact that the security of the augmented polynomial evaluation (C, B) holds even for fixed inputs (w, c_B, τ, y_A, z_A) and with advice string (w', Π_A) to the distinguisher.

We will now prove that the hash index w' is $(1-\epsilon)$-pseudorandom with respect to the inputs given to $\mathrm{C_{hash}}'$. This will imply that the value $h_{w'}(f^{n+\ell}(Q_C(w)))$ is $(1-\epsilon)$-indistinguishable from uniform. Thus $h_{w'}(f^{n+\ell}(Q_C(w)))$ will be predicted by $\mathrm{C_{hash}}'$ with probability at most $\epsilon + 2^{-m}$.

Lemma 7.2 *For every PPT adversary C' interacting with $A(Q_A)$ who halts after the augmented polynomial evaluation,* $\{w, Q_A(w), C'^{A(Q_A)}\} \stackrel{\epsilon}{\equiv} \{w, U_n, C'^{A(Q_A)}\}$

Proof: C' receives a commitment $c_A = \mathrm{Commit}(Q_A, r_A)$ from A before executing the secure protocol for augmented polynomial evaluation. By security of the augmented polynomial evaluation, we know that there exists an ideal adversary $\mathrm{C}'_{\mathrm{ideal}}$ such that for every Q_A, c_A, r_A, we have $C'^{A(Q_A,c_A,r_A)} \stackrel{c}{\equiv} \mathrm{C}'_{\mathrm{ideal}}{}^{A(Q_A,c_A,r_A)}(c_A)$. Without loss of generality, we will assume that we are in the correct input case so that $\mathrm{C}'_{\mathrm{ideal}}$ always receives $Q_A(w_C)$ for some input $w_C = \mathrm{C}'_{\mathrm{ideal}}(c_A)$. Hence for every w, Q_A, c_A, r_A, we have $C'^{A(Q_A,c_A,r_A)} \stackrel{c}{\equiv} \mathrm{C}'_{\mathrm{ideal}}(c_A, w_C, Q_A(w_C))$.

We want to show that

$$\{w, Q_A(w), \mathrm{Commit}(Q_A), w_C, Q_A(w_C)\} \stackrel{\epsilon}{\equiv} \{w, U_n, \mathrm{Commit}(Q_A), w_C, Q_A(w_C)\}$$

where $w_C = \mathrm{C}'_{\mathrm{ideal}}(\mathrm{Commit}(Q_A))$.

- By the hiding property of the commitment scheme, we can replace the commitment to Q_A by a commitment to 0^{2n} in the distributions. This makes $w_C = \mathrm{C}'_{\mathrm{ideal}}(\mathrm{Commit}(0^{2n}))$, which is independent of Q_A.
- Since w_C is independent of w, the probability that $w_C = w$ is at most $\epsilon = \frac{1}{|\mathcal{D}|}$.
- If $w \neq w_C$, $Q_A(w)$ is within 2^{-n} statistical distance of U_n and independent of $Q_A(w_C)$ by pairwise independence of (non-constant linear) polynomials.

$$\{w, Q_A(w), \mathrm{Commit}(0^{2n}), w_C, Q_A(w_C) | w_C \neq w\} \stackrel{c}{\equiv} \{w, U_n, \mathrm{Commit}(0^{2n}), w_C, Q_A(w_C) | w_C \neq w\}$$

■

By Lemma 7.2, we have:

$$\{w, \Pi_A, \tau\} \stackrel{\epsilon}{\equiv} \{w, U_n, \tau\}$$

Note that w' is independent of all the above variables hence we have:

$$\{w, w', \Pi_A, \tau\} \stackrel{\epsilon}{\equiv} \{w, w', U_n, \tau\}$$

We can then apply the deterministic polytime function $(h_{w'}(f^{n+\ell}(.)), k_1(.))$ using the second component w' to the third component of each distribution to obtain:

$$\{w, w', \tau, h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A)\} \stackrel{\epsilon}{\equiv} \{w, w', \tau, h_{w'}(f^{n+\ell}(U_n)), k_1(U_n)\}$$

By applying the polytime function $\mathrm{C}_{2,\mathrm{ideal}}(.)$ to the last three components of each distribution, we have:

$$\{w, w', \tau, h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A), Q_C\} \stackrel{\epsilon}{\equiv} \{w, w', \tau, h_{w'}(f^{n+\ell}(U_n)), k_1(U_n), \tilde{Q}_C\}$$

where $Q_C = \mathrm{C}_{2,\mathrm{ideal}}(\tau, y_A, z_A)$ and $\tilde{Q}_C = \mathrm{C}_{2,\mathrm{ideal}}(\tau, h_{w'}(f^{n+\ell}(U_n)), k_1(U_n))$.

$$\begin{aligned}
&\Pr\left[\mathrm{C_{hash}}'(\tau, h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A), Q_C) = h_{w'}(f^{n+\ell}(Q_C(w)))\right] \\
&\leq \Pr\left[\mathrm{C_{hash}}'(\tau, h_{w'}(f^{n+\ell}(U_n)), k_1(U_n), \tilde{Q}_C) = h_{w'}(f^{n+\ell}(\tilde{Q}_C(w)))\right] + \epsilon + \mathrm{neg}(n) \\
&\leq \Pr\left[\mathrm{C_{hash}}'(\tau, U_m, U_\ell, \tilde{Q}_C) = h_{w'}(f^{n+\ell}(\tilde{Q}_C(w)))\right] + \epsilon + \mathrm{neg}(n) \\
&\leq \epsilon + 2^{-m} + \mathrm{neg}(n)
\end{aligned}$$

The last inequality follows because the inputs to $\mathrm{C_{hash}}'$ are independent of w'. ■

8 Key-Match Property for the warm-up case

Before proving the key-match property for the second scheduling, we will consider a warm-up case for intuition. The warm-up case is defined as "C sends its last message in the (C, B) polynomial evaluation *before* A sends the validation message y_A".

The intuition for this case is that even though the polynomial evaluations are concurrent, the values Π_A and Π_B are defined and computed before $y_A = h_{w'}(f^{n+\ell}(\Pi_B))$ is sent so that we can apply almost pairwise independence. The almost pairwise independence property means that for fixed values $x_1, x_2 \in \{0,1\}^n$, if the index w' is chosen at random and independently of x_1 and x_2, then given the value $h_{w'}(x_1)$, one cannot do much better than randomly guess the value $h_{w'}(x_2)$. Before y_A is sent, the hash index w' is random (since it has not been used by A). Thus, if we show that the values Π_A and Π_B can be computed from the adversary's view before y_A is sent, then w' is independent of $x_1 = f^{n+\ell}(\Pi_A)$ and $x_2 = f^{n+\ell}(\Pi_B)$ and the adversary cannot guess $h_{w'}(x_2)$ from $h_{w'}(x_1)$.

Proposition 8.1 *For every PPT real adversary C and all sufficiently large values of n*

$$\Pr[\mathrm{dec}_B = \mathtt{accept} \wedge \Pi_A \neq \Pi_B] < \mu + \mathrm{neg}(n)$$

From the discussion in Section 5, recall that:

$$\Pr[\mathrm{dec}_B = \mathtt{accept} \wedge \Pi_A \neq \Pi_B] \leq \Pr\left[\mathrm{C_{hash}}(C^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B\right]$$

Let τ denote the adversary's view of the (possibly concurrent) executions (A, C) and (C, B) until $y_A = h_{w'}(f^{n+\ell}(\Pi_A))$ is sent.

Since Q_A represents all of A's input in the protocol up to this point and w represents all of B's input in the protocol up to this point, there exists a PPT C' which on input (Q_A, w) simulates the (possibly concurrent) executions (A, C) and (C, B) on its own until the validation message y_A is sent. In particular, C' computes $\Pi_A = Q_A(w)$ and $\Pi_B = \mathrm{output}(B^C(w))$.

$$C'(Q_A, w) \equiv \{\tau, \Pi_A, \Pi_B\}$$

Note that w' is random and independent of Q_A and w.

Then we have :

$$\begin{aligned}
&\Pr\left[\mathrm{C_{hash}}(C^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B\right] \\
&= \Pr\left[\mathrm{C_{hash}}(\tau, h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A)) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B\right] \\
&\leq \Pr\left[\mathrm{C_{hash}}(C'(Q_A, w), h_{w'}(f^{n+\ell}(\Pi_A))) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B\right] \\
&\leq \mu
\end{aligned}$$

where the last inequality follows from almost pairwise independence, since the index of the hash function, w', is random and independent of the points $f^{n+\ell}(\Pi_A)$ and $f^{n+\ell}(\Pi_B)$.

9 Key-Match Property for the Second Scheduling

Scheduling 2 is defined as " the validation message y_A is sent by A *after* C sends the commitment c_B to B".

The intuition for this case is similar to the warm-up case. The almost pairwise independence property means that for fixed values $x_1, x_2 \in \{0,1\}^n$, if the index w' is chosen at random and independently of x_1 and x_2, then given the value $h_{w'}(x_1)$, one cannot do much better than randomly guess the value $h_{w'}(x_2)$. Before y_A is sent, the hash index w' is random (since it has not been used by A). Thus, if we show that the values Π_A and Π_B can be computed before y_A is sent, then w' is independent of $x_1 = f^{n+\ell}(\Pi_A)$ and $x_2 = f^{n+\ell}(\Pi_B)$ and the adversary cannot guess $h_{w'}(x_2)$. To show that Π_A and Π_B can be computed before y_A is sent, we used an *ideal* polynomial evaluation (C, B) to extract an opening of the adversary's commitment c_B. (The adversary must input such an opening in the ideal evaluation, else B will reject).

9.1 Mental Experiment Protocol

In order to prove the key-match property for this scheduling case, we will need to consider a "mental experiment protocol" which is related to the initial protocol we are analyzing.

Protocol 9.1 (Mental experiment)

1. **Inputs** There are three parties A, B, C_m involved in the protocol. A and B have a joint password $(w, w') \xleftarrow{R} \mathcal{D} \times \mathcal{D}'$. In addition, A is given a random non-constant linear polynomial Q_A.
2. A sends Q_A to C_m.
3. C_m computes $Q_C = C_m(Q_A)$ and sends it to B.
4. B sends w to C.
5. A computes $Q_A(w)$ and sends $y_A = h_{w'}(f^{n+\ell}(Q_A(w)))$ to C. Note that the scheduling "Q_C is sent before y_A" is enforced.
6. C_m sends a string y_B to B.

Proposition 9.2 *In the above mental experiment protocol, for every adversary C_m,*

$$\Pr\left[\mathrm{C_{hash}}(C_m^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(Q_C(w))) \wedge Q_A(w) \neq Q_C(w)\right] \leq \mu$$

Proof: By definition of the mental experiment, $(Q_A(w), Q_C(w))$ can be computed from the view of the adversary C_m *before* $y_A = h_{w'}(f^{n+\ell}(Q_A(w)))$ is sent. Thus the values $(Q_A(w), Q_C(w))$ are independent of the hash index w'. Hence we obtain:

$$\begin{aligned}
&\Pr\left[\mathrm{C_{hash}}(C_m^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(Q_C(w))) \wedge Q_A(w) \neq Q_C(w)\right] \\
\leq\ &\Pr\left[\mathrm{C_{hash}}(C_m(Q_A, Q_C, w, h_{w'}(f^{n+\ell}(Q_A(w))))) = h_{w'}(f^{n+\ell}(Q_C(w))) \wedge Q_A(w) \neq Q_C(w)\right] \\
\leq\ &\mu
\end{aligned}$$

where the last inequality follows from almost pairwise independence (the index of the hash function, w', is random and independent from the points $f^{n+\ell}(Q_A(w))$ and $f^{n+\ell}(Q_C(w))$). ■

9.2 Key-match Property for the Second Scheduling

Proposition 9.3 *For every PPT real adversary C,*

$$\Pr\left[\mathrm{C_{hash}}(C^{A(Q_A,w,w'),B(w,w')} = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B \wedge \mathrm{Sch2}\right] \leq \sqrt{\mu} + \mathrm{neg}(n)$$

This proposition is proved via a reduction to the above mental experiment. We want to show that if an adversary succeeds in computing the correct hash value y_B in the original protocol, then we can build an adversary which computes the correct hash value in the mental experiment (and we know how to upper bound this success probability).

In the mental experiment, the adversary C_m is forced to send (hence commit to) the value Q_C before receiving $y_A = h_{w'}(f^{n+\ell}(\Pi_A))$. This is analogous to forcing the adversary C to open its commitment $c_B = \text{Commit}(Q_C)$ in the original protocol. Roughly speaking, the only difference between the original protocol and the mental experiment is that we force the adversary to open its commitment c_B. Thus, a natural way to build the adversary C_m in the mental experiment is to run the original adversary until the commitment c_B must be opened, find a way to open the commitment, then continue to run the original adversary.

Our main problem is to open the commitment c_B. Intuitively, if the original adversary C can compute the correct hash value y_B, it is because C inputs an opening of c_B in the augmented polynomial evaluation (C, B) (otherwise C would have no idea what Π_B is). However, we cannot formally talk about C's input, since the polynomial evaluation (C, B) does not occur in the stand alone setting. Indeed, C receives from A the validation messages depending on Q_A, w and w' while B_2 occurs. Nevertheless, there exists an ideal adversary C_{ideal} for which the input to the (C, B) polynomial evaluation is well-defined if this ideal adversary is given (Q_A, w, w') as input to simulate the validation messages on its own. The input of $\text{C}_{\text{ideal}}(Q_A, w, w')$ to the augmented polynomial evaluation is well-defined and is intuitively an opening of c_B.

We show that at this stage of the protocol, w and w' are pseudorandom so that given a random password $\tilde{w}$ and a random index $\tilde{w}'$, $\text{C}_{\text{ideal}}(Q_A, \tilde{w}, \tilde{w}')$ will also input to the augmented polynomial evaluation an opening of c_B. So to open the commitment c_B, we will just run the input function of C_{ideal} on a random password and a random index. The details of the proof follow.

Proof: Let us assume that there exists an adversary C in the initial protocol such that

$$\Pr\left[\text{C}_{\text{hash}}(C^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B \wedge \text{Sch2}\right] \geq \eta$$

We will build an adversary C_m in the mental experiment such that

$$\Pr\left[\text{C}_{\text{hash}}(C_m^{A(Q_A,w,w'),B(w,w')}) = h_{w'}(f^{n+\ell}(Q_C(w))) \wedge Q_A(w) \neq Q_C(w)\right] \geq \eta^2 - \text{neg}(n)$$

By Lemma 9.2, we know that the latter probability is bounded by μ which means that $\eta \leq \sqrt{\mu}$.

Once the commitment $c_B = \text{Commit}(Q_C, r_C)$ is sent to B, the state of the three parties is described by $s \stackrel{\text{def}}{=} (Q_A, R_A, R_C, \tau)$ where:

- Q_A is the polynomial used by A
- R_A are the coin tosses of A other than Q_A
- R_C are the coin tosses of C
- τ is C's view up to and including the sending of c_B (this view includes c_A and c_B)

Note that we do not need to include the coin tosses of B since B has not sent any messages yet.

Once the commitment c_B is sent, we can think of an adversary C' as being given $s = (Q_A, R_A, R_C, \tau)$ so that C' simulates the continuation of the augmented polynomial evaluation (A, C) on its own. Given a state s, we define $p_1(s)$ to be the probability that the adversary C' computes the value of the hash function $h_{w'}(f^{n+\ell}(.))$ on a point $\Pi_B \neq \Pi_A$. Formally, we write

$$p_1(s) \stackrel{\text{def}}{=} \Pr\left[\text{C}_{\text{hash}}(C'^{h_{w'}(f^{n+\ell}(\Pi_A)),k_1(\Pi_A),B(w,c_B)}(s) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B | s\right]$$

Note that we have included oracle access to the MAC-key rather than the MAC itself, since the latter is easily computable from the MAC-key and the transcript so far. By definition of p_1 and η, $\mathrm{E}_s[p_1] \geq \eta$.

Consider the adversary $\tilde{C}$ which on input (s, w, w') simulates the hash $h_{w'}(f^{n+\ell}(\Pi_A))$ and the MAC-key $k_1(\Pi_A)$ on its own. We have

$$\Pr\left[\mathrm{C}_{\mathrm{hash}}(\tilde{C}^{B(w,c_B)}(s,w,w') = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_A \neq \Pi_B | s\right] = p_1(s)$$

where the probabilities are taken over the coin tosses of B, which we denote by R_B. Since $(\tilde{C}, B)$ are executing the protocol for augmented polynomial evaluation in the stand-alone setting, there exists an ideal adversary $\tilde{\mathrm{C}}_{\mathrm{ideal}}$ such that the ideal and real distributions will be computationally indistinguishable. Recall that the computational indistinguishability holds even for fixed inputs (s, w, w') and (w, c_B). Hence, for every s, w, w', c_B,

$$(\Pi_{B,\mathrm{ideal}}, \tilde{\mathrm{C}}_{\mathrm{ideal}}^{B(w,c_B)}(s,w,w')) \stackrel{c}{\equiv} (\Pi_B, \tilde{C}^{B(w,c_B)}(s,w,w'))$$

where $\Pi_{B,\mathrm{ideal}} \stackrel{\mathrm{def}}{=} \mathrm{output}(B^{\tilde{\mathrm{C}}_{\mathrm{ideal}}(s,w,w')}(w, c_B))$ and $\Pi_B \stackrel{\mathrm{def}}{=} \mathrm{output}(B^{\tilde{C}(s,w,w')}(w, c_B))$.

Given a state s and $\tilde{\mathrm{C}}_{\mathrm{ideal}}$, we define $p_2(s)$ to be the probability that the ideal adversary $\tilde{\mathrm{C}}_{\mathrm{ideal}}$, given (s, w, w'), opens c_B successfully and computes the value of the hash function $h_{w'}(f^{n+\ell}(.))$ on a point $\Pi_B = Q_C(w) \neq Q_A(w)$. Formally, we write

$$\begin{aligned} p_2(s) \quad \stackrel{\mathrm{def}}{=} \quad & \Pr\left[\mathrm{input}(\tilde{\mathrm{C}}_{\mathrm{ideal}}(s,w,w')) = (Q_C, R) \ s.t.\ c_B = \mathrm{Commit}(Q_C, R)\right. \\ & \left. \wedge\ \mathrm{C}_{\mathrm{hash}}(\tilde{\mathrm{C}}_{\mathrm{ideal}}^{B(w,c_B)}(s,w,w') = h_{w'}(f^{n+\ell}(Q_C(w))) \wedge Q_A(w) \neq Q_C(w) | s\right] \end{aligned}$$

where the probability is taken over the coin tosses of B and the coin tosses of $\tilde{\mathrm{C}}_{\mathrm{ideal}}$, which we denote by R_B and $\mathrm{R}_{\tilde{\mathrm{C}}_{\mathrm{ideal}}}$ respectively. Since w and w' are independent of the state s, the probability $p_2(s)$ is the same if we replace (w, w') by a random pair $(\tilde{w}, \tilde{w}') \in \mathcal{D} \times \mathcal{D}'$.

Intuitively, the adversary C computes the correct hash value $h_{w'}(f^{n+\ell}(\Pi_B))$ because it can open its commitment c_B (otherwise, we are in the case of incorrect inputs and B rejects immediately). This relationship between p_1 and p_2 is formalized in the following lemma.

Lemma 9.4 *For every state $s = (Q_A, R_A, R_C, \tau)$, we have $p_2(s) \geq p_1(s) - \mathrm{neg}(n)$*

Proof:

$$\begin{aligned} p_2(s) \quad = \quad & \Pr\left[\mathrm{input}(\tilde{\mathrm{C}}_{\mathrm{ideal}}(s,w,w')) = (Q_C, R) \ s.t.\ c_B = \mathrm{Commit}(Q_C, R)\right. \\ & \left. \wedge\ \mathrm{C}_{\mathrm{hash}}(\tilde{\mathrm{C}}_{\mathrm{ideal}}^{B(w,c_B)}(s,w,w')) = h_{w'}(f^{n+\ell}(Q_C(w))) \wedge Q_A(w) \neq Q(w) | s\right] \\ = \quad & \Pr\left[\mathrm{C}_{\mathrm{hash}}(\tilde{\mathrm{C}}_{\mathrm{ideal}}^{B(w,c_B)}(s,w,w')) = h_{w'}(f^{n+\ell}(\Pi_{B,\mathrm{ideal}})) \wedge \Pi_A \neq \Pi_{B,\mathrm{ideal}} | s\right] \end{aligned} \tag{8}$$

where we define Π_A to be $Q_A(w)$. Indeed, recall that if $\tilde{\mathrm{C}}_{\mathrm{ideal}}$ does not input an opening to the commitment c_B, then B has no output and

$$\Pr\left[\mathrm{C}_{\mathrm{hash}}(\tilde{\mathrm{C}}_{\mathrm{ideal}}^{B(w,c_B)}(s,w,w')) = h_{w'}(f^{n+\ell}(\Pi_{B,\mathrm{ideal}})) \wedge \Pi_{B,\mathrm{ideal}} = \perp | s\right] = 0$$

By indistinguishability of $\tilde{\mathrm{C}}_{\mathrm{ideal}}^B$ and $\tilde{C}^B$, we have:

Claim 9.5 *For every state s, $|p_2(s) - p_1(s)| \leq \mathrm{neg}(n)$.*

Indeed, recall that for every s, w, w', c_B, $\tilde{\mathrm{C}}_{\mathrm{ideal}}^B$ and $\tilde{C}^B$ are computationally indistinguishable. Hence if we replace $\Pr\left[\mathrm{C}_{\mathrm{hash}}(\tilde{\mathrm{C}}_{\mathrm{ideal}}^{B(w,c_B))}(s,w,w')) = h_{w'}(f^{n+\ell}(\Pi_B)) \wedge \Pi_B \neq \Pi_A | s\right]$ by $p_1(s)$ in Equation (8), there will be only a negligible difference. ■

Assume that there exists an adversary C (described by three algorithms C_1, C' and $\mathrm{C_{hash}}$) in the initial protocol which breaks the key-match property with success probability η. We now describe the adversary C_m in the mental experiment:

1. C_m generates coin tosses R_A, R_B, R_C and $\mathrm{R}_{\tilde{\mathrm{C}}_{\mathrm{ideal}}}$.
2. **Simulating the** (A, C) **polynomial evaluation** Once A sends Q_A to C_m, C_m simulates on its own the beginning of the augmented polynomial evaluation between $A(Q_A, R_A)$ and C_1 and obtains a view τ and a commitment c_B.
3. **Opening algorithm**
   - C_m chooses a random pair $(\tilde{w}, \tilde{w}')$ and runs $\mathrm{input}(\tilde{\mathrm{C}}_{\mathrm{ideal}}(s, \tilde{w}, \tilde{w}'; \mathrm{R}_{\tilde{\mathrm{C}}_{\mathrm{ideal}}})) = (Q_C, R)$, where $s \stackrel{\mathrm{def}}{=} (Q_A, R_A, R_C, \tau)$.
   - If $c_B = \mathrm{Commit}(Q_C, R)$, C_m sends Q_C to B. Otherwise C_m aborts.
4. **Simulating the** (C, B) **polynomial evaluation**
   - C_m receives w from B and $y_A = h_{w'}(f^{n+\ell}(\Pi_A))$ from A.
   - C_m computes $k_1(\Pi_A) = k_1(Q_A(w))$.
   - C simulates $C'^{h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A), B(w, c_B)}$.
5. C_m runs $\mathrm{C_{hash}}(C'^{h_{w'}(f^{n+\ell}(\Pi_A)), k_1(\Pi_A), B(w, c_B)}(s)) = y_B$ and sends y_B to B.

C_m succeeds in this execution of the mental experiment if and only if:

- in this execution of the initial protocol, the adversary C follows scheduling 2.
- the algorithm $\tilde{\mathrm{C}}_{\mathrm{ideal}}$ succeeds in opening the commitment c_B
- $\mathrm{C_{hash}}$ succeeds in sending the correct validation message $h_{w'}(f^{n+\ell}(\Pi_B))$ for $\Pi_B \neq \Pi_A$

More precisely:

$$\begin{aligned}
\Pr&\left[\mathrm{C_{hash}}(C_m^{A(Q_A, w, w'), B(w, w')}) = h_{w'}(f^{n+\ell}(Q_C(w))) \wedge \Pi_A \neq Q_C(w)\right] \\
&= \mathop{\mathrm{E}}_{s}\left[\mathop{\mathrm{E}}_{\tilde{w}, \tilde{w}', \mathrm{R}_{\tilde{\mathrm{C}}_{\mathrm{ideal}}}}[\tilde{\mathrm{C}}_{\mathrm{ideal}}(s, \tilde{w}, \tilde{w}'; \mathrm{R}_{\tilde{\mathrm{C}}_{\mathrm{ideal}}}) \text{ opens } c_B] \mathop{\mathrm{E}}_{w, w', R_B}[\mathrm{C_{hash}} \text{ computes } h_{w'}(f^{n+\ell}(Q_C(w)))]\right] \\
&\geq \mathop{\mathrm{E}}_{s}[p_2(s) \cdot p_1(s)] \\
&\geq \mathop{\mathrm{E}}_{s}\left[(p_1(s))^2\right] - \mathrm{neg}(n) \\
&\geq \eta^2 - \mathrm{neg}(n)
\end{aligned}$$

We built an adversary C_m which succeeds with probability $\eta^2 - \mathrm{neg}(n)$. By Proposition 9.2, we obtain that $\eta \leq \sqrt{\mu} + \mathrm{neg}(n)$. ■

10 Adapting the GL Techniques to our Protocol

Now that we have established the key-match property, we will adapt the proofs of [15] to our protocol for the following steps:

- Simulation of the (C, B) interaction: we show that the interaction (C, B) can be simulated by an adversary C' interacting only with A, even if the interaction (A, C) is concurrent.
- Simulation of the (A, C') interaction: we show that the interaction (A, C') as a stand-alone can be simulated.
- Combining the above steps, we obtain a proof of security against active adversaries.

For the sake of clarity we will first present the simulation of the (A, C') interaction. For each step, the modifications necessary to take into account the session-key challenge and Definition 2.2 are given.

10.1 Simulation of the (A, C') Execution

Proposition 10.1 *For the dictionary $\mathcal{D} \times \mathcal{D}' = \mathcal{D} \times \{0,1\}^{d'}$, for every polytime channel C' interacting with A only, there exists a non-interactive C'' such that for every auxiliary input σ,*

$$\{w, w', k_2(\Pi_A), \mathrm{output}(C'^{A(Q_A,w,w')}(\sigma))\} \overset{\epsilon}{\equiv} \{w, w', U_n, \mathrm{output}(C''(\sigma))\}$$

Proof:

By Lemma 7.2, we know that after the augmented polynomial evaluation, $\{w, \Pi_A\}$ is $(1-\epsilon)$ indistinguishable from $\{w, U_n\}$ with respect to C''s view:

$$\{w, \Pi_A, C'^{A(Q_A)}\} \overset{\epsilon}{\equiv} \{w, U_n, C'^{A(Q_A)}\}$$

Note that w' is independent of all the above variables hence

$$\{w, w', \Pi_A, C'^{A(Q_A)}\} \overset{\epsilon}{\equiv} \{w, w', U_n, C'^{A(Q_A)}\}$$

Lemma 10.2

$\{w, w', k_2(\Pi_A), k_1(\Pi_A), \mathrm{MAC}_{k_1(\Pi_A)}(t_A), h_{w'}(f^{n+\ell}(\Pi_A)), C'^{A(Q_A)}\} \overset{\epsilon}{\equiv} \{w, w', U_n^1, U_\ell, \mathrm{MAC}_{U_\ell}(t_A), U_m, C'^{A(Q_A)}\}$

Proof: From the above discussion, we know that

$$\begin{aligned}
\{w, w', \Pi_A, C'^{A(Q_A)}\} &\overset{\epsilon}{\equiv} \{w, w', U_n, C'^{A(Q_A)}\} \\
\Rightarrow \{w, w', k_2(\Pi_A), k_1(\Pi_A), f^{n+\ell}(\Pi_A), C'^{A(Q_A)}\} &\overset{\epsilon}{\equiv} \{w, w', k_2(U_n), k_1(U_n), f^{n+\ell}(U_n), C'^{A(Q_A)}\} \\
\{w, w', k_2(\Pi_A), k_1(\Pi_A), f^{n+\ell}(\Pi_A), C'^{A(Q_A)}\} &\overset{\epsilon}{\equiv} \{w, w', U_n^1, U_\ell, U_n^2, C'^{A(Q_A)}\} \\
\{w, w', k_2(\Pi_A), \mathrm{MAC}_{k_1(\Pi_A)}(t_A), h_{w'}(f^{n+\ell}(\Pi_A)), C'^{A(Q_A)}\} &\overset{\epsilon}{\equiv} \{w, w', U_n^1, \mathrm{MAC}_{U_\ell}(t_A), U_m, C'^{A(Q_A)}\}
\end{aligned}$$

where t_A is $A(Q_A)$'s transcript of the commitment and the augmented polynomial evaluation, which can be computed from $C'^{A(Q_A)}$. ■

The non-interactive adversary $C''(\sigma)$ will do the following:

1. Generate a random non-constant linear polynomial Q_A.
2. Simulate the interaction between $C'(\sigma)$ and $A(Q_A)$, from which it can compute the transcript t_A.
3. Generate random strings U_ℓ and U_m
4. Output $(C'^{A(Q_A)}, U_m, \mathrm{MAC}_{U_\ell}(t_A))$.

■

Augmented definition We know that

$$\{w, w', k_2(\Pi_A), C'^{A(Q_A,w,w')}(\sigma)\} \overset{\epsilon}{\equiv} \{w, w', U_n, C''(\sigma)\}$$

The session-key challenge is given only after the entire execution (A, C') has been completed (recall that in our protocol A always accepts). The session-key challenge can be generated from each distribution by the distinguisher. We define $C'^{A(Q_A,w,w')}(\sigma, K_\beta) \overset{\mathrm{def}}{=} (C'^{A(Q_A,w,w')}(\sigma), K_\beta)$ and $C''(\sigma, K_\beta) \overset{\mathrm{def}}{=} (C''(\sigma), K_\beta)$. By the above discussion we have:

$$\{w, w', k_2(\Pi_A), C'^{A(Q_A,w,w')}(\sigma, K_\beta), \beta\} \overset{\epsilon}{\equiv} \{w, w', U_n, C''(\sigma, K_\beta), \beta\}$$

where on the left-hand side K_β is given when A concludes and is defined as:

$$K_\beta = \begin{cases} k_2(\Pi_A) & \text{if } \beta = 1 \\ U_n' & \text{if } \beta = 0 \end{cases}$$

and on the right-hand side K_β is defined as

$$K_\beta = \begin{cases} U_n & \text{if } \beta = 1 \\ U_n' & \text{if } \beta = 0 \end{cases}$$

10.2 Simulation of the (C, B) Execution

Proposition 10.3 *For the dictionary $\mathcal{D} \times \mathcal{D}' = \mathcal{D} \times \{0,1\}^{d'}$, for every real adversary C interacting with A and B, there exists a PPT C' interacting only with A such that for every auxiliary input $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$*

$$\{w, w', k_2(\Pi_A), \mathrm{output}(C'^{A(Q_A,w,w')}(\sigma))\} \stackrel{\epsilon+\eta}{\equiv} \{w, w', k_2(\Pi_A), \mathrm{output}(C^{A(Q_A,w,w'),B(w,w')}(\sigma))\}$$

where $\eta = \sqrt{\mu} + \mu + \epsilon$.

The proof of this proposition relies on two facts:

- it is easy to simulate B in the augmented polynomial evaluation by security of two-party computation (see Lemma 10.4)
- B's decision bit can be simulated with high probability because of the key-match property (see Lemma 10.6). We need for C' to simulate B's decision bit because the view of the real adversary $C^{A(Q_A,w,w'),B(w,w')}$ includes B's decision bit.

Lemma 10.4 *Let $\tilde{C}$ be a real adversary interacting with A and a modified party $\mathrm{B}_{\not{d}\mathrm{ec}}$ ($\mathrm{B}_{\not{d}\mathrm{ec}}$ does the same as B except that it does not output a decision bit). There exists C' interacting only with A such that:*

$$\{w, w', k_2(\Pi_A), \mathrm{output}(C'^{A(Q_A,w,w')}(\sigma))\} \stackrel{c}{\equiv} \{w, w', k_2(\Pi_A), \mathrm{output}(\tilde{C}^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma))\}$$

where on the left-hand side $k_2(\Pi_A)$ refers to the output of A in the execution $C'^{A(Q_A,w,w')}(\sigma)$ and on the right-hand side $k_2(\Pi_A)$ refers to the output of A in the execution $\tilde{C}^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma)$.

Proof: Note that these distributions do not refer to $\mathrm{B}_{\not{d}\mathrm{ec}}$'s output from the polynomial evaluation, hence we can switch $\mathrm{B}_{\not{d}\mathrm{ec}}$'s input from w to a random password $\tilde{w} \in \mathcal{D}$ via the following claim.

Claim 10.5 *For every $w, w', Q_A, \tilde{w}$ and auxiliary input $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$,*

$$\{\mathrm{output}(A), C^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma)\} \stackrel{c}{\equiv} \{\mathrm{output}(A), C^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(\tilde{w})}(\sigma)\}$$

where on the left-hand side output(A) *refers to the output of A in the execution $C^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma)$ and on the right-hand side* output(A) *refers to the output of A in the execution $C^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(\tilde{w})}(\sigma)$.*

Proof of claim: Define C' which on input (w, w', Q_A) simulates the entire (A, C) execution, including computing output(A), on its own:

$$\begin{aligned} C'^{\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(w, w', Q_A, \sigma) &\equiv \{\mathrm{output}(A), C^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma)\} \\ C'^{\mathrm{B}_{\not{d}\mathrm{ec}}(\tilde{w})}(w, w', Q_A, \sigma) &\equiv \{\mathrm{output}(A), C^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(\tilde{w})}(\sigma)\} \end{aligned}$$

Since C' and $\mathrm{B}_{\not{d}\mathrm{ec}}$ are executing the secure polynomial evaluation protocol in the stand-alone setting, there exists an ideal adversary $\mathrm{C}'_{\mathrm{ideal}}$ such that for every $w, w', Q_A, \tilde{w}, \sigma$,

$$\begin{aligned} \mathrm{C}'_{\mathrm{ideal}}(w, w', Q_A, \sigma) &\stackrel{c}{\equiv} C'^{\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(w, w', Q_A, \sigma) \\ \mathrm{C}'_{\mathrm{ideal}}(w, w', Q_A, \sigma) &\stackrel{c}{\equiv} C'^{\mathrm{B}_{\not{d}\mathrm{ec}}(\tilde{w})}(w, w', Q_A, \sigma) \end{aligned}$$

By transitivity of indistinguishability, we obtain the lemma. □

By Claim 10.5, we have:

$$\{w, w', \Pi_A, \mathrm{output}(\tilde{C}^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(\tilde{w})}(\sigma))\} \stackrel{c}{\equiv} \{w, w', \Pi_A, \mathrm{output}(\tilde{C}^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma))\}$$

Hence for any adversary $\tilde{C}$ interacting with A and $\mathrm{B}_{\not{d}\mathrm{ec}}$, we build an adversary C' which will simulate on its own $\mathrm{B}_{\not{d}\mathrm{ec}}$ by using an arbitrary element $\tilde{w}$ for the polynomial evaluation (C, B). ■

Augmented definition In the case of the augmented definition, the proof of the first lemma still holds because Claim 10.5 will hold for every session-key challenge given by A. Hence we have

$$\{w, w', k_2(\Pi_A), \mathrm{output}(C'^{A(Q_A,w,w')}(\sigma, K_\beta)), \beta\} \stackrel{c}{\equiv} \{w, w', k_2(\Pi_A), \mathrm{output}(\tilde{C}^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma, K_\beta)), \beta\}$$

where K_β is given when A concludes and is defined as:

$$K_\beta = \begin{cases} k_2(\Pi_A) & \text{if } \beta = 1 \\ U'_n & \text{if } \beta = 0 \end{cases}$$

Lemma 10.6 *Let C be a real adversary interacting with A and B. There exists $\tilde{C}$ interacting with A and $\mathrm{B}_{\not{d}\mathrm{ec}}$ such that*

$$\{w, w', k_2(\Pi_A), \mathrm{output}(\tilde{C}^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma))\} \stackrel{\epsilon+\eta}{\equiv} \{w, w', k_2(\Pi_A), \mathrm{output}(C^{A(Q_A,w,w'),B(w,w')}(\sigma))\}$$

The proof of this lemma relies on the fact that the decision bit of B can be predicted by C with high probability because of the key-match property and the following claim.

Claim 10.7 *For every C interacting with A and $\mathrm{B}_{\not{d}\mathrm{ec}}$, the probability that $t_B \neq t_A$ and C computes $\mathrm{MAC}_{k_1(\Pi_A)}(t_B)$ is at most $\epsilon + \mathrm{neg}(n)$.*

Proof of claim: First, we will remove B by modifying C into C' from Lemma 10.4, which simulates B in the polynomial evaluation phase. We know from Lemma 10.2 that:

$$\begin{aligned} &\{w, w', k_2(\Pi_A), k_1(\Pi_A), \mathrm{MAC}_{k_1(\Pi_A)}(t_A), h_{w'}(f^{n+\ell}(\Pi_A)), C'^{A(Q_A)}\} \\ \stackrel{\epsilon}{\equiv} \quad &\{w, w', U_n^1, U_\ell, \mathrm{MAC}_{U_\ell}(t_A), U_m, C'^{A(Q_A)}\} \end{aligned}$$

We will bound from above, for $t \neq t_A$, the probability

$$\begin{aligned} &\Pr\left[\mathrm{C_{mac}}(C'^{A(Q_A)}, h_{w'}(f^{n+\ell}(\Pi_A)), \mathrm{MAC}_{k_1(\Pi_A)}(t_A)) = \mathrm{MAC}_{k_1(\Pi_A)}(t)\right] \\ \leq \quad &\Pr\left[\mathrm{C_{mac}}(C'^{A(Q_A)}, U_m, \mathrm{MAC}_{U_\ell}(t_A)) = \mathrm{MAC}_{U_\ell}(t)\right] + \epsilon + \mathrm{neg}(n) \\ \leq \quad &\epsilon + \mathrm{neg}(n) \end{aligned}$$

where the last inequality comes from the one-time MAC property. □

Using Claim 10.7, we obtain the following adversary $\tilde{C}$: $\tilde{C}$ interacts with A and $\mathrm{B}_{\not{d}\mathrm{ec}}$ by passing their messages to C. Since $\tilde{C}$ has the transcript of the interactions (A, C) and (C, B), $\tilde{C}$ can tell whether C was reliable or not. If C was reliable, $\tilde{C}$ predicts that $\mathrm{dec}_B = \texttt{accept}$ (since B always accepts if C is reliable), otherwise, it predicts $\mathrm{dec}_B = \texttt{reject}$. We know that $\Pr\left[\tilde{C} \text{ predicts incorrectly}\right] = \Pr[\mathrm{dec}_B = \texttt{accept} \wedge \mathit{reliable}_C = \texttt{false}]$. In order to prove Lemma 10.6, it remains to show that for any C,

$$\Pr[\mathrm{dec}_B = \texttt{accept} \wedge \mathit{reliable}_C = \texttt{false}] < \epsilon + \eta + \mathrm{neg}(n)$$

$$\begin{aligned} &\Pr[\mathrm{dec}_B = \texttt{accept} \wedge \mathit{reliable}_C = \texttt{false}] \\ &\quad = \quad \Pr[\mathrm{dec}_B = \texttt{accept} \wedge \mathit{reliable}_C = \texttt{false} \wedge \Pi_A \neq \Pi_B] \\ &\qquad\quad + \Pr[\mathrm{dec}_B = \texttt{accept} \wedge \mathit{reliable}_C = \texttt{false} \wedge \Pi_A = \Pi_B] \\ &\quad \leq \quad (\eta + \mathrm{neg}(n)) \qquad \text{from the key-match property} \\ &\qquad\quad +(\epsilon + \mathrm{neg}(n)) \qquad \text{from Claim 10.7} \\ &\quad \leq \quad \epsilon + \eta + \mathrm{neg}(n) \end{aligned}$$

Augmented definition

C is not reliable and B concludes first: $\tilde{C}$ will set B's simulated decision bit to be $\mathrm{dec}_B = \texttt{reject}$ and its simulated session-key challenge to be $\perp$. Note that if B concludes first, then with high probability B would indeed reject (which follows from the fact that if C is not reliable, then with high probability B will reject as shown above).

A concludes first: Lemma 10.6 must be slightly modified. One can show using Lemma 10.2 that the probability that $t_B \neq t_A$ and C computes $\mathrm{MAC}_{k_1(\Pi_A)}(t_B)$ is at most $\epsilon + \mathrm{neg}(n)$ even if $k_2(\Pi_A) = k_A$ is given.

From the above two arguments, we have:

$$\{w, w', k_2(\Pi_A), \mathrm{output}(\tilde{C}^{A(Q_A,w,w'),\mathrm{B}_{\not{d}\mathrm{ec}}(w)}(\sigma, K_\beta)), \beta\} \overset{\epsilon+\eta}{\equiv} \{w, w', k_2(\Pi_A), \mathrm{output}(C^{A(Q_A,w,w'),B(w,w')}(\sigma, K_\beta)), \beta\}$$

where on left-hand side K_β is given when A concludes and is defined as:

$$K_\beta = \begin{cases} k_2(\Pi_A) & \text{if } \beta = 1 \\ U_n' & \text{if } \beta = 0 \end{cases}$$

and on the right-hand side the session-key challenge K_β is given once the first party (either A or B) concludes with output L_1 and is defined as:

$$K_\beta = \begin{cases} L_1 & \text{if } \beta = 1 \text{ or } L_1 = \perp \\ U_n' & \text{if } \beta = 0 \text{ and } L_1 \neq \perp \end{cases}$$

10.3 Security Theorem

Theorem 10.8 *For the dictionary $\mathcal{D} \times \mathcal{D}' = \mathcal{D} \times \{0,1\}^{d'}$, for every PPT real adversary C, there exists a polytime ideal model channel $\hat{C}$ such that for any $\sigma \in \{0,1\}^{\mathrm{poly}(n)}$*

$$\{\mathrm{IDEAL}_{\hat{C}}(\mathcal{D}, \sigma)\} \overset{3\epsilon+2\eta}{\equiv} \{\mathrm{REAL}_C(\mathcal{D}, \sigma)\}$$

where $\eta = \sqrt{\mu} + \mu + \epsilon$.

Proof: From the previous two sections, we know that there exists a non-interactive C'' such that

$$\{w, w', U_n, C''(\sigma)\} \overset{2\epsilon+\eta}{\equiv} \{w, w', k_2(\Pi_A), \mathrm{output}(C^{A,B}(\sigma))\}$$

The ideal model adversary $\hat{C}$ does the following:

- $\hat{C}$ decides that A will conclude first and accept in the ideal model.
- C invokes C'', which is non-interactive.
- According to the view output by C'', $\hat{C}$ will decide whether B accepts or not in the ideal execution.
- $\hat{C}$ outputs the output of C''.

$$\Rightarrow \{w, w', U_n, \hat{C}(\sigma)\} \overset{2\epsilon+\eta}{\equiv} \{w, w', k_2(\Pi_A), \mathrm{output}(C^{A,B}(\sigma))\} \tag{9}$$

We now need to include B's output in the above distributions. Let D be a distinguisher for $\mathrm{IDEAL}_{\hat{C}}$ and REAL_C. We will consider the different cases, whether B accepts or not.

If B rejects

$$\begin{aligned}\Pr\left[D(\text{IDEAL}_{\hat{C}}) = 1 \wedge \text{dec}_B = \texttt{reject}\right] &= \Pr\left[D(w, w', U_n, \perp, \hat{C}) = 1 \wedge \text{dec}_B = \texttt{reject}\right] \\ \Pr\left[D(\text{REAL}_C) = 1 \wedge \text{dec}_B = \texttt{reject}\right] &= \Pr\left[D(w, w', k_2(\Pi_A), \perp, C^{A,B}) = 1 \wedge \text{dec}_B = \texttt{reject}\right]\end{aligned}$$

In the ideal model, we are guaranteed that when B rejects, $\hat{C}$ will send $b = 0$ to the trusted party, causing B to output $\perp$. In the real protocol, when B rejects, it always outputs $\perp$. But B's decision bit is contained in the view $\hat{C}$ (as simulated by C'') and in the view $C^{A,B}$ so by Equation (9) the difference between $\Pr\left[D(\text{IDEAL}_{\hat{C}}) = 1 \wedge \text{dec}_B = \texttt{reject}\right]$ and $\Pr\left[D(\text{REAL}_C) = 1 \wedge \text{dec}_B = \texttt{reject}\right]$ is at most $2\epsilon + \eta + \text{neg}(n)$.

If B accepts

- suppose C was reliable: in the real model, B always accepts and outputs $k_2(\Pi_A)$; in the ideal model, B outputs U_n. C is acting like a passive adversary, so we know that $\text{IDEAL}_{\hat{C}} \overset{c}{\equiv} \text{REAL}_C$.
- suppose C was not reliable, but B accepts. From the proof of Theorem 10.3, we know that $\Pr\left[\text{dec}_B = \texttt{accept} \wedge \textit{reliable}_C = \texttt{false}\right] \leq \epsilon + \eta + \text{neg}(n)$, whether in the real model or in the one simulated by $\hat{C}$.

$$\begin{aligned} & \left|\Pr\left[D(\text{IDEAL}_{\hat{C}}) = 1 \wedge \text{dec}_B = \texttt{accept} \wedge \textit{reliable}_C = \texttt{false}\right]\right. \\ - & \left.\Pr\left[D(\text{REAL}_C) = 1 \wedge \text{dec}_B = \texttt{accept} \wedge \textit{reliable}_C = \texttt{false}\right]\right| \\ \leq & \quad \epsilon + \eta + \text{neg}(n)\end{aligned}$$

Combining all the above cases, we have that the ideal distribution and the real distribution are distinguishable with probability at most $3\epsilon + 2\eta$. ■

Augmented definition From the previous sections, we know that

$$\{w, w', U_n, C''(\sigma, K_\beta), \beta\} \overset{2\epsilon+\eta}{\equiv} \{w, w', k_2(\Pi_A), C^{A,B}(\sigma, K_\beta), \beta\}$$

where on the left-hand side K_β is defined as

$$K_\beta = \begin{cases} U_n & \text{if } \beta = 1 \\ U'_n & \text{if } \beta = 0 \end{cases}$$

and on the right-hand side the session-key challenge K_β is given once the first party (either A or B) concludes with output L_1:

$$K_\beta = \begin{cases} L_1 & \text{if } \beta = 1 \text{ or } L_1 = \perp \\ U_n & \text{if } \beta = 0 \text{ and } L_1 \neq \perp \end{cases}$$

The ideal adversary $\hat{C}$ does the following:

- $\hat{C}$ decides that A will conclude first and accept. The trusted party chooses $\beta \overset{R}{\leftarrow} \{0,1\}$ and gives $\hat{C}$ the string K_β where
$$K_\beta = \begin{cases} U_n & \text{if } \beta = 1 \\ U'_n & \text{if } \beta = 0 \end{cases}$$
- C invokes $C''(\sigma, K_\beta)$, which is non-interactive.
- According to the view output by C'', $\hat{C}$ will decide whether B accepts or not in the ideal execution.
- $\hat{C}$ outputs the output of $C''(\sigma, K_\beta)$.

Acknowledgments

We thank Oded Goldreich and Yehuda Lindell for their encouragement and an inspiring discussion which led to a substantial simplification of our protocol. We are also grateful to Mihir Bellare for pointing out the extension of our protocol to arbitrary dictionaries in the common random string model.

References

[1] Barak, B.: Constant-Round Coin-Tossing With a Man in the Middle or Realizing the Shared Random String Model. IEEE Symposium on Foundations of Computer Science (2002) 345–355

[2] Bellare, M., Pointcheval, D., Rogaway, P.: Authenticated Key Exchange Secure against Dictionary Attacks. Advances in Cryptology - Eurocrypt 2000 Proceedings, Lecture Notes in Computer Science **1807** (2000) 139–155

[3] Bellare, M., Rogaway, P.: Entity Authentication and Key Distribution. Advances in Cryptology - Crypto 93 Proceedings, Lecture Notes in Computer Science **773** (1994) 232–249

[4] Bellovin, S., Merritt, M.: Encrypted Key Exchange: Password-Based Protocols Secure Against Dictionary Attacks. ACM/IEEE Symposium on Research in Security and Privacy (1992) 72–84

[5] Bellovin, S., Merritt, M.: Augmented Encrypted Key Exchange: A Password-Based Protocol Secure against Dictionary Attacks and Password File Compromise. ACM Conference on Computer and Communications Security (1993) 244–250

[6] Boyarsky, M.: Public-Key Cryptography and Password Protocols: The Multi-User Case. ACM Conference on Computer and Communications Security (1999) 63–72

[7] Boyko, V., MacKenzie, P., Patel, S.: Provably Secure Password-Authenticated Key Exchange Using Diffie-Hellman. Advances in Cryptology - Eurocrypt 2000 Proceedings, Lecture Notes in Computer Science **1807** (2000) 156–171

[8] Canetti, R.: Universally Composable Security: A New Paradigm for Cryptographic Protocols. IEEE Symposium on Foundations of Computer Science (2001) 136–145

[9] Chor, B., Goldreich, O.: Unbiased Bits from Sources of Weak Randomness and Probabilistic Communication Complexity. SIAM Journal on Computing **17:2** (1988) 230–261

[10] Diffie, W., Hellman, M.: New Directions in Cryptography. IEEE Transactions on Information Theory **22:6** (1976) 644–654

[11] Dodis, Y., Oliveira, R.: On Extracting Private Randomness over a Public Channel. Approximation, Randomization, and Combinatorial Optimization. Proc. of APPROX 2003 and RANDOM 2003, Lecture Notes in Computer Science **2764** (2003) 252–263

[12] Dodis, Y., Elbaz, A., Raz, R., Oliveira, R.: Improved Randomness Extraction from Two Independent Sources. Approximation, Randomization, and Combinatorial Optimization. Proc. of APPROX 2004 and RANDOM 2004, Lecture Notes in Computer Science (2004) To appear.

[13] Gennaro, R., Lindell, Y.: A Framework for Password-Based Authenticated Key Exchange. Advances in Cryptology - Eurocrypt 2003 Proceedings, Lecture Notes in Computer Science **2656** (2003) 524–543

[14] Goldreich, O.: Foundations of Cryptography, Volume 2. Cambridge University Press (2004)

[15] Goldreich, O., Lindell, Y.: Session-Key Generation Using Human Passwords Only. Advances in Cryptology - Crypto 2001 Proceedings, Lecture Notes in Computer Science **2139** (2001) 408–432

[16] Goldreich, O., Wigderson, A.: Tiny Families of Functions with Random Properties: A Quality-Size Trade-off for Hashing. Random Structures and Algorithms **11:4** (1997) 315–343

[17] Goldwasser, S., Micali, S.: Probabilistic Encryption. Journal of Computer and System Sciences **28:2** (1984) 270–299

[18] Halevi, S., Krawczyk, H.: Public-Key Cryptography and Password Protocols. ACM Conference on Computer and Communications Security (1998) 122–131

[19] Katz, J., Ostrovsky, R., Yung, M.: Efficient Password-Authenticated Key Exchange Using Human-Memorable Passwords. Advances in Cryptology - Eurocrypt 2001 Proceedings, Lecture Notes in Computer Science **2045** (2001) 475–494

[20] Kobara, K., Imai, H.: Pretty-Simple Password-Authenticated Key-Exchange Under Standard Assumptions. IECIE Trans. **E85-A:10** (2002) 2229–2237

[21] Micali, S., Rackoff, C., Sloan, B.: The Notion of Security for Probabilistic Cryptosystems. SIAM Journal on Computing **17** (1988) 412–426

[22] Naor, M., Pinkas, B.: Oblivious Transfer and Polynomial Evaluation. ACM Symposium on Theory of Computing (1999) 245–254

[23] Nguyen, M.-H., Vadhan, S.: Simpler Session-Key Generation from Short Random Passwords. Proceedings of the First Theory of Cryptography Conference (TCC '04), Lecture Notes in Computer Science **2951** (2004) 428–445

[24] Nisan, N., Zuckerman, D.: Randomness is Linear in Space. Journal of Computer and System Sciences **52:1** (1996) 43–52

[25] Richardson, R., Kilian, J.: On the Concurrent Composition of Zero-Knowledge Proofs. Advances in Cryptology - Eurocrypt 99 Proceedings, Lecture Notes in Computer Science **1592** (1999) 415–431

[26] Shoup, V.: On Formal Models for Secure Key Exchange. Cryptology ePrint Archive (1999) Report 1999/012

[27] Srinivasan, A., Zuckerman, D.: Computing with Very Weak Random Sources. SIAM Journal on Computing **28:4** (1999) 1453–1459

[28] Steiner, M., Tsudik, G., Waidner, M.: Refinement and Extension of Encrypted Key Exchange. Operating Systems Review **29:3** (1995) 22–30

[29] Yao, A.: How to Generate and Exchange Secrets. IEEE Symposium on Foundations of Computer Science (1986) 162–167

A Secure Two-Party Computation

This presentation is taken from [14]. We will describe secure two-party computation for the special case of single-output functionalities, i.e. functionalities where only one party obtains an output. Indeed, we will only use tools from secure two-party computation when dealing with the augmented polynomial evaluation functionality, which is a single-output functionality. Furthermore, for simplicity, we will restrict our description to the case where none of the parties aborts and at least one of the two parties is honest.

Let $f : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^* \times \{0,1\}^*$ be a deterministic single-output functionality, i.e. f is of the form $f(x, y) = (f_1(x, y), \lambda)$.

We first define the ideal model:

Inputs Each party obtains an input, denoted u.

Sending inputs to the trusted party An honest party will always send its input u to the trusted party. A malicious party will send some input u', which may depend on its initial input and auxiliary input.

Answer of the trusted party Upon obtaining (x, y), the trusted party will reply $f_1(x, y)$ to the first party.

Output An honest party will always output the message obtained from the trusted party. A malicious party may output a polytime computable function of its initial input, its auxiliary input and the message obtained from the trusted party.

Let (B_1, B_2) be a pair of PPT representing strategies in the ideal model, such that at least one of the two parties is honest. The joint distribution of f under (B_1, B_2) in the ideal model, on input pair (x, y) and auxiliary input z, denoted by $\mathrm{IDEAL}_{f,B_1(z),B_2(z)}$, is:

- in the case where B_1 is honest, $\mathrm{IDEAL}_{f,B_1(z),B_2(z)}(x, y) = (f_1(x, B_2(y, z)), B_2(y, z, \lambda))$
- in the case where B_2 is honest, $\mathrm{IDEAL}_{f,B_1(z),B_2(z)}(x, y) = (B_1(x, z, f_1(B_1(x, z), y)), \lambda)$

We now describe the real model. Let Π be a two-party protocol for computing f. Let (A_1, A_2) be a pair of PPT representing strategies in the real model, such that at least one of two parties is honest (i.e. follows the strategy specified by Π). The joint execution of Π under (A_1, A_2) in the real model, on input pair (x, y) and auxiliary input z, denoted by $\mathrm{REAL}_{\Pi,A_1(z),A_2(z)}$ is defined as the output pair resulting from the interaction between $A_1(x, z)$ and $A_2(x, z)$.

Definition A.1 *Let $f : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^* \times \{0,1\}^*$ be a deterministic single-output functionality and Π be a two-party protocol for computing f. Protocol Π* securely computes *f if for every PPT pair (A_1, A_2) (such that at least one party follows the strategy specified by Π), there exists a PPT pair (B_1, B_2) (such that the corresponding party is honest in the ideal model) such that:*

$$\{\mathrm{IDEAL}_{f,B_1(z),B_2(z)}(x, y)\}_{x,y,z} \stackrel{c}{\equiv} \{\mathrm{REAL}_{\Pi,A_1(z),A_2(z)}(x, y)\}_{x,y,z}$$

Assuming the existence of trapdoor permutations, it is known how to obtain a secure protocol for any two-party computation ([29]).

B Almost Pairwise Independence

Lemma B.1 *For a given dictionary $\mathcal{D}' = \{0,1\}^{d'} \subseteq \{0,1\}^n$, there exists a family of almost pairwise independent hash functions $\mathcal{H} = \{h_{w'} : \{0,1\}^n \to \{0,1\}^m\}$ for $\mu = O\left(\frac{n}{|\mathcal{D}'|^{1/3} \log |\mathcal{D}'|}\right) = O\left(\frac{n}{d' 2^{d'/3}}\right)$.*

Proof Sketch: Let $\mathbb{F}$ be a finite field of characteristic 2 (the size of $\mathbb{F}$ will be determined below). Let k be an integer (its value will be determined below). An element p of $\mathbb{F}^k$ can be seen as a polynomial of degree $(k-1)$ over $\mathbb{F}$. Given an index $(\alpha, \beta, \gamma) \in \mathbb{F}^\star \times \mathbb{F} \times \mathbb{F}$, we define the hash function $h_{\alpha,\beta,\gamma}$ as follows: $h_{\alpha,\beta,\gamma} : \mathbb{F}^k \to \{0,1\}^m, p \mapsto \lfloor \alpha p(\beta) + \gamma \rfloor_m$, i.e. we consider the m-bit prefix of $\alpha p(\beta) + \gamma \in \mathbb{F}$.

We will now verify that $\mathcal{H} = \{h_{\alpha,\beta,\gamma} : \mathbb{F}^k \to \{0,1\}^m\}$ is a family of almost pairwise independent hash functions. First, note that when we choose $(\alpha, \beta, \gamma) \overset{R}{\leftarrow} \mathbb{F}^\star \times \mathbb{F} \times \mathbb{F}$, for every $p \in \mathbb{F}^k$, $h_{\alpha,\beta,\gamma}(p)$ is uniform over $\{0,1\}^m$ just by randomness of γ. Also, for a fixed triple $(\alpha, \beta, \gamma) \in \mathbb{F}^\star \times \mathbb{F} \times \mathbb{F}$ and a fixed element $y \in \{0,1\}^m$, $\Pr_{p\in\mathbb{F}^k}[h_{\alpha,\beta,\gamma}(p) = y] = 2^{-m}$, hence the function $h_{\alpha,\beta,\gamma}$ is regular.

We will now prove that $\forall p \neq q \in \mathbb{F}^k$, $\forall y_1, y_2 \in \{0,1\}^m$,

$$\Pr_{\alpha,\beta,\gamma}[h_{\alpha,\beta,\gamma}(p) = y_1 \wedge h_{\alpha,\beta,\gamma}(q) = y_2] \leq \frac{k-1}{|\mathbb{F}|2^m} + \frac{1}{2^{2m}}$$

Consider the case where $y_1 \neq y_2$:

- Suppose that β is a root of the polynomial $p - q$. It is impossible to have $h_{\alpha,\beta,\gamma}(p) = y_1$ and $h_{\alpha,\beta,\gamma}(p) = y_2$.
- Suppose that β is not a root of the polynomial $p - q$.

$$\Pr_{\alpha,\beta,\gamma}[h_{\alpha,\beta,\gamma}(p) = y_1 \wedge h_{\alpha,\beta,\gamma}(q) = y_2 | p(\beta) \neq q(\beta)] \leq \frac{|\mathbb{F}|}{2^{2m}(|\mathbb{F}|-1)}$$

Consider the case where $y_1 = y_2$:

- Suppose that β is a root of the polynomial $p - q$ (which happens with probability at most $\frac{k-1}{|\mathbb{F}|}$). Then by the regularity property,

$$\Pr_{\alpha,\beta,\gamma}[h_{\alpha,\beta,\gamma}(p) = y_1 \wedge h_{\alpha,\beta,\gamma}(q) = y_1 | p(\beta) = q(\beta)] = 2^{-m}$$

- Suppose that β is not a root of the polynomial $p - q$.

$$\Pr_{\alpha,\beta,\gamma}[h_{\alpha,\beta,\gamma}(p) = y_1 \wedge h_{\alpha,\beta,\gamma}(q) = y_1 | p(\beta) \neq q(\beta)] \leq \frac{1}{2^{2m}}$$

By choosing $|\mathbb{F}| = \sqrt[3]{|\mathcal{D}'|}$, $k = \frac{3n}{\log|\mathcal{D}'|}$, we obtain the parameters $\mu = O\left(\frac{n}{|\mathcal{D}'|^{1/3}\log|\mathcal{D}'|}\right) = O\left(\frac{n}{d' 2^{d'/3}}\right)$ and $m \geq \frac{1}{3}d' + \log d' - \log n - O(1)$. $\square$