

Tighter Proofs for the SIGMA and TLS 1.3 Key Exchange Protocols

HANNAH DAVIS*

FELIX GÜNTHER†

August 26, 2020

Abstract

We give new proofs that justify the SIGMA and TLS 1.3 key exchange protocols not just in principle, but in practice. By this we mean that, for standardized elliptic curve group sizes, the overall protocol actually achieves the intended security level.

Prior work gave reductions of both protocols' security to the underlying building blocks that were loose (in the number of users and/or sessions), so loose that they gave no guarantees for practical parameters. Adapting techniques by Cohn-Gordon et al. (Crypto 2019), we give reductions for SIGMA and TLS 1.3 to the strong Diffie–Hellman problem which are tight, and prove that this problem is as hard as solving discrete logarithms in the generic group model. Leveraging our tighter and fully-quantitative bounds, we meet the protocols' targeted security levels when instantiated with standardized curves and improve over prior bounds by up to over 80 bits of security across a range of real-world parameters.

* Department of Computer Science & Engineering, University of California San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. Email: h3davis@eng.ucsd.edu. URL: <https://cseweb.ucsd.edu/~h3davis/>. Supported in part by NSF grant CNS-1717640.

† Department of Computer Science, ETH Zürich, Zürich, Switzerland. Email: mail@felixguenther.info. URL: <https://www.felixguenther.info>. Supported in part by Research Fellowship grant GU 1859/1-1 of the German Research Foundation (DFG) and National Science Foundation (NSF) grant CNS-1717640.

Contents

1	Introduction	3
1.1	Qualitative and Quantitative Bounds	3
1.2	Contributions	5
1.3	Optimizations, Limitations, and Possible Extensions	6
1.4	Concurrent Work	7
2	AKE Security Model	7
2.1	Key Exchange Protocols	8
2.2	Key Exchange Security	8
2.3	Security Properties	9
3	Assumptions, Building Blocks, and Multi-User Security	11
3.1	Decisional and Strong Diffie–Hellman	11
3.2	Multi-User PRF Security	12
3.3	Multi-User Unforgeability with Adaptive Corruptions of Signatures and MACs . . .	12
3.4	Hash Function Collision Resistance	15
4	The SIGMA Protocol	15
5	Tighter Security Proof for SIGMA-I	16
6	The TLS 1.3 Handshake Protocol	29
6.1	Protocol Description	29
6.2	Handling the TLS 1.3 Key Schedule	31
7	Tighter Security Proof for the TLS 1.3 Handshake	32
8	Evaluation	39
A	Evaluation Details	46
A.1	Fully-quantitative CK SIGMA Bound	46
A.2	Fully-quantitative DFGS TLS 1.3 Bound	47
A.3	Numerical bounds for SIGMA and TLS 1.3	47
B	Proof of the Strong Diffie–Hellman GGM Bound	51

1 Introduction

The Transport Layer Security (TLS) protocol [47] is responsible for securing billions of Internet connections every day. Usage statistics for Google Chrome¹ and Mozilla Firefox² report that 78–97% of all web page accesses are encrypted. At the heart of TLS is an authenticated key exchange (AKE) protocol, the so-called handshake protocol, responsible for providing the parties (client and server) with a shared, symmetric key that is fresh, private and authenticated. The ensuing record layer secures data using this key. The AKE protocol of TLS is based on the SIGMA (“SIGn-and-MAC”) design of Krawczyk [37] for the Internet Key Exchange (IKE) protocol [32] of IPsec [36], which generically augments an unauthenticated, ephemeral Diffie–Hellman (DH) key exchange with authenticating signatures and MACs.

Naturally, the SIGMA AKE protocol, and its incarnation in TLS, have been the recipients of proofs of security. We contend that these, largely, justify the AKE protocols in principle, but not in practice, meaning not for the parameters in actual use and at the desired or expected level of security. Our work takes steps towards filling this gap.

1.1 Qualitative and Quantitative Bounds

Let us expand on this. The protocols KE we consider are built from a cyclic group \mathbb{G} in which some DH problem P is assumed to be hard, a pseudorandom function PRF and unforgeable signature and MAC schemes S and M . The target for KE is session-key security with explicit authentication as originating from [12, 10, 18]. A proof of security has both a qualitative and quantitative dimension. Qualitatively, a proof of security for the AKE protocol KE says that KE meets its target definition assuming the building blocks meet theirs, where, in either case, meeting the definition means any poly-time adversary has negligible advantage in violating it.

The quantitative dimension associates to each adversary in the security game of KE a set of resources r , representing its runtime and attack surface (e.g., the number of users and executed protocol sessions the adversary has access to). It then relates the maximum advantage of any r -resource adversary in breaking KE’s security to likewise advantage functions for the building blocks through an equation of the (simplified) form

$$\text{Adv}_{\text{KE}}(r) \leq f_{\mathbb{G}} \cdot \text{Adv}_{\mathbb{G}}^P(r_{\mathbb{G}}) + f_S \cdot \text{Adv}_S^{\text{EUF-CMA}}(r_S) + \dots,$$

deriving quantitative factors f_X and resources r_X for the advantage of each building block X .

Speaking asymptotically again, when f_X and r_X are polynomial functions in r , then $\text{Adv}_{\text{KE}}(r)$ is negligible whenever all building blocks’ advantages are. Due to the complexity of key exchange models and the challenging task of combining the right components in a secure manner, key exchange analyses (including those for SIGMA [19] and TLS 1.3 [24, 40, 26, 28]) indeed often remain abstract and consider only qualitative, asymptotic security bounds.

Standardized protocols like TLS in contrast have to define concrete choices for each cryptographic building block. This involves considering reasonable estimates for adversarial resources (like runtime t and number of key-exchange model queries q) and specific instances and parameters for the underlying components X . One would hope that key exchange proofs can provide guidance in making sound choices that result in the desired overall security level. Unfortunately, AKE security bounds regularly are highly non-tight, meaning that f_X and/or r_X for some components X are so large that reasonable stand-alone parameters for X yield vacuous key exchange advantages for

¹<https://transparencyreport.google.com/https/>

²<https://telemetry.mozilla.org/>

Adv. resources					SIGMA		TLS 1.3	
t	$\#U$	$\#S$	Curve	Target	CK [19]	Us (Thm 5.1)	DFGS [24]	Us (Thm 7.1)
2^{60}	2^{20}	2^{35}	secp256r1	2^{-68}	$\approx 2^{-59}$	$\approx 2^{-114}$	$\approx 2^{-62}$	$\approx 2^{-114}$
2^{60}	2^{30}	2^{55}	secp256r1	2^{-68}	$\approx 2^{-19}$	$\approx 2^{-104}$	$\approx 2^{-22}$	$\approx 2^{-104}$
2^{60}	2^{20}	2^{35}	x25519	2^{-68}	$\approx 2^{-57}$	$\approx 2^{-112}$	$\approx 2^{-60}$	$\approx 2^{-112}$
2^{60}	2^{30}	2^{55}	x25519	2^{-68}	$\approx 2^{-17}$	$\approx 2^{-102}$	$\approx 2^{-20}$	$\approx 2^{-102}$
2^{80}	2^{20}	2^{35}	secp256r1	2^{-48}	$\approx 2^{-19}$	$\approx 2^{-74}$	$\approx 2^{-22}$	$\approx 2^{-74}$
2^{80}	2^{30}	2^{55}	secp256r1	2^{-48}	1	$\approx 2^{-64}$	1	$\approx 2^{-64}$
2^{80}	2^{20}	2^{35}	x25519	2^{-48}	$\approx 2^{-17}$	$\approx 2^{-72}$	$\approx 2^{-20}$	$\approx 2^{-72}$
2^{80}	2^{30}	2^{55}	x25519	2^{-48}	1	$\approx 2^{-62}$	1	$\approx 2^{-62}$
2^{80}	2^{20}	2^{35}	secp384r1	2^{-112}	$\approx 2^{-149}$	$\approx 2^{-204}$	$\approx 2^{-152}$	$\approx 2^{-204}$
2^{80}	2^{30}	2^{55}	secp384r1	2^{-112}	$\approx 2^{-109}$	$\approx 2^{-194}$	$\approx 2^{-112}$	$\approx 2^{-194}$

Table 1: Exemplary concrete advantages of a key exchange adversary with given resources t (running time), $\#U$ (number of users), $\#S$ (number of sessions), in breaking the security of the SIGMA and TLS 1.3 protocols when instantiated with curve `secp256r1`, `secp384r1`, or `x25519`, based on the prior bounds by Canetti-Krawczyk [19] resp. Dowling et al. [24, 26], and the bounds we establish (Theorem 5.1 and 7.1). Target indicates the maximal advantage $t/2^b$ tolerable when aiming for the respective curve’s security level ($b = 128$ resp. 192 bits); entries in red-shaded cells miss that target. See Section 8 and Appendix A for full details and curves `secp521r1` and `x448`.

practical parameters. While the asymptotic bound tells us that scaling up the parameters for X (say, the DDH problem) will at some point result in a secure overall advantage, this causes efficiency concerns (e.g., doubling elliptic curve DH security parameters means quadrupling the cost for group operations) and hence does not happen in practice.

We illustrate in Table 1 the effects of the non-tight bounds for SIGMA and TLS 1.3 when instantiating the protocols with NIST curves `secp256r1`, `secp384r1` [45], or curve `x25519` [42] and idealizing protocols’ other components (cf. Section 8 for full details). Following the curves’ security, we aim at a security level of 128 bits, resp. 192 bits, meaning the ratio of an adversary’s runtime to its advantage should be bounded by 2^{-128} , resp. 2^{-192} . When considering the advantage of key exchange adversaries running in time t , interacting in the key exchange security game with $\#U$ users and $\#S$ sessions, we can see that previous security bounds fail to meet the targeted level of 128/192-bit security for real-world-scale parameters ($\#U$ ranging in 2^{20} – 2^{30} based on 2^{27} active certificates on the Internet³, $\#S$ ranging in 2^{35} – 2^{55} based on 2^{32} Internet users and 2^{33} daily Google searches⁴). In the security analysis by Canetti and Krawczyk [19] (CK) for SIGMA, the factor associated to the Decisional Diffie–Hellman problem is $f_{\text{DDH}}(t, \#U, \#S) = \#U \cdot \#S$, where $\#U$ is the number of users and $\#S$ is the number of sessions accessible by the adversary. The analysis by Dowling et al. [24, 26] (DFGS) for TLS 1.3 reduces to the strong Diffie–Hellman problem [1] with factor $f_{\text{stDH}}(t, \#U, \#S) = (\#S)^2$. In contrast, we reduce to the strong Diffie–Hellman problem with a constant factor for both SIGMA and TLS 1.3.

Let us discuss three data points from Table 1:

1. Already with medium-sized resources, investing time $t = 2^{60}$ and interacting with a million users ($\#U = 2^{20}$) and a few billion sessions ($\#S = 2^{35}$), the CK [19] and DFGS [24, 26] advantage bounds for SIGMA and TLS 1.3 with curves `secp256r1` and `x25519` fall 6–11 bits below the target of 2^{-68} for 128-bit security.

³<https://letsencrypt.org/stats/>

⁴<https://www.internetlivestats.com/>

2. When considering a more powerful, global-scale adversary ($t = 2^{80}$, $\#U = 2^{30}$, $\#S = 2^{55}$), both CK and DFGS bounds for `secp256r1/x25519` become fully vacuous; the upper bound on the probability of the adversary breaking the protocol is 1. We stress that `secp256r1` is the mandatory-to-implement curve for TLS 1.3; `secp256r1` and `x25519` together make up for 90% of the TLS 1.3 ECDHE handshakes reported through Firefox Telemetry.
3. Finally, and notably, even switching to the higher-security curve `secp384r1` helps only marginally in the latter case: the resulting advantage against SIGMA falls 3 bits short of the 192-bit security target of 2^{-112} , and the TLS advantage bound only barely meets that target.

For all curves and choices of parameters, our bounds do better.

1.2 Contributions

Most prior results in tightly secure key exchange (e.g., [4, 29]) apply only to bespoke protocols, carefully designed to allow for tighter proof techniques, at the cost of requiring more complex primitives which, in the end, eat up the gained practical efficiency. Recently, Cohn-Gordon et al. [21, 22] established a proof strategy for a simple and efficient DH key exchange with reasonable tightness loss (only linear in the number of users $\#U$), achieving implicit authentication through static DH keys through careful key derivation via a random oracle [11] with an optional explicit-authentication step.

Our work is the first to apply the techniques by Cohn-Gordon et al. to existing and standardized AKE protocols. We give tight reductions for the security of SIGMA and TLS 1.3 to the strong Diffie–Hellman problem [1], which in addition we prove is as hard as the discrete logarithm problem in the generic group model (GGM) [48, 43]. Instantiating our bounds shows that, with standardized real-world parameters, we achieve the intended security levels even when considering powerful, globally-scaled attackers.

Code-based security model and proofs. For our proofs, we provide detailed proof steps and reductions using the code-based game-playing framework of Bellare and Rogaway [13]. Our security model is similar to the one applied by Cohn-Gordon et al. [21], but formalized also as a code-based game (in Section 2) and stronger in that it captures explicit authentication and regular (“perfect”) forward secrecy (instead of only weak forward secrecy in [21]).

Tighter security proof of SIGMA(-I). We establish fully quantitative security bounds for SIGMA and its identity-protecting variant SIGMA-I [37] in Sections 4 and 5. Our result is for BR-like [12] key exchange security and gives a tight reduction to the strong Diffie–Hellman problem [1] in the used DH group, and to the multi-user (mu) security of the employed pseudorandom function (PRF), signature scheme, and MAC scheme, adapting the techniques by Cohn-Gordon et al. [21] in the random oracle model [11]. The latter mu-security bounds are essentially equivalent to the corresponding bounds by CK [19]. Our improvement comes from shaving off a factor of $\#U \cdot \#S$ (number of users times number of sessions) on the DH problem advantage compared to CK. While we move to the interactive strong Diffie–Hellman problem (compared to the decisional DH (DDH) problem [16] used in [19]), we prove (in Appendix B) that the strong DH problem, like DDH, is as hard as solving discrete logarithms in the generic group model [48, 43], reflecting the (only generic) algorithms known for solving discrete logarithms in elliptic curve groups.

Tighter security proof for the TLS 1.3 DH handshake. We likewise establish fully quantitative security bounds for the key exchange of the recently standardized newest version of the Transport Layer Security protocol, TLS 1.3 [47], in Sections 6 and 7. The main quantitative improvement in our reduction is again a tight reduction to the strong DH problem, whereas prior bounds by DFGS [24, 26] incurred a quadratic loss to the PRF-ODH assumption [34], a loss which translates directly to strong DH [17]. While TLS 1.3 roughly follows the SIGMA-I design, its cascading key schedule impedes the precise technique of Cohn-Gordon et al. [21] and a direct application of our results on SIGMA-I, as no single function (to be modeled as a random oracle) binds the Diffie–Hellman values to the session context. We therefore have to carefully adapt the proof to accommodate the more complex key schedule and other core variations in TLS 1.3’s key exchange, achieving conceptually similar tightness results as for SIGMA-I.

Evaluation. In Section 8, we evaluate the concrete security implications of our improved bounds for SIGMA and TLS 1.3 for a wide range of real-world resource parameters and all five elliptic curves (`secp256r1`, `secp384r1`, `secp521r1`, `x25519`, `x448`) standardized for use in TLS 1.3 [47], a summary of which is displayed in Table 1. Leveraging our GGM bound for the strong Diffie–Hellman problem, we focus on the hardness of solving discrete logarithms in the respective elliptic curve groups, instantiating signatures based on ECDSA [45] resp. EdDSA [15]. We idealized the symmetric PRF, MAC, and hash function primitives (in two variants, with key and output sizes twice as large as the curve’s security level, or fixed at 256 bits corresponding to the choice in most TLS 1.3 cipher suites).

We report that our tighter proofs indeed materialize for a wide range of real-world resource parameters (adversary runtime $t \in \{2^{40}, 2^{60}, 2^{80}\}$, number of users $\#U \in \{2^{20}, 2^{30}\}$, and number of sessions $\#S \in \{2^{35}, 2^{45}, 2^{55}\}$). The resulting attacker advantages meet the targeted security levels of all five curves. In comparison to the prior CK [19] SIGMA and DFGS [24, 26] TLS 1.3 bounds, our results improve the obtained security across these real-world parameters by up to 82 bits for SIGMA and 85 bits for TLS 1.3, respectively.

1.3 Optimizations, Limitations, and Possible Extensions

SIGMA being a generic AKE design, the signature, PRF, and MAC schemes may be instantiated with primitives optimized for multi-user security. While we focus on standardized and deployed schemes in our evaluation without assuming tight mu-security, our SIGMA bound allows to directly leverage such optimization. For PRFs and MACs, efficient candidates exist (e.g., AMAC [6]). For signatures, tight mu-security is more challenging [5] and often involves computationally much more expensive constructions [4].

Like Cohn-Gordon et al. [21], our key exchange security model considers exposure of long-term secrets and session keys, but does not allow revealing internal session state or randomness (as in the (e)CK model [18, 41]). This is appropriate for protocols like TLS 1.3 not aiming to protect against such threats. The original SIGMA proof [19] did establish security in the CK model [18] allowing exposure of session state; in that sense our results are qualitatively weaker. Establishing a tight security reduction for a DH-based AKE protocol which can handle adaptive compromises of session state (including DH exponents) remains a challenging open problem.

In our proofs, we crucially rely on the ability to observe and program a random oracle used for key derivation in the AKE protocol, borrowing from [21]. Notably, the approach of Cohn-Gordon et al. is tailored to an AKE protocol achieving authenticity implicitly through mixing long-term DH keys into the key derivation. Our proofs can hence be seen as translating and adapting their technique to the setting of SIGMA and TLS 1.3, where an unauthenticated ephemeral DH exchange

is explicitly authenticated through signatures and MACs, confirming that the generic SIGMA design as well as the standardized TLS 1.3 protocol bind enough context to their DH shares for this proof technique to work. Leveraging the random oracle model [11] is another qualitative difference compared to the original SIGMA proof [19] in the standard model. Interestingly, this distinction vanishes in comparison to the provable security results for the TLS 1.3 handshake protocol [24, 26, 28] which employ the PRF-ODH assumption [34, 17], an interactive assumption which plausibly can only be instantiated in the random oracle model (from the strong DH assumption).

1.4 Concurrent Work

In concurrent and independent work, Diemert and Jager [23] studied the tight security of the main TLS 1.3 handshake. Despite the use of different security models, their bounds and ours provide similarly tight reductions to the strong Diffie–Hellman problem. Their analysis is carried out in the multi-stage key exchange model [27], proving security not only of the final session key, but also of intermediate handshake encryption keys and further secrets. In their proof, the full cascading derivation of each intermediate key is modeled as an independent, programmable random oracle (cf. [23, Theorem 6]) to tackle the challenge that TLS 1.3’s key schedule does not bind DH values and session context in one function (or random oracle). We instead model the key schedule’s inner HKDF [39] extraction and expansion functions as two individual random oracles, carefully connected via efficient look-up tables. This yields more compact bounds and allows our analysis to stay closer to the use of HKDF in TLS 1.3, where the output of one extraction call is used to derive multiple keys. Beyond the strong DH problem, both their and our proofs further reduce to multi-user security of signatures and PRFs, applying random oracle bounds for the latter (cf. [23, Section 5] and our Section 3.2). In addition, our model captures explicit authentication (which we prove via HMAC’s unforgeability) and our bounds are fully parameterized enabling the evaluation of concrete practical advantages (cf. Section 8). Finally, while our work additionally treats the more generic SIGMA-I protocol and proves GGM bounds for the strong DH problem, Diemert and Jager [23] further study composition of the TLS 1.3 handshake with the nonce-randomized AES-GCM encryption in the record protocol, connecting tight multi-user bounds for the latter [14, 33] with a tighter version of prior MSKE composition results [31] they establish.

2 AKE Security Model

We provide our results in a game-based key exchange model formalized in Figure 1, at its core following the seminal work by Bellare and Rogaway [12] considering an active network adversary that controls all communication (initiating sessions and determining their next inputs through SEND queries) and is able to corrupt long-term secrets (REVLONGTERMKEY) as well as session keys (REVSESSIONKEY). The adversary’s goal is then to (a) distinguish a session key in a “fresh” (not trivially compromised, captured through a `Fresh` predicate) session from a uniformly random key obtained through TEST queries (breaking *key secrecy*), or (b) make a session accept without matching communication partner (breaking *explicit authentication*).

Following Cohn-Gordon et al. [21], we formalize our model in a real-or-random version (following Abdalla, Fouque, and Pointcheval [3] with added forward secrecy [2]) with *many* TEST queries which all answer with a real or uniformly random session key based on the *same* random bit b . In contrast to [21], our model additionally captures explicit authentication through the `ExplicitAuth` predicate in Figure 1, ensuring sessions with non-corrupted peer accept with an honest partner session. Our model further captures protocols where the communication partner’s identity of a session may be

unknown at the outset and only learned during the protocol execution; this setting of “post-specified peers” [19] particularly applies to the SIGMA protocol family [37] as well as TLS 1.3 [47].

2.1 Key Exchange Protocols

We begin with formalizing the syntax of key exchange protocols. A key exchange protocol KE consists of three algorithms (KGen , Activate , Run) and an associated key space KE.KS (where most commonly $\text{KE.KS} = \{0, 1\}^n$ for some $n \in \mathbb{N}$). The key generation algorithm $\text{KGen}() \xrightarrow{\$} (pk, sk)$ generates new long-term public/secret key pairs. In the security model, we will associate key pairs to distinct *users* (or *parties*) with some identity $u \in \mathbb{N}$ running the protocol, and log the public long-term keys associated with each user identity in a list peerpk . (The adversary will be in control of initializing new users, identified by an increasing counter, and we assume it only references existing user identities.) The activation algorithm $\text{Activate}(id, sk, \text{peerid}, \text{peerpk}, \text{role}) \xrightarrow{\$} (st', m')$ initiates a new session for a given user identity id (and associated long-term secret key sk) acting in a given role $\text{role} \in \{\text{initiator}, \text{responder}\}$ and aiming to communicate with some peer user identity peerid . Activate also takes as input the list peerpk of all users’ public keys; protocols may use this list to look up their own and their peers’ public keys. We provide the entire list instead of just the user’s and peers’ public keys to accommodate protocols with post-specified peer. These protocols may leave peerid unspecified at the time of session activation; when the peer identity is set at some later point, the list can be used to find the corresponding long-term key. Activation outputs a session state and (if $\text{role} = \text{initiator}$) first protocol message m' , and will be invoked in the security model to create a new session π_u^i at a user u (where the label i distinguishes different sessions of the same user). Finally, $\text{Run}(id, sk, st, \text{peerpk}, m) \xrightarrow{\$} (st', m')$ delivers the next incoming key exchange message m to the session of user id with secret key sk and state st , resulting in an updated state st' and a response message m' . Like Activate , it relies on the list peerpk to look up its own and its peer’s long-term keys.

The state of each session in a key exchange protocol contains at least the following variables, beyond possibly further, protocol-specific information:

$\text{peerid} \in \mathbb{N}$. Reflects the (intended) partner identity of the session; in protocols with post-specified peers this is learned and set (once) by the session during the protocol execution.

$\text{role} \in \{\text{initiator}, \text{responder}\}$. The session’s role, determined upon activation.

$\text{status} \in \{\text{running}, \text{accepted}, \text{rejected}\}$. The session’s status; initially $\text{status} = \text{running}$, a session accepts when it switches to $\text{status} = \text{accepted}$ (once).

$\text{sk}_{\text{key}} \in \text{KE.KS}$. The derived session key (in the protocol-specific key space KE.KS), set upon acceptance.

sid . The session identifier used to define partnered session in the security model; initially unset, sid is determined (once) during protocol execution.

2.2 Key Exchange Security

We formalize our key exchange security game $G_{\text{KE}, \mathcal{A}}^{\text{KE-SEC}}$ based on the concepts introduced above in Figure 1, following the framework for code-based game playing by Bellare and Rogaway [13]. After initializing the game, the adversary \mathcal{A} is given access to queries NEWUSER (generating a new user’s public/secret key pair), SEND (controlling activation and message processing of sessions),

REVSESSIONKEY (revealing session keys), REVLONGTERMKEY (corrupting user’s long-term secrets), and TEST (providing challenge real-or-random session keys), as well as a FINALIZE query to which it will submit its guess b' for the challenge bit b , ending the game.

The game $G_{\text{KE},\mathcal{A}}^{\text{KE-SEC}}$ then (in FINALIZE) determines whether \mathcal{A} was successful through the following predicates, formalized in pseudocode in Figure 1:

Sound. The soundness predicate **Sound** checks that (a) no three session identifiers collide (hence the session identifier properly serves to identify two partnered sessions). Furthermore, it ensures that (b) accepted sessions with the same session identifier, agreeing partner identities, and distinct roles derive the same session key. The adversary breaks soundness if it violates either of these properties.

ExplicitAuth. The predicate **ExplicitAuth** captures explicit authentication in that it requires that for any session of some user id that accepted while its partner $peerid$ was not corrupted (captured through logging relative acceptance time t_{acc} and long-term reveal time $\text{rev}t_{k_{peerid}}$) has (a) a partnered session run by the intended peer identity and in an opposite role, and (b) if that partnered session accepts, it will do so with peer identity id . The adversary breaks explicit authentication if this predicate evaluates to false.

Fresh. Finally, to capture key secrecy, we have to restrict the adversary to testing so-called *fresh* sessions in order to exclude trivial attacks, which the freshness predicate **Fresh** ensures. A tested session is non-fresh, if (a) its session key has been revealed (in which case \mathcal{A} knows the real key), (b) its partnered session (through sid) has been revealed or tested (in which case \mathcal{A} knows the real key or may see two different random keys), or (c) its intended peer identity was compromised prior to accepting (in which case \mathcal{A} may fully control the communication partner). If the adversary violates freshness, we invalidate its guess by overwriting $b' \leftarrow 0$.

We call two distinct sessions π_u^i and π_v^j *partnered* if $\pi_u^i.sid = \pi_v^j.sid$. We refer to sessions generated by **Activate** (i.e., controlled by the game) as *honest* sessions to reflect that their behavior is determined honestly by the game and not the adversary. The long-term key of an honest session may still be corrupted, or its session key may be revealed without affecting this notion of “honesty”.

Definition 2.1 (Key exchange security). *Let KE be a key exchange protocol and $G_{\text{KE},\mathcal{A}}^{\text{KE-SEC}}$ be the key exchange security game defined in Figure 1. We define*

$$\text{Adv}_{\text{KE}}^{\text{KE-SEC}}(t, q_N, q_S, q_{RS}, q_{RL}, q_T) := 2 \cdot \max_{\mathcal{A}} \Pr \left[G_{\text{KE},\mathcal{A}}^{\text{KE-SEC}} \Rightarrow 1 \right] - 1,$$

where the maximum is taken over all adversaries, denoted $(t, q_N, q_S, q_{RS}, q_{RL}, q_T)$ -KE-SEC-adversaries, running in time at most t and making at most q_N, q_S, q_{RS}, q_{RL} , resp. q_T queries to their oracles **NEWUSER**, **SEND**, **REVSESSIONKEY**, **REVLONGTERMKEY**, resp. **TEST**.

2.3 Security Properties

Let us briefly revisit some core security properties captured in our key exchange security model.

First, we capture regular *key secrecy* through **TEST** queries, incorporating *forward secrecy* (sometimes called “perfect” forward secrecy) by allowing the adversary to corrupt any user as long as all tested sessions accept prior to corrupting their intended peer. This strengthens our model compared to that of Cohn-Gordon et al. [21] which only captures weak forward secrecy where the adversary has to be passive in sessions where it corrupts long-term secrets.

$G_{KE,A}^{KE-SEC}$

INITIALIZE:

1 $time \leftarrow 0$; $users \leftarrow 0$
 2 $b \xleftarrow{\$} \{0, 1\}$

NEWUSER:

3 $users \leftarrow users + 1$
 4 $(pk_{users}, sk_{users}) \xleftarrow{\$} KGen()$
 5 $revltk_{users} \leftarrow \infty$
 6 $peerpk[users] \leftarrow pk_{users}$
 7 return pk_{users}

SEND(u, i, m):

8 if $\pi_u^i = \perp$ then
 9 $(peerid, role) \leftarrow m$
 10 $(\pi_u^i, m') \xleftarrow{\$} Activate(u, sk_u, peerid, peerpk, role)$
 11 $\pi_u^i.t_{acc} \leftarrow 0$
 12 else
 13 $(\pi_u^i, m') \xleftarrow{\$} Run(u, sk_u, \pi_u^i, peerpk, m)$
 14 if $\pi_u^i.status = accepted$ then
 15 $time \leftarrow time + 1$
 16 $\pi_u^i.t_{acc} \leftarrow time$
 17 return m'

REVSSESSIONKEY(u, i):

18 if $\pi_u^i = \perp$ or $\pi_u^i.status \neq accepted$ then
 19 return \perp
 20 $\pi_u^i.revealed \leftarrow true$
 21 return $\pi_u^i.skey$

REVLONGTERMKEY(u):

22 $time \leftarrow time + 1$
 23 $revltk_u \leftarrow time$
 24 return sk_u

TEST(u, i):

25 if $\pi_u^i = \perp$ or $\pi_u^i.status \neq accepted$ or $\pi_u^i.tested$ then
 26 return \perp
 27 $\pi_u^i.tested \leftarrow true$
 28 $T \leftarrow T \cup \{\pi_u^i\}$
 29 $k_0 \leftarrow \pi_u^i.skey$
 30 $k_1 \xleftarrow{\$} KE.KS$
 31 return k_b

FINALIZE(b'):

32 if $\neg Sound$ then
 33 return 1
 34 if $\neg ExplicitAuth$ then
 35 return 1
 36 if $\neg Fresh$ then
 37 $b' \leftarrow 0$
 38 return $[[b = b']]$

Sound:

1 if \exists distinct $\pi_u^i, \pi_v^j, \pi_w^k$ with $\pi_u^i.sid = \pi_v^j.sid = \pi_w^k.sid$
 then // no triple sid match
 2 return false
 3 if $\exists \pi_u^i, \pi_v^j$ with
 $\pi_u^i.status = \pi_v^j.status = accepted$
 and $\pi_u^i.sid = \pi_v^j.sid$
 and $\pi_u^i.peerid = v$ and $\pi_v^j.peerid = u$
 and $\pi_u^i.role \neq \pi_v^j.role$, but $\pi_u^i.skey \neq \pi_v^j.skey$ then
 // partnering implies same key
 4 return false
 5 return true

ExplicitAuth:

1 return
 $\forall \pi_u^i : \pi_u^i.status = accepted$
 and $\pi_u^i.t_{acc} < revltk_{\pi_u^i.peerid}$
 // all sessions accepting with a non-corrupted peer ...
 $\implies \exists \pi_v^j : \pi_u^i.peerid = v$
 and $\pi_u^i.sid = \pi_v^j.sid$
 and $\pi_u^i.role \neq \pi_v^j.role$
 // ... have a partnered session ...
 and $(\pi_v^j.status = accepted \implies \pi_v^j.peerid = u)$
 // ... agreeing on the peerid (upon acceptance)

Fresh:

1 for each $\pi_u^i \in T$
 2 if $\pi_u^i.revealed$ then
 3 return false // tested session may not be revealed
 4 if $\exists \pi_v^j \neq \pi_u^i : \pi_v^j.sid = \pi_u^i.sid$
 and $(\pi_v^j.tested$ or $\pi_v^j.revealed)$ then
 5 return false // tested session's partnered session may
 not be tested or revealed
 6 if $revltk_{\pi_u^i.peerid} < \pi_u^i.t_{acc}$ then
 7 return false // tested session's peer may not be cor-
 rupted prior to acceptance
 8 return true

Figure 1: Key exchange security game.

Our model encodes *explicit authentication* (via `ExplicitAuth`), a strengthening compared to the implicit-authentication model in [21].

Like [21], our model captures *key-compromise impersonation* attacks by allowing the session owner’s secret key of tested sessions to be corrupted at any point in time. Similarly, we do *not* capture *session-state or randomness reveals* [18, 41] or *post-compromise security* [20].

3 Assumptions, Building Blocks, and Multi-User Security

Before we continue to our main technical results, let us briefly introduce notation and discuss the multi-user security of the involved building blocks: strong Diffie–Hellman (including the GGM bound we prove), PRFs, digital signatures, MAC schemes, and hash functions.

3.1 Decisional and Strong Diffie–Hellman

The classical decisional Diffie–Hellman assumption [16] states that, when only observing the two Diffie–Hellman shares g^x , g^y , the resulting secret g^{xy} is indistinguishable from a random group element.

Definition 3.1 (Decisional Diffie–Hellman (DDH) assumption). *Let $\mathbb{G} = \langle g \rangle$ be a cyclic group of prime order q . We define*

$$\text{Adv}_{\mathbb{G}}^{\text{DDH}}(t) := \max_{\mathcal{A}} \left| \Pr [\mathcal{A}(\mathbb{G}, g, g^x, g^y, g^{xy}) \Rightarrow 1 \mid x, y \xleftarrow{\$} \mathbb{Z}_q] - \Pr [\mathcal{A}(\mathbb{G}, g, g^x, g^y, g^z) \Rightarrow 1 \mid x, y, z \xleftarrow{\$} \mathbb{Z}_q] \right|,$$

where the maximum is taken over all adversaries, denoted (t) -DDH-adversaries running in time at most t .

The strong Diffie–Hellman assumption, a weakening of the gap Diffie–Hellman assumption [46], states that solving the computational Diffie–Hellman problem given a restricted decisional Diffie–Hellman oracle is hard.

Definition 3.2 (Strong Diffie–Hellman assumption [46]). *Let $\mathbb{G} = \langle g \rangle$ be a cyclic group of prime order q . Let $\text{DDH}(X, Y, Z) := [[X^{\log_g(Y)} = Z]]$ be a decisional Diffie–Hellman oracle. We define*

$$\text{Adv}_{\mathbb{G}}^{\text{stDH}}(t, q_{\text{SDH}}) := \max_{\mathcal{A}} \Pr \left[\mathcal{A}^{\text{DDH}(g^x, \cdot, \cdot)}(\mathbb{G}, g, g^x, g^y) = g^{xy} \mid x, y \xleftarrow{\$} \mathbb{Z}_q \right],$$

where the maximum is taken over all adversaries, denoted (t, q_{SDH}) -stDH-adversaries running in time at most t and making at most q_{SDH} queries to their DDH oracle.

To argue that it is reasonable to rely on the strong Diffie–Hellman assumption, we turn to the generic group model [48, 43]. Although some known attacks like index calculus fall outside the generic group model, the best known algorithms for elliptic curve groups are generic. Shoup [48] proved that, in the generic group model, any adversary computing at most t group operations in a group of prime order p has advantage at most $\mathcal{O}(t^2/p)$ in solving the computational or decisional Diffie–Hellman problem. We claim, and prove in Appendix B, that any adversary in the generic group model making at most t group operations and DDH oracle queries, also has advantage at most $\mathcal{O}(t^2/p)$ in solving the strong Diffie–Hellman problem.

Theorem 3.3. *Let \mathbb{G} be a group with prime order p . In the generic group model, $\text{Adv}_{\mathbb{G}}^{\text{stDH}}(t, q) \leq 4t^2/p$.*

$\underline{G_{\text{PRF}, \mathcal{A}}^{\text{mu-PRF}}}$ <p><u>INITIALIZE:</u></p> <ol style="list-style-type: none"> 1 $b \xleftarrow{\\$} \{0, 1\}$ 2 $u \leftarrow 0$ <p><u>FN(i, x):</u></p> <ol style="list-style-type: none"> 3 return $f_i(x)$ <p><u>FINALIZE(b^*):</u></p> <ol style="list-style-type: none"> 4 return $[[b = b^*]]$ 	$\underline{\text{NEW:}}$ <ol style="list-style-type: none"> 5 $u \leftarrow u + 1$ 6 if $b = 1$ then 7 $K_u \xleftarrow{\\$} \{0, 1\}^k$ 8 $f_u := \text{PRF}(K_u, \cdot)$ 9 else 10 $f_u \xleftarrow{\\$} \text{FUNC}$
---	---

Figure 2: Multi-user PRF security of a pseudorandom function $\text{PRF}: \{0, 1\}^k \times \{0, 1\}^m \rightarrow \{0, 1\}^n$. FUNC is the space of all functions $\{0, 1\}^m \rightarrow \{0, 1\}^n$.

3.2 Multi-User PRF Security

Let us recap the multi-user security notion for pseudorandom functions (PRFs) [7].

Definition 3.4 (Multi-user PRF security). *Let $\text{PRF}: \{0, 1\}^k \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ be a function (for $k, n \in \mathbb{N}$ and $m \in \mathbb{N} \cup \{*\}$) and $G_{\text{PRF}, \mathcal{A}}^{\text{mu-PRF}}$ be defined as in Figure 2. We define*

$$\text{Adv}_{\text{PRF}}^{\text{mu-PRF}}(t, q_{\text{NW}}, q_{\text{FN}}, q_{\text{FN/U}}) := 2 \cdot \max_{\mathcal{A}} \Pr \left[G_{\text{PRF}, \mathcal{A}}^{\text{mu-PRF}} \Rightarrow 1 \right] - 1,$$

where the maximum is taken over all adversaries, denoted $(t, q_{\text{NW}}, q_{\text{FN}}, q_{\text{FN/U}})$ -mu-PRF-adversaries, running in time at most t and making at most q_{NW} queries to their NEW oracle, at most q_{FN} total queries to their FN oracle, and at most $q_{\text{FN/U}}$ queries $\text{FN}(i, \cdot)$ for any user i .

Generically, the multi-user security of PRFs reduces to single-user security (formally, $G_{\text{PRF}, \mathcal{A}}^{\text{mu-PRF}}$ with \mathcal{A} restricted to $q_{\text{NW}} = 1$ queries to NEW) with a factor in the number of users via a hybrid argument [7], i.e.,

$$\text{Adv}_{\text{PRF}}^{\text{mu-PRF}}(t, q_{\text{NW}}, q_{\text{FN}}, q_{\text{FN/U}}) \leq q_{\text{NW}} \cdot \text{Adv}_{\text{PRF}}^{\text{mu-PRF}}(t', 1, q_{\text{FN/U}}, q_{\text{FN/U}}),$$

where $t \approx t'$. (Note that the total number q_{FN} of queries to the FN oracle across all users does not affect the reduction.) There exist simple and efficient constructions, like AMAC [6], that however achieve multi-user security tightly.

If we use a random oracle RO as a PRF with key length kl , then

$$\text{Adv}_{\text{RO}}^{\text{mu-PRF}}(t, q_{\text{NW}}, q_{\text{FN}}, q_{\text{FN/U}}, q_{\text{RO}}) \leq \frac{q_{\text{NW}} \cdot q_{\text{RO}}}{2^{kl}}.$$

3.3 Multi-User Unforgeability with Adaptive Corruptions of Signatures and MACs

We recap the definition of digital signature schemes and message authentication codes (MACs) as well as the natural extension of classical *existential unforgeability under chosen-message attacks* [30] to the *multi-user* setting with *adaptive corruptions*. For signatures, this notion was considered by Bader et al. [4] and, without corruptions, by Menezes and Smart [44].

Definition 3.5 (Signature scheme). *A signature scheme $S = (\text{KGen}, \text{Sign}, \text{Vrfy})$ consists of three efficient algorithms defined as follows.*

- $\text{KGen}() \xrightarrow{\$} (pk, sk)$. This probabilistic algorithm generates a public verification key pk and a secret signing key sk .
- $\text{Sign}(sk, m) \xrightarrow{\$} \sigma$. On input a signing key sk and a message m , this (possibly) probabilistic algorithm outputs a signature σ .
- $\text{Vrfy}(vk, m, \sigma) \rightarrow d$. On input a verification key pk , a message m , and a signature σ , this deterministic algorithm outputs a decision bit $d \in \{0, 1\}$ (where $d = 1$ indicates validity of the signature).

Definition 3.6 (Signature mu-EUF-CMA security). Let S be a signature scheme and $G_{S, \mathcal{A}}^{\text{mu-EUF-CMA}}$ be defined as in Figure 3. We define

$$\text{Adv}_S^{\text{mu-EUF-CMA}}(t, q_{\text{NW}}, q_{\text{SG}}, q_{\text{SG/U}}, q_{\text{C}}) := \max_{\mathcal{A}} \Pr \left[G_{S, \mathcal{A}}^{\text{mu-EUF-CMA}} \Rightarrow 1 \right],$$

where the maximum is taken over all adversaries, denoted $(t, q_{\text{NW}}, q_{\text{SG}}, q_{\text{SG/U}}, q_{\text{C}})$ -mu-EUF-CMA-adversaries, running in time at most t and making at most q_{NW} , q_{SG} , resp. q_{C} total queries to their NEW, SIGN, resp. CORRUPT oracle, and making at most $q_{\text{SG/U}}$ queries $\text{SIGN}(i, \cdot)$ for any user i .

Multi-user EUF-CMA security of signature schemes (with adaptive corruptions) can be reduced to classical, single-user EUF-CMA security (formally, $G_{S, \mathcal{A}}^{\text{mu-EUF-CMA}}$ with \mathcal{A} restricted to $q_{\text{NW}} = 1$ queries to NEW) by a standard hybrid argument, losing a factor of number of users. Formally, this yields

$$\text{Adv}_S^{\text{mu-EUF-CMA}}(t, q_{\text{NW}}, q_{\text{SG}}, q_{\text{SG/U}}, q_{\text{C}}) \leq q_{\text{NW}} \cdot \text{Adv}_S^{\text{mu-EUF-CMA}}(t', 1, q_{\text{SG/U}}, q_{\text{SG/U}}, 0),$$

where $t \approx t'$. (Note that the reduction is not affected by the total number of signature queries q_{SG} across all users.) In many cases, such loss is indeed unavoidable [5].

Definition 3.7 (MAC scheme). A MAC scheme $M = (\text{KGen}, \text{Tag}, \text{Vrfy})$ consists of three efficient algorithms defined as follows.

- $\text{KGen}() \xrightarrow{\$} K$. This probabilistic algorithm generates a key K .
- $\text{Tag}(K, m) \xrightarrow{\$} \tau$. On input a key K and a message m , this (possibly) probabilistic algorithm outputs a message authentication code (MAC) τ .
- $\text{Vrfy}(K, m, \tau) \rightarrow d$. On input a key K , a message m , and a MAC τ , this deterministic algorithm outputs a decision bit $d \in \{0, 1\}$ (where $d = 1$ indicates validity of the MAC).

Definition 3.8 (MAC mu-EUF-CMA security). Let M be a MAC scheme and $G_{M, \mathcal{A}}^{\text{mu-EUF-CMA}}$ be defined as in Figure 3. We define

$$\text{Adv}_M^{\text{mu-EUF-CMA}}(t, q_{\text{NW}}, q_{\text{TG}}, q_{\text{TG/U}}, q_{\text{VF}}, q_{\text{VF/U}}, q_{\text{C}}) := \max_{\mathcal{A}} \Pr \left[G_{M, \mathcal{A}}^{\text{mu-EUF-CMA}} \Rightarrow 1 \right],$$

where the maximum is taken over all adversaries, denoted $(t, q_{\text{NW}}, q_{\text{TG}}, q_{\text{TG/U}}, q_{\text{VF}}, q_{\text{VF/U}}, q_{\text{C}})$ -mu-EUF-CMA-adversaries, running in time at most t and making at most q_{NW} , q_{TG} , q_{VF} , resp. q_{C} queries to their NEW, SIGN, VRFY, resp. CORRUPT oracle, and making at most $q_{\text{TG/U}}$ queries $\text{TAG}(i, \cdot)$, resp. $q_{\text{VF/U}}$ queries $\text{VRFY}(i, \cdot)$ for any user i .

$G_{S,A}^{\text{mu-EUF-CMA}}$

INITIALIZE:

- 1 $Q \leftarrow \emptyset$
- 2 $\mathcal{C} \leftarrow \emptyset$
- 3 $u \leftarrow 0$

NEW:

- 4 $u \leftarrow u + 1$
- 5 $(pk_u, sk_u) \xleftarrow{\$} \text{KGen}()$
- 6 return pk_u

SIGN(i, m):

- 7 $\sigma \xleftarrow{\$} \text{Sign}(sk_i, m)$
- 8 $Q \leftarrow Q \cup \{(i, m)\}$
- 9 return σ

CORRUPT(i):

- 10 $\mathcal{C} \leftarrow \mathcal{C} \cup \{i\}$
- 11 return sk_i

FINALIZE(i^*, m^*, σ^*):

- 12 $d^* \leftarrow \text{Vrfy}(pk_{i^*}, m^*, \sigma^*)$
- 13 return $[[d^* = 1 \wedge i^* \notin \mathcal{C} \wedge (i^*, m^*) \notin Q]]$

$G_{M,A}^{\text{mu-EUF-CMA}}$

INITIALIZE:

- 1 $Q \leftarrow \emptyset$
- 2 $\mathcal{C} \leftarrow \emptyset$
- 3 $u \leftarrow 0$

NEW:

- 4 $u \leftarrow u + 1$
- 5 $K_u \xleftarrow{\$} \text{KGen}()$

TAG(i, m):

- 6 $\tau \xleftarrow{\$} \text{Tag}(K_i, m)$
- 7 $Q \leftarrow Q \cup \{(i, m)\}$
- 8 return τ

VERFY(i, m, τ):

- 9 $d \leftarrow \text{Vrfy}(K_i, m, \tau)$
- 10 return d

CORRUPT(i):

- 11 $\mathcal{C} \leftarrow \mathcal{C} \cup \{i\}$
- 12 return K_i

FINALIZE(i^*, m^*, τ^*):

- 13 $d^* \leftarrow \text{Vrfy}(K_{i^*}, m^*, \tau^*)$
- 14 return $[[d^* = 1 \wedge i^* \notin \mathcal{C} \wedge (i^*, m^*) \notin Q]]$

Figure 3: Multi-user existential unforgeability (mu-EUF-CMA) of signature schemes (left) and MAC schemes (right).

As for signature schemes, multi-user EUF-CMA security of MACs reduces to the single-user case ($q_{\text{NW}} = 1$) by a standard hybrid argument:

$$\begin{aligned} & \text{Adv}_{\text{M}}^{\text{mu-EUF-CMA}}(t, q_{\text{NW}}, q_{\text{TG}}, q_{\text{TG/U}}, q_{\text{VF}}, q_{\text{VF/U}}, q_{\text{C}}) \\ & \leq q_{\text{NW}} \cdot \text{Adv}_{\text{M}}^{\text{mu-EUF-CMA}}(t, 1, q_{\text{TG/U}}, q_{\text{TG/U}}, q_{\text{VF/U}}, q_{\text{VF/U}}, 0), \end{aligned}$$

where $t \approx t'$. (Note that the reduction is not affected by the total number of tagging and verification queries q_{TG} resp. q_{VF} across all users.)

Our multi-user definition of MACs provides a verification oracle, which is non-standard (and in general not equivalent to a definition with a single forgery attempts, as Bellare, Goldreich and Mityagin [9] showed). For PRF-based MACs (which in particular includes HMAC used in TLS 1.3), it however is equivalent and the reduction from multi-query to single-query verification is tight [9].

In our key exchange reductions, we can actually dispense with corruptions of MAC keys, i.e., work in the restricted setting of $q_{\text{C}} = 0$. This in particular allows specific constructions like AMAC [6] achieving tight multi-user security (without corruptions).

If we use a random oracle RO as PRF-like MAC with key length kl and output length ol , then

$$\text{Adv}_{\text{RO}}^{\text{mu-EUF-CMA}}(t, q_{\text{NW}}, q_{\text{TG}}, q_{\text{TG/U}}, q_{\text{VF}}, q_{\text{VF/U}}, q_{\text{C}}, q_{\text{RO}}) \leq \frac{q_{\text{VF}}}{2^{ol}} + \frac{(q_{\text{NW}} - q_{\text{C}}) \cdot q_{\text{RO}}}{2^{kl}}.$$

3.4 Hash Function Collision Resistance

Finally, let us define collision resistance of hash functions.

Definition 3.9 (Hash function collision resistance). *Let $H: \{0, 1\}^* \rightarrow \{0, 1\}^{ol}$ for $ol \in \mathbb{N}$ be a function. For a given adversary \mathcal{A} running in time at most t , we can consider*

$$\text{Adv}_H^{\text{CR}}(t) := \Pr [(m, m') \xleftarrow{\$} \mathcal{A} : m \neq m' \text{ and } H(m) = H(m')].$$

4 The SIGMA Protocol

The SIGMA family of key exchange protocols introduced by Krawczyk [37, 38] describes several variants for building authenticated Diffie–Hellman key exchange using the “SIGn-and-MAC” approach. Its design has been adopted in several Internet security protocols, including, e.g., the Internet Key Exchange protocol [32, 35] as part of the IPsec Internet security protocol [36] and the newest version 1.3 of the Transport Layer Security protocol [47].

Beyond the basic SIGMA design, we are particularly interested in the SIGMA-I variant which forms the basis of the TLS 1.3 key exchange and aims at hiding the protocol participants’ identities as additional feature. We here present an augmented version of the basic SIGMA/SIGMA-I protocols which includes explicit exchange of session-identifying random numbers (nonces) to be closer to SIGMA(-like) protocols in practice, somewhat following the “full fledged” SIGMA variant [38, Appendix B]. Figure 4 illustrates the (augmented) SIGMA and SIGMA-I protocol flows and Figure 5 formalizes both as key exchange protocols according to the syntax of Section 2.1.

The SIGMA and SIGMA-I protocols make use of a signature scheme $S = (\text{KGen}, \text{Sign}, \text{Vrfy})$, a MAC scheme $M = (\text{KGen}, \text{Tag}, \text{Vrfy})$, a pseudorandom function PRF, and a function RO which we model as a random oracle. The parties’ long-term secret keys consist of one signing key, i.e., $\text{KE.KGen} = S.\text{KGen}$. The protocols consists of three messages exchanged and accordingly two steps performed by both initiator and responder, which we describe in more detail now.

Initiator Step 1. The initiator picks a Diffie–Hellman exponent $x \xleftarrow{\$} \mathbb{Z}_q$ and a random nonce n_I of length nl and sends n_I and g^x .

Responder Step 1. The responder also picks a random DH exponent y and a random nonce n_R . It then derives a master key as $mk \leftarrow \text{RO}(n_I, n_R, X, Y, X^y)$ from nonces, DH shares, and the joint DH secret $g^{xy} = (g^x)^y$. From mk , keys are derived via PRF with distinct labels: the session key k_s , the MAC key k_t , and (only in SIGMA-I) the encryption key k_e .

The responder computes a signature σ with sk_R over nonces and DH shares (and a unique label L_{rs}) and a MAC value τ under key k_t over the nonces and its identity R (and unique label L_{rm}). It sends n_I , g^y , as well as R , σ , and τ to the initiator. In SIGMA-I the last three elements are encrypted using k_e to conceal the responder’s identity against passive adversaries.

Initiator Step 2. The initiator also computes mk and keys k_s , k_t , and (in SIGMA-I, used to decrypt the second message part) k_e . It ensures both the received signature σ and MAC τ verify, and aborts otherwise.

It computes its own signature σ' under sk_I on nonces and DH shares (with a different label L_{is}) and a MAC τ' under k_t over the nonces and its identity I (with yet another label L_{im}). It sends I , σ' , and τ' to the responder (in SIGMA-I encrypted under k_e) and accepts with session key k_s using the nonces $n_I || n_R$ as session identifier.

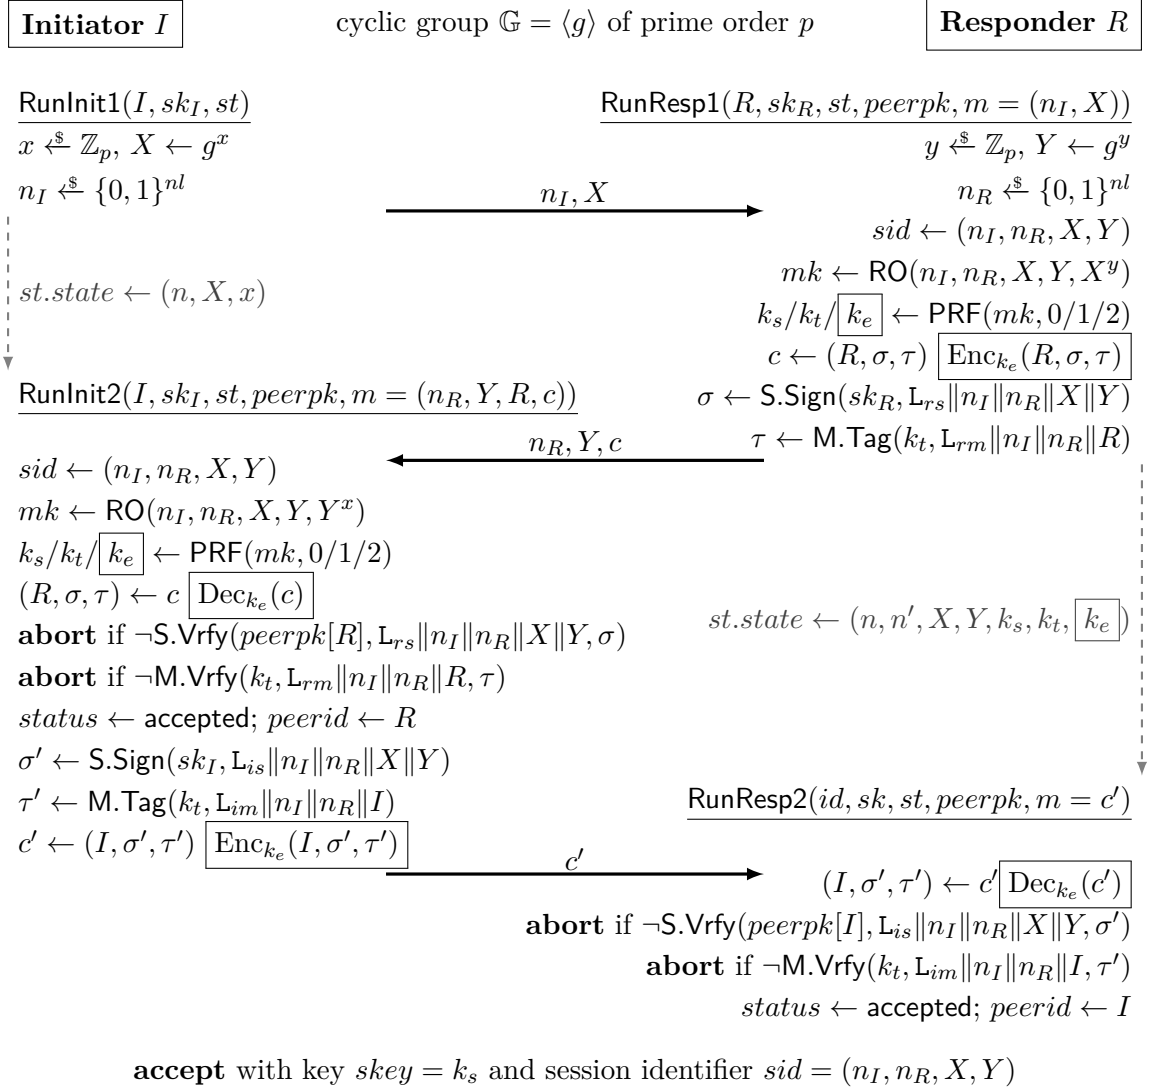


Figure 4: The SIGMA/SIGMA-I protocol flow diagram. Boxed code is only performed in the SIGMA-I variant. Values L_x indicate label strings (distinct per x).

Responder Step 2. The responder finally checks the initiator's signature σ' and MAC τ' (aborting if either fails) and then accepts with session key $skey = k_s$ and session identifier $sid = n_I \| n_R$.

5 Tighter Security Proof for SIGMA-I

We now come to our main result, a tighter security proof for the SIGMA-I protocol. Note that by omitting message encryption (cf. Figure 5) our proof similarly applies to the basic SIGMA protocol.

Theorem 5.1. *Let the SIGMA-I protocol be as specified in Figure 5 based on a group \mathbb{G} of prime order p , a PRF PRF, a signature scheme S, and a MAC M, and let RO in the protocol be modeled as a random oracle. For any $(t, q_N, q_S, q_{RS}, q_{RL}, q_T)$ -KE-SEC-adversary against SIGMA-I making at most q_{RO} queries to RO, we give algorithms $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$, and \mathcal{B}_4 in the proof, with running*

Activate($id, sk, peerid, peerpk, role$):

```

1  $st'.role \leftarrow role$ 
2  $st'.status \leftarrow running$ 
3 if  $role = initiator$  then
4    $(st', m') \leftarrow RunInit1(id, sk, st')$ 
5 else  $m' \leftarrow \perp$ 
6 return  $(st', m')$ 

```

Run($id, sk, st, peerpk, m$):

```

1 if  $st.status \neq running$  then
2 return  $\perp$ 
3 if  $st.role = initiator$  then
4    $(st', m') \leftarrow RunInit2(id, sk, st, peerpk, m)$ 
5 else if  $st.sid = \perp$ 
6    $(st', m') \leftarrow RunResp1(id, sk, st, peerpk, m)$ 
7 else
8    $(st', m') \leftarrow RunResp2(id, sk, st, peerpk, m)$ 
9 return  $st', m'$ 

```

RunInit1(id, sk, st):

```

1  $n \xleftarrow{\$} \{0, 1\}^{nl}$ 
2  $x \xleftarrow{\$} \mathbb{Z}_q$ 
3  $X \leftarrow g^x$ 
4  $st'.state \leftarrow (n, X, x)$ 
5  $m' \leftarrow (n, X)$ 
6 return  $(st', m')$ 

```

RunResp1($id, sk, st, peerpk, m$):

```

1  $(n, X) \leftarrow m$ 
2  $n' \xleftarrow{\$} \{0, 1\}^{nl}$ 
3  $y \xleftarrow{\$} \mathbb{Z}_q$ 
4  $Y \leftarrow g^y$ 
5  $st'.sid \leftarrow (n, n', X, Y)$ 
6  $\sigma \leftarrow S.Sign(sk, L_{rs} \| n \| n' \| X \| Y)$ 
7  $mk \leftarrow RO(n \| n' \| X \| Y \| X^y)$ 
8  $k_s \leftarrow PRF(mk, 0)$ 
9  $k_t \leftarrow PRF(mk, 1)$ 
10  $k_e \leftarrow PRF(mk, 2)$ 
11  $\tau \leftarrow M.Tag(k_t, L_{rm} \| n \| n' \| id)$ 
12  $st'.state \leftarrow (n, n', X, Y, k_s, k_t)$ 
    $st'.state \leftarrow (n, n', X, Y, k_s, k_t, k_e)$ 
13  $m' \leftarrow (n', Y, id, \sigma, \tau)$ 
    $m' \leftarrow (n', Y, Enc(k_e, (id, \sigma, \tau)))$ 

```

RunInit2($id, sk, st, peerpk, m$):

```

1  $(n', Y, peerid, \sigma, \tau) \leftarrow m$ 
    $(n', Y, c) \leftarrow m$ 
2  $(n, X, x) \leftarrow st.state$ 
3  $st'.sid \leftarrow (n, n', X, Y)$ 
4  $mk \leftarrow RO(n \| n' \| X \| Y \| Y^x)$ 
5  $k_s \leftarrow PRF(mk, 0)$ 
6  $k_t \leftarrow PRF(mk, 1)$ 
7  $k_e \leftarrow PRF(mk, 2)$ 
8  $(peerid, \sigma, \tau) \leftarrow Dec(k_e, c)$ 
9  $st'.peerid \leftarrow peerid$ 
10 if  $S.Vrfy(peerpk[peerid], L_{rs} \| n \| n' \| X \| Y, \sigma)$ 
   and  $M.Vrfy(k_t, L_{rm} \| n \| n' \| peerid, \tau)$  then
11    $st'.status \leftarrow accepted$ 
12    $st'.skey \leftarrow k_s$ 
13    $\sigma' \leftarrow S.Sign(sk, L_{is} \| n \| n' \| X \| Y)$ 
14    $\tau' \leftarrow M.Tag(k_t, L_{im} \| n \| n' \| id)$ 
15    $m' \leftarrow (id, \sigma', \tau')$ 
    $m' \leftarrow Enc(k_e, (id, \sigma', \tau'))$ 
16 else
17    $m' \leftarrow \perp$ 
18    $st'.status \leftarrow rejected$ 

```

RunResp2($id, sk, st, peerpk, m$):

```

1  $(n, n', X, Y, k_s, k_t) \leftarrow st.state$ 
    $(n, n', X, Y, k_s, k_t, k_e) \leftarrow st.state$ 
2  $(peerid, \sigma', \tau') \leftarrow m$ 
    $(peerid, \sigma', \tau') \leftarrow Dec(k_e, m)$ 
3  $st'.peerid \leftarrow peerid$ 
4 if  $S.Vrfy(peerpk[peerid], L_{is} \| n \| n' \| X \| Y, \sigma')$ 
   and  $M.Vrfy(k_t, L_{im} \| n \| n' \| peerid, \tau')$  then
5    $st'.status \leftarrow accepted$ 
6    $st'.skey \leftarrow k_s$ 
7 else  $st'.status \leftarrow rejected$ 
8  $m' \leftarrow \varepsilon$ 
9 return  $st', m'$ 

```

Figure 5: The formalized SIGMA/SIGMA-I key exchange protocols (cf. Section 2.1). Boxed code is only performed in the SIGMA-I variant.

times $t_{\mathcal{B}_1} \approx t + 2q_{\text{RO}} \log_2 p$ and $t_{\mathcal{B}_i} \approx t$ (for $i = 2, \dots, 4$) close to that of \mathcal{A} , such that

$$\begin{aligned} & \text{Adv}_{\text{SIGMA-I}}^{\text{KE-SEC}}(t, q_N, q_S, q_{\text{RS}}, q_{\text{RL}}, q_T) \\ & \leq \frac{q_S^2}{2^{nl+1}} + \frac{q_S^2}{2p} + \frac{q_S^2}{2^{nl} \cdot p} + \text{Adv}_{\mathbb{G}}^{\text{stDH}}(t_{\mathcal{B}_1}, q_{\text{RO}}) + \text{Adv}_{\text{PRF}}^{\text{mu-PRF}}(t_{\mathcal{B}_2}, q_S, 3q_S, 3) \\ & \quad + \text{Adv}_{\text{S}}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_3}, q_N, q_S, q_S, q_{\text{RL}}) + \text{Adv}_{\text{M}}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_4}, q_S, q_S, 1, q_S, 1, 0), \end{aligned}$$

where nl is the nonce length in SIGMA-I and \mathbb{G} is the used Diffie–Hellman group of prime order p .

In terms of single-user security for the employed primitives, multi-user PRF and MAC security can be obtained tightly, e.g., via the efficient AMAC construction [6], and multi-user signature security can be generically reduced to single-user security of any signature scheme with a loss in the number of users, here parties (not sessions) in the key exchange game.

Proof. Our proof of key exchange security for SIGMA-I proceeds via a sequence of code-based games [13]. For the first half, the proof conceptually follows the strategy put forward by Cohn-Gordon et al. [21].

Game 0. The initial game, G_0 , is the key exchange security game played by \mathcal{A} (cf. Figure 1), using the KGen, Activate, and Run routines of SIGMA-I defined in Figure 5. Therefore,

$$\Pr[G_0 \Rightarrow 1] = \Pr[G_{\text{KE}, \mathcal{A}}^{\text{KE-SEC}} \Rightarrow 1].$$

Game 1. Between G_0 and G_1 (Figure 6), we make internal changes to the record-keeping of the game, namely we track the nonces and group elements chosen and received by honest sessions. Whenever two honest sessions pick the same nonce or group element, we set a flag bad_C . Whenever an honest responder session picks a nonce and group element that has already been received by an initiator session, we set a flag bad_O . This change is unobservable by the adversary, hence

$$\Pr[G_0 \Rightarrow 1] = \Pr[G_1 \Rightarrow 1].$$

Game 2. In Game G_2 (Figure 6), we abort whenever nonces or group elements collide among honest sessions (i.e., the bad_C or the bad_O flag is set). By the identical-until-bad lemma [13],

$$\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1] \leq \Pr[\text{bad}_C \text{ or } \text{bad}_O \leftarrow \text{true in } G_1].$$

Since both nonces and group elements are chosen at random, by the union bound and the birthday bound, we have that, for m calls to RunInit1 and RunResp1 , $\Pr[\text{bad}_C \leftarrow \text{true in } G_1] \leq \frac{m^2}{2^{nl+1}} + \frac{m^2}{2p}$ (where nl is the nonce length and p the order of the Diffie–Hellman group). There are at most q_S pairs received by initiator sessions, so the probability that a responder session randomly chooses one of these pairs is at most $\frac{q_S}{2^{nl} \cdot p}$; then by the union bound we have that $\Pr[\text{bad}_O \leftarrow \text{true in } G_1] \leq \frac{q_S^2}{2^{nl} \cdot p}$. Since each of RunInit1 and RunResp1 is called at most once per SEND query, if an adversary makes q_S queries to its SEND oracle, then

$$\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1] \leq \frac{q_S^2}{2^{nl+1}} + \frac{q_S^2}{2p} + \frac{q_S^2}{2^{nl} \cdot p}.$$

In all subsequent games, we are now sure that each honest session chooses a unique nonce and group element. Since the session identifier $\text{sid} = (n_I, n_R, X, Y)$ contains exactly one initiator and

$G_1, \boxed{G_2}$

RunInit1(id, sk, st):

```

1  $n \xleftarrow{\$} \{0, 1\}^{nl}$ 
2 if  $n \in N$  then  $\text{bad}_C \leftarrow \text{true} \boxed{; \text{abort}}$ 
3  $N \leftarrow N \cup \{n\}$ 
4  $x \xleftarrow{\$} \mathbb{Z}_q$ 
5  $X \leftarrow g^x$ 
6 if  $X \in E$  then  $\text{bad}_C \leftarrow \text{true} \boxed{; \text{abort}}$ 
7  $E \leftarrow E \cup \{X\}$ 
8  $st'.state \leftarrow (n, X, x)$ 
9  $m' \leftarrow n \| X$ 
10 return ( $st', m'$ )

```

RunInit2($id, sk, st, peerpk, m$):

```

11  $(n', Y, c) \leftarrow m$ 
12  $Recv \leftarrow Recv \cup \{n', Y\}$ 
13  $(n, X, x) \leftarrow st.state$ 
14 ...

```

RunResp1($id, sk, st, peerpk, m$):

```

15  $(n, X) \leftarrow m$ 
16  $n' \xleftarrow{\$} \{0, 1\}^{nl}$ 
17 if  $n' \in N$  then  $\text{bad}_C \leftarrow \text{true} \boxed{; \text{abort}}$ 
18  $N \leftarrow N \cup \{n'\}$ 
19  $y \xleftarrow{\$} \mathbb{Z}_q$ 
20  $Y \leftarrow g^y$ 
21 if  $n', Y \in Recv$  then  $\text{bad}_O \leftarrow \text{true} \boxed{; \text{abort}}$ 
22  $E \leftarrow E \cup \{Y\}$ 
23 if  $Y \in E$  then  $\text{bad}_C \leftarrow \text{true} \boxed{; \text{abort}}$ 
24  $st'.sid \leftarrow (n, n', X, Y)$ 
25 ...

```

RO(m):

```

101 if  $H[m] = \perp$  then  $H[m] \xleftarrow{\$} \{0, 1\}^{kl}$ 
102 return  $H[m]$ 

```

Figure 6: Games G_1 (changes highlighted in gray) and G_2 (changes highlighted in $\boxed{\text{frames}}$) of the SIGMA-I proof; with the explicit (lazy-sampled) random oracle RO.

one responder nonce, this furthermore implies that when two honest sessions are partnered, they must have different roles.

Game 3. In Game G_3 (Figure 7), we remove the now superfluous collision flags bad_C and bad_O and add additional bookkeeping. All honest initiator sessions now log their outgoing messages in an internal table Sent . Honest responder sessions use this table to check if the message they received was sent by an honest initiator session. If so, they log their keys k_t , k_e , and k_s in a second internal table, \mathbf{S} , indexed by their session identifier. These changes are unobservable by the adversary, so

$$\Pr[G_2 \Rightarrow 1] = \Pr[G_3 \Rightarrow 1].$$

Game 4. In Game G_4 (Figure 7), we require that initiator sessions whose key material has already been computed by an honest partner session simply copy their partners' key material. When an honest initiator session π_u^i with nonce n and group element X receives a message $m \leftarrow (n', Y, c)$, it sets its session identifier $sid \leftarrow (n, n', X, Y)$. It then checks if $\mathbf{S}[sid] \neq \perp$ (which is only the case if π_u^i has an honest partner), and if so uses the stored key material $k_s, k_t, k_e \leftarrow \mathbf{S}[st'.sid]$.

Recall that both partnered sessions agree on the DH shares X and Y as components of sid . They hence also agree on the shared DH secret $Z = g^{xy}$ and thus on the master key derived as $\text{RO}(n \| n' \| X \| Y \| Z)$ as well as the derived key k_s, k_t , and k_e . For the adversary \mathcal{A} it is hence unobservable if initiators with honest partner actually compute their keys themselves or copy their partners' key material in Game G_4 , so

$$\Pr[G_3 \Rightarrow 1] = \Pr[G_4 \Rightarrow 1].$$

Game 5. In Game G_5 (Figure 8), all honest sessions sample their master keys uniformly at random (as long as the random oracle has not been queried on the corresponding input) and

$G_3, \boxed{G_4}$

RunInit1(id, sk, st):

```

1  $n \xleftarrow{\$} \{0, 1\}^{nl}$ 
2 if  $n \in N$  then abort
3  $N \leftarrow N \cup \{n\}$ 
4  $x \xleftarrow{\$} \mathbb{Z}_q$ 
5  $X \leftarrow g^x$ 
6 if  $X \in E$  then abort
7  $E \leftarrow E \cup \{X\}$ 
8  $st'.state \leftarrow (n, X, x)$ 
9  $m' \leftarrow n \| X$ 
10 Sent  $\leftarrow$  Sent  $\cup m'$ 
11 return ( $st', m'$ )

```

RunInit2($id, sk, st, peerpk, m$):

```

12 ( $n', Y, c$ )  $\leftarrow$   $m$ 
13  $Recv \leftarrow Recv \cup \{n', Y\}$ 
14 ( $n, X, x$ )  $\leftarrow$   $st.state$ 
15  $st'.sid \leftarrow (n, n', X, Y)$ 
16 if S[ $st'.sid$ ]  $\neq \perp$  then
17    $mk \leftarrow RO(n \| n' \| X \| Y \| Y^x)$ 
18    $k_s \leftarrow PRF(mk, 0)$ 
19    $k_t \leftarrow PRF(mk, 1)$ 
20    $k_e \leftarrow PRF(mk, 2)$ 
21    $\boxed{k_s, k_t, k_e \leftarrow S[st'.sid]}$ 
22 else
23    $mk \leftarrow RO(n \| n' \| X \| Y \| Y^x)$ 
24    $k_s \leftarrow PRF(mk, 0)$ 
25    $k_t \leftarrow PRF(mk, 1)$ 
26    $k_e \leftarrow PRF(mk, 2)$ 
27 ( $peerid, \sigma, \tau$ )  $\leftarrow$  Dec( $k_e, c$ )
28  $st'.peerid \leftarrow peerid$ 
29 ...

```

RunResp1($id, sk, st, peerpk, m$):

```

30 ( $n, X$ )  $\leftarrow$   $m$ 
31  $n' \xleftarrow{\$} \{0, 1\}^{nl}$ 
32 if  $n \in N$  then abort
33  $N \leftarrow N \cup \{n\}$ 
34  $y \xleftarrow{\$} \mathbb{Z}_q$ 
35  $Y \leftarrow g^y$ 
36 if  $Y \in E$  then abort
37  $E \leftarrow E \cup \{Y\}$ 
38 if  $n', Y \in Recv$  then abort
39  $st'.sid \leftarrow (n, n', X, Y)$ 
40  $\sigma \leftarrow S.Sign(sk, L_{rs} \| n \| n' \| X \| Y)$ 
41  $mk \leftarrow RO(n \| n' \| X \| Y \| X^y)$ 
42  $k_s \leftarrow PRF(mk, 0)$ 
43  $k_t \leftarrow PRF(mk, 1)$ 
44  $k_e \leftarrow PRF(mk, 2)$ 
45 if  $m \in$  Sent then
46   S[ $st'.sid$ ]  $\leftarrow$  ( $k_s, k_t, k_e$ )
47  $\tau \leftarrow M.Tag(k_t, L_{rm} \| n \| n' \| id)$ 
48  $st'.state \leftarrow (n, n', X, Y, k_s, k_t)$ 
49  $m' \leftarrow (n', Y, Enc(k_e, (id, \sigma, \tau)))$ 
50 return ( $st', m'$ )

```

Figure 7: Games G_3 (changes highlighted in gray) and G_4 (changes highlighted in frames) of the SIGMA-I proof.

program the random oracle to that value (through RO's internal table $H[n\|n'\|X\|Y\|Y^x] \leftarrow mk$). This is equivalent to RO performing the same checks and uniform sampling, and hence undetectable for \mathcal{A} :

$$\Pr[G_4 \Rightarrow 1] = \Pr[G_5 \Rightarrow 1].$$

Game 6. In Game G_6 (Figure 8), responder sessions whose first message came from an honest initiator stop programming the random oracle on their uniformly chosen master key mk . This is undetectable for adversary \mathcal{A} unless it makes a query $\text{RO}(n\|n'\|X\|Y\|Z)$, where $sid = (n, n', X, Y)$ is the session identifier shared by two honest partnered sessions, and Z is the Diffie–Hellman secret corresponding to the pair (X, Y) . We call this event F , and bound the probability of F by giving a reduction \mathcal{B}_1 (specified in Figure 9) to the strong Diffie–Hellman assumption in the DH group \mathbb{G} . The reduction makes at most as many queries to its stDH oracle as \mathcal{A} makes to its RO oracle, as follows.

Given its strong DH challenge $(A = g^a, B = g^b)$ and having access to an oracle $\text{stDH}_a(U, V)$ which outputs 1 if $U^a = V$ and 0 otherwise, \mathcal{B}_1 simulates G_6 for an adversary \mathcal{A} as follows. In all honest initiator sessions, \mathcal{B}_1 embeds its challenge into the sent DH value as $X \leftarrow A \cdot g^r$, where $r \in \mathbb{Z}_q$ is sampled uniformly at random for each session. Furthermore, in all responder sessions receiving their first message from an honest initiator, \mathcal{B}_1 embeds its challenge as $Y \leftarrow B \cdot g^{r'}$, where $r' \in \mathbb{Z}_q$ is sampled uniformly at random for each session.

Let us first observe that if event F occurs, then the value Z in the random oracle query $\text{RO}(n\|n'\|X\|Y\|Z)$ will equal $g^{(a+r)(b+r')}$ for some $r, r' \in \mathbb{Z}_q$ chosen by \mathcal{B}_1 , and consequently

$$Z \cdot Y^{-r} = g^{(a+r)(b+r') - (b+r') \cdot r} = g^{a(b+r')} = Y^a.$$

This equality can be tested for by \mathcal{B}_1 by calling its stDH_a oracle on the pair $(Y, Z \cdot Y^{-r})$. We let \mathcal{B}_1 do so whenever \mathcal{A} queries RO on some value $(n\|n'\|X\|Y\|Z)$ where $(n, X = A \cdot g^r)$ was output by an honest initiator session and $(n', Y = g^{(b+r')})$ was output by a responder session with an honest initiator; the responder stores (n, n', X, Y) in a look-up table \mathbf{Q} so this can be checked efficiently. If $\text{stDH}_a(Y, Z \cdot Y^{-r}) = 1$ on such occasion, i.e., event F occurs, \mathcal{B}_1 stops with output $Z \cdot Y^{-r} \cdot A^{-r'} = g^{(a+r)(b+r')} \cdot g^{-(b+r') \cdot r} \cdot g^{-ar'} = g^{ab}$ and wins. Therefore,

$$\Pr[F] \leq \text{Adv}_{\mathbb{G}, \mathcal{B}_1}^{\text{stDH}}.$$

One subtlety in this step is ensuring that \mathcal{B}_1 can correctly simulate answers to REVSESSIONKEY queries to any initiator or responder session. We do so by accordingly programming the random oracle on the sampled master key, where needed. First of all observe that responder sessions without honest initiator keep picking their own Y share and compute mk regularly. Initiator and responder sessions with honest partner have the challenge embedded and sample an independent master key which is not programmed to the random oracle. However, \mathcal{B}_1 stops and wins (as described above) if \mathcal{A} ever queries the random oracle on the correct DH secret; i.e., \mathcal{A} never sees the (inconsistent) random oracle output for these master keys. The interesting case is when an initiator session (which always embeds the challenge in its DH share as $X = A \cdot g^r$) obtains a message (n', Y, c) *not* originating from an honest responder: Here, Y may well have been picked by the adversary who could furthermore have corrupted the initiator's peer and hence make the initiator accept—with a master key it cannot compute itself.

We therefore let \mathcal{B}_1 attempt to copy the adversary's master key, if it has been computed. The RO oracle logs all queries it receives by their putative session id (n, n', X, Y) in a look-up table H' ,

$G_5, \boxed{G_6}$

RunInit2($id, sk, st, peerpk, m$):

```

1  ( $n', Y, c$ )  $\leftarrow m$ 
2   $Recv \leftarrow Recv \cup \{n', Y\}$ 
3  ( $n, X, x$ )  $\leftarrow st.state$ 
4   $st'.sid \leftarrow (n, n', X, Y)$ 
5  if  $S[st'.sid] \neq \perp$  then
6     $k_s, k_t, k_e \leftarrow S[st'.sid]$ 
7  else
8     $mk \xleftarrow{\$} \{0, 1\}^{kl}$ 
9    if  $H[n\|n'\|X\|Y\|Y^x] \neq \perp$ 
10      $mk \leftarrow H[n\|n'\|X\|Y\|Y^x]$ 
11      $H[n\|n'\|X\|Y\|Y^x] \leftarrow mk$ 
12      $k_s \leftarrow PRF(mk, 0)$ 
13      $k_t \leftarrow PRF(mk, 1)$ 
14      $k_e \leftarrow PRF(mk, 2)$ 
15  ( $peerid, \sigma, \tau$ )  $\leftarrow Dec(k_e, c)$ 
16   $st'.peerid \leftarrow peerid$ 
17  if  $S.Vrfy(peerpk[peerid], L_{rs}\|n\|n'\|X\|Y, \sigma)$ 
    and  $M.Vrfy(k_t, L_{rm}\|n\|n'\|peerid)$  then
18     $st'.status \leftarrow accepted$ 
19     $st'.skey \leftarrow k_s$ 
20     $\sigma' \leftarrow S.Sign(sk, L_{is}\|n\|n'\|X\|Y)$ 
21     $\tau' \leftarrow M.Tag(k_t, L_{im}\|n\|n'\|id)$ 
22     $m' \leftarrow Enc(k_e, (id, \sigma', \tau'))$ 
23  else
24     $m' \leftarrow \perp$  ;  $st'.status \leftarrow rejected$ 

```

RunResp1($id, sk, st, peerpk, m$):

```

25  ( $n, X$ )  $\leftarrow m$ 
26   $n' \xleftarrow{\$} \{0, 1\}^{nl}$ 
27  if  $n \in N$  then  $bad \leftarrow true$ ; abort
28   $N \leftarrow N \cup \{n\}$ 
29   $y \xleftarrow{\$} \mathbb{Z}_q$ 
30   $Y \leftarrow g^x$ 
31  if  $Y \in E$  then  $bad \leftarrow true$ ; abort
32   $E \leftarrow E \cup \{Y\}$ 
33  if  $n', Y \in Recv$  then abort
34   $st'.sid \leftarrow (n, n', X, Y)$ 
35   $\sigma \leftarrow S.Sign(sk, L_{rs}\|n\|n'\|X\|Y)$ 
36   $mk \xleftarrow{\$} \{0, 1\}^{kl}$ 
37  if  $m \notin Sent$  then
38    if  $H[n\|n'\|X\|Y\|X^y] \neq \perp$ 
39      $mk \leftarrow H[n\|n'\|X\|Y\|X^y]$ 
40      $H[n\|n'\|X\|Y\|X^y] \leftarrow mk$ 
41     $k_s \leftarrow PRF(mk, 0)$ 
42     $k_t \leftarrow PRF(mk, 1)$ 
43     $k_e \leftarrow PRF(mk, 2)$ 
44    if  $m \in Sent$  then
45      $S[st'.sid] \leftarrow (k_s, k_t, k_e)$ 
46     $\tau \leftarrow M.Tag(k_t, L_{rm}\|n\|n'\|id)$ 
47     $st'.state \leftarrow (n, n', X, Y, k_s, k_t)$ 
48     $m' \leftarrow (n', Y, id, \sigma, \tau)$ 

```

Figure 8: Games G_5 (changes highlighted in gray) and G_6 (changes highlighted in frames) of the SIGMA-I proof.

```

 $\mathcal{B}_1(A, B)^{\text{stDH}_a(\cdot, \cdot)}$ 
RunInit1( $id, sk, st$ ):
1  $n \xleftarrow{\$} \{0, 1\}^{nl}$ 
2 if  $n \in N$  then abort
3  $N \leftarrow N \cup \{n\}$ 
4  $r \xleftarrow{\$} \mathbb{Z}_q$ 
5  $X \leftarrow A \cdot g^r$ 
6 if  $X \in E$  then abort
7  $E \leftarrow E \cup \{X\}$ 
8  $st'.state \leftarrow (n, X, r)$ 
9  $m' \leftarrow n \| X$ 
10  $\text{Sent}[m'] \leftarrow x$ 
11 return  $(st', m')$ 

RunInit2( $id, sk, st, \text{peerpk}, m$ ):
12  $(n', Y, c) \leftarrow m$ 
13  $\text{Recv} \leftarrow \text{Recv} \cup \{n', Y\}$ 
14  $(n, X, r) \leftarrow st.state$ 
15  $st'.sid \leftarrow (n, n', X, Y)$ 
16 if  $\mathcal{S}[st'.sid] \neq \perp$  then
17    $k_s, k_t, k_e \leftarrow \mathcal{S}[st'.sid]$ 
18 else
19    $mk \xleftarrow{\$} \{0, 1\}^{kl}$ 
20   for each  $Z \in H'[n \| n' \| X \| Y]$ 
21     if  $\text{stDH}_a(Y, Z \cdot Y^{-r}) = 1$  then
22        $mk \leftarrow H[n \| n' \| X \| Y \| Z]$ 
23    $\mathcal{Q}[st'.sid] \leftarrow (r, \perp, mk)$ 
24    $k_s \leftarrow \text{PRF}(mk, 0)$ 
25    $k_t \leftarrow \text{PRF}(mk, 1)$ 
26    $k_e \leftarrow \text{PRF}(mk, 2)$ 
27    $(\text{peerid}, \sigma, \tau) \leftarrow \text{Dec}(k_e, c)$ 
28    $st'.\text{peerid} \leftarrow \text{peerid}$ 
29   if  $\mathcal{S}.\text{Vrfy}(\text{L}_{rs} \| n \| n' \| X \| Y, \sigma, pk_{\text{peerid}})$ 
    and  $\mathcal{M}.\text{Vrfy}(k_t, \text{L}_{rm} \| n \| n' \| \text{peerid})$  then
30      $st'.status \leftarrow \text{accepted}$ 
31      $st'.skey \leftarrow k_s$ 
32      $\sigma' \leftarrow \mathcal{S}.\text{Sign}(sk, \text{L}_{is} \| n \| n' \| R \| W)$ 
33      $\tau' \leftarrow \mathcal{M}.\text{Tag}(k_t, \text{L}_{im} \| n \| n' \| id)$ 
34      $m' \leftarrow \text{Enc}(k_e, (id, \sigma', \tau'))$ 
35   else
36      $m' \leftarrow \perp$ 
37      $st'.status \leftarrow \text{rejected}$ 

RunResp1( $id, sk, st, \text{peerpk}, m$ ):
38  $(n, X) \leftarrow m$ 
39  $n' \xleftarrow{\$} \{0, 1\}^{nl}$ 
40 if  $n' \in N$  then abort
41  $N \leftarrow N \cup \{n'\}$ 
42  $r' \xleftarrow{\$} \mathbb{Z}_q$ 
43  $mk \xleftarrow{\$} \{0, 1\}^{kl}$ 
44 if  $m \in \text{Sent}$  then
45    $r \leftarrow \text{Sent}[m]$ 
46    $Y \leftarrow B \cdot g^{r'}$ 
47   if  $n', Y \in \text{Recv}$  then abort
48    $st'.sid \leftarrow (n, n', X, Y)$ 
49   for each  $Z \in H'[n \| n' \| X \| Y]$ 
50     if  $\text{stDH}_a(Y, Z \cdot Y^{-r}) = 1$  then
51        $\text{FINALIZE}(Z \cdot Y^{-r} \cdot A^{-r'})$ 
52    $\mathcal{Q}[st'.sid] \leftarrow (r, r', mk)$ 
53 else
54    $Y \leftarrow g^{r'}$ 
55    $st'.sid \leftarrow (n, n', X, Y)$ 
56   if  $H[n \| n' \| X \| Y \| X^y] \neq \perp$ 
57      $mk \leftarrow H[n \| n' \| X \| Y \| X^{r'}]$ 
58    $H[n \| n' \| X \| Y \| X^y] \leftarrow mk$ 
59   if  $Y \in E$  then abort
60    $E \leftarrow E \cup \{Y\}$ 
61    $\sigma \leftarrow \mathcal{S}.\text{Sign}(sk, \text{L}_{rs} \| n \| n' \| X \| Y)$ 
62    $k_s \leftarrow \text{PRF}(mk, 0)$ 
63    $k_t \leftarrow \text{PRF}(mk, 1)$ 
64    $k_e \leftarrow \text{PRF}(mk, 2)$ 
65   if  $m \in \text{Sent}$  then
66      $\mathcal{S}[st'.sid] \leftarrow (k_s, k_t, k_e)$ 
67      $\tau \leftarrow \mathcal{M}.\text{Tag}(k_t, \text{L}_{rm} \| n \| n' \| id)$ 
68      $st'.state \leftarrow (n, n', X, Y, k_s, k_t, k_e)$ 
69      $m' \leftarrow (n', Y, \text{Enc}(k_e, (id, \sigma, \tau)))$ 

RO( $m$ ):
70 if  $H[m] = \perp$  then
71    $H[m] \xleftarrow{\$} \{0, 1\}^{kl}$ 
72   parse  $n \| n' \| X \| Y \| Z \leftarrow m$ 
73    $H'[n \| n' \| X \| Y] \leftarrow H'[n \| n' \| X \| Y] \cup \{Z\}$ 
74   if  $\mathcal{Q}[(n, n', X, Y)] \neq \perp$  then
75      $(r, r', mk) \leftarrow \mathcal{Q}[n, n', X, Y]$ 
76     if  $\text{stDH}_a(Y, Z \cdot Y^{-r}) = 1$  then
77       if  $r' = \perp$  then  $H[m] \leftarrow mk$ 
78       else  $\text{FINALIZE}(Z \cdot Y^{-r} \cdot A^{-r'})$ 
79   return  $H[m]$ 

```

Figure 9: Reduction \mathcal{B}_1 to the strong Diffie–Hellman assumption of the SIGMA-I proof. Sections highlighted in gray have been significantly altered compared to Game G_6 .

so \mathcal{B}_1 can efficiently access all Z such that (n, n', X, Y, Z) has been queried to RO. Since the DH secret corresponding to the pair (X, Y) equals Y^{a+r} , if Z is this DH secret, then

$$Z \cdot Y^{-r} = Y^{(a+r)-r} = Y^a.$$

The reduction can check this equality using its stDH_a oracle and in that case use the response to $\text{RO}(n, n', X, Y, Z)$ as mk . Otherwise, \mathcal{B}_1 samples mk at random and stores it in the table \mathbf{Q} (Line 52 of Figure 9), indicating it should be programmed in the random oracle later if queried on a matching Z value (Line 77). This ensures all responses to REVSESSIONKEY queries are consistent with \mathcal{A} 's queries to the random oracle RO .

Observe that, in all this, \mathcal{B}_1 calls its stDH oracle at most once for each entry $H[n\|n'\|X\|Y\|Z] = mk$ in the RO table H . In RO , stDH is called (once) only for entries that were not present when $\mathbf{Q}[(n, n', X, Y)]$ was set, then H' is set. In RunInit2 and RunResp1 , stDH is called only for matching H' entries established prior to setting \mathbf{Q} . Therefore, if stDH is called in RO for an entry, it was not called in either RunInit2 or RunResp1 . If stDH is called on an entry in RunResp1 , then the responder session is partnered, so its partner will copy its keys in RunInit2 and not call stDH . Furthermore, due to uniqueness of nonces and DH shares (by Game \mathbf{G}_2), no RunInit2 or RunResp1 call makes stDH be invoked twice for the same H' entry.

Since the total time to iterate through the for loops over all Run and RO queries is at most $O(q_{\text{RO}})$, the running time of \mathcal{B}_1 is roughly that of \mathcal{A} , plus the time needed to compute the arguments of the stDH queries. Each of these arguments requires one group operation and one exponentiation. (All other operations performed by \mathcal{B}_1 add only a small constant amount of time per SEND query, which is dominated by the runtime of \mathcal{A} .) The exponentiation can be computed using $2 \log_2 p$ group operations using the square-and-multiply (or double-and-add) algorithm, so $t_{\mathcal{B}_1} \approx t + 2q_{\text{RO}} \log_2 p$. The runtime t of \mathcal{A} already includes the computation of $2q_{\text{S}} \log_2 p$ group operations, so this is a significant but not prohibitive increase in runtime.

Having \mathcal{B}_1 perfectly simulate Game \mathbf{G}_5 for \mathcal{A} up to the point when F happens, and \mathbf{G}_6 and \mathbf{G}_5 differing only when F happens, we have

$$\Pr[\mathbf{G}_5 \Rightarrow 1] = \Pr[\mathbf{G}_6 \Rightarrow 1] + \Pr[F] \leq \Pr[\mathbf{G}_6 \Rightarrow 1] + \text{Adv}_{\mathbb{G}}^{\text{stDH}}(t_{\mathcal{B}_1}, q_{\text{RO}}),$$

and $t_{\mathcal{B}_1} \approx t + 2q_{\text{RO}} \log_2 p$.

Game 7. In Game \mathbf{G}_7 (Figure 10), responder oracles responding to honest messages samples session, MAC, and encryption keys k_s , k_t , and k_e randomly instead of computing them through a PRF. (Initiator oracles partnered with an honest responder will continue to copy those, now randomly sampled keys.)

Since the PRF key mk in this case is sampled independently of the random oracle and the rest of the game, this reduces straightforwardly to the multi-user security of the PRF via the reduction \mathcal{B}_2 we give in Figure 10. The adversary \mathcal{B}_2 makes one NEW and two FUNC queries for each RunResp1 query, or three FUNC queries in SIGMA-I . Notably, it makes at most three FUNC queries per user, and no CORRUPT queries because mk is never revealed to the adversary. Outside of the oracle calls, its running time exactly equals that of \mathcal{A} in Game \mathbf{G}_6 , as their pseudocode is identical, so $t_{\mathcal{B}_2} \approx t$. Using its FN oracle of the PRF game, \mathcal{B}_2 perfectly simulates \mathbf{G}_6 if the oracle gives real-PRF answers and \mathbf{G}_7 if it returns uniformly random values. Therefore,

$$\Pr[\mathbf{G}_6 \Rightarrow 1] \leq \Pr[\mathbf{G}_7 \Rightarrow 1] + \text{Adv}_{\text{PRF}}^{\text{mu-PRF}}(t_{\mathcal{B}_2}, q_{\text{S}}, 3q_{\text{S}}, 3, 0).$$

Observe that from now on, session and MAC keys of responder oracles that received honest initiator's messages are chosen independently at random, and that initiator oracles with matching sid will copy those keys. Notably, this is the case even for sessions whose (own or peer's) long-term secret have been revealed to the adversary. We will use these properties in the following to argue authentication of sessions as well as forward security of the session keys.

G₇

RunResp1(*id, sk, st, peerpk, m*):

```
1 (n, X) ← m
2 n' ←s {0, 1}nl
3 if n ∈ N then bad ← true; abort
4 N ← N ∪ {n}
5 y ←s ℤq
6 Y ← gy
7 if Y ∈ E then bad ← true; abort
8 E ← E ∪ {Y}
9 if n', Y ∈ Recv then abort
10 st'.sid ← (n, n', X, Y)
11 σ ← S.Sign(sk, Lrs||n||n'||X||Y)
12 mk ←s {0, 1}kl
13 if m ∉ Sent then
14   if H[Yx||n||n'||X||Y] ≠ ⊥
15     mk ← H[Yx||n||n'||X||Y]
16 ks ← PRF(mk, 0)
17 kt ← PRF(mk, 1)
18 ke ← PRF(mk, 2)
19 if m ∈ Sent
20   ks ←s {0, 1}kl
21   kt ←s {0, 1}kl
22   ke ←s {0, 1}kl
23   S[st'.sid] ← (ks, kt, ke)
24 τ ← M.Tag(kt, Lrm||n||n'||id)
25 st'.state ← (n, n', X, Y, ks, kt)
26 m' ← (n', Y, Enc(ke, (id, σ, τ)))
```

B₂^{FN(·,·)}

RunResp1(*id, sk, st, peerpk, m*):

```
1 (n, X) ← m
2 n' ←s {0, 1}nl
3 if n ∈ N then bad ← true; abort
4 N ← N ∪ {n}
5 y ←s ℤq
6 Y ← gy
7 if Y ∈ E then bad ← true; abort
8 E ← E ∪ {Y}
9 if n', Y ∈ Recv then abort
10 st'.sid ← (n, n', X, Y)
11 σ ← S.Sign(sk, Lrs||n||n'||X||Y)
12 mk ←s {0, 1}kl
13 if m ∉ Sent then
14   if H[n||n'||X||Y||Xy] ≠ ⊥
15     mk ← H[n||n'||X||Y||Xy]
16 ks ← PRF(mk, 0)
17 kt ← PRF(mk, 1)
18 ke ← PRF(mk, 2).
19 if m ∈ Sent
20   NEW(); i ++
21   ks ← FN(i, 0)
22   kt ← FN(i, 1)
23   ke ← FN(i, 2)
24   S[st'.sid] ← (ks, kt, ke)
25 τ ← M.Tag(kt, Lrm||n||n'||id)
26 st'.state ← (n, n', X, Y, ks, kt)
27 m' ← (n', Y, Enc(ke, (id, σ, τ)))
```

Figure 10: Game G₇ and reduction B₂ to PRF security of the SIGMA-I proof. Changes from G₆ resp. compared to G₇ highlighted in gray.

Our final game hops are concerned with the explicit authentication performed through signatures and MACs in the SIGMA-I protocol, and as such extend those proof steps for implicit authentication of the main protocols in [21].

Game 8. In Game G_8 (Figure 11), we log all messages for which signatures are generated by an honest session, and set a bad flag bad_S if the adversary submits a valid signature under an uncorrupted signing key for a message which was not produced by an honest session. This internal bookkeeping does not affect the adversary’s advantage, so

$$\Pr[G_7 \Rightarrow 1] = \Pr[G_8 \Rightarrow 1].$$

Game 9. In Game G_9 (Figure 11), we abort if the bad_S flag is set. By the identical-until-bad lemma, the difference in advantage between G_8 and G_9 is bounded by the probability that this event occurs, which we reduce via an algorithm \mathcal{B}_3 to the multi-user security of the digital signature scheme S .

In the reduction, \mathcal{B}_3 obtains all long-term public keys from the multi-user signature game and uses its signing oracles for any honest signature to be produced. It therefore makes q_N queries to `NEW` and one `Sign` query for each call to `RunResp1` or `RunInit2`, for at most q_S such queries. It relays `REVLONGTERMKEY` queries as corruptions in its multi-user game, making q_{RL} corruption queries in total. When bad_S is triggered, \mathcal{B}_3 submits the triggering message and signature under the targeted (uncorrupted) public key as its forgery. As the triggering message was not signed before under the corresponding secret key (and hence not queried to the signing oracle by \mathcal{B}_3), the forgery is valid and \mathcal{B}_3 wins if bad_S is set. It follows that

$$\Pr[G_8 \Rightarrow 1] \leq \Pr[G_9 \Rightarrow 1] + \text{Adv}_{S, \mathcal{B}_3}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_3}, q_N, q_S, q_S, q_{RL}).$$

Except for the replacement of key generation, signatures, corruptions with oracle queries, the pseudocode of \mathcal{B}_3 is identical to that of \mathcal{A} in game G_8 , so $t_{\mathcal{B}_3} \approx t$.

Game 10. In Game G_{10} (Figure 11), we remove the now redundant bad_S flag again, and log all MAC tags generated by honest sessions with honest partners in a list Q_M (using, as before, the table S to determine whether a session has an honest partner). We set a flag bad_M if a session with an honest partner receives a valid MAC tag which was not computed by any honest oracle. This bookkeeping is similar to the changes from G_7 to G_8 , but noting MAC tags instead of signatures. As before, the bookkeeping itself does not affect the adversary’s advantage:

$$\Pr[G_9 \Rightarrow 1] = \Pr[G_{10} \Rightarrow 1].$$

Game 11. In Game G_{11} (Figure 11), we abort if the bad_M flag is set to true. Again applying the identical-until-bad lemma, we need to bound the probability of bad_M being set in G_{10} , which we do via the following reduction \mathcal{B}_4 to the multi-user EUF-CMA security of the MAC scheme M .

The reduction \mathcal{B}_4 simulates G_{10} truthfully, except that for any session with honest origin partner (i.e., session with state st where $S[st.sid] \neq \perp$), \mathcal{B}_4 does not compute k_t itself, but instead assigns an incremented user identifier i to this session’s sid and computes any calls to `Tag` or `Vrfy` using its corresponding oracles for user i . There is at most one query to `NEWUSER`, and one each to `Tag` and `Vrfy` for each of \mathcal{A} ’s queries to `SEND`. Hence \mathcal{B}_4 makes at most q_S queries to each of these three oracles, and at most one query to `Tag` and `Vrfy` per user in the mu-EUF-CMA game. When bad_M is triggered, \mathcal{B}_4 submits the triggering message and MAC tag under user identifier i as its forgery. In

$G_8, \boxed{G_9}$

RunInit2($id, sk, st, peerpk, m$):

```
1 ...
2 if  $S.Vrfy(peerpk[peerid], L_{rs} || n || n' || X || Y, \sigma)$ 
   and  $M.Vrfy(k_t, L_{rm} || n || n' || peerid, \tau)$  then
3   if  $revltk_{peerid} = \infty$  and
   ( $peerid, L_{rs} || n || n' || X || Y$ )  $\notin Q_S$  then
4      $bad_S \leftarrow true$   $\boxed{abort}$ 

7    $st'.status \leftarrow accepted$ 
8    $st'.skey \leftarrow k_s$ 
9    $\sigma' \leftarrow S.Sign(sk, L_{is} || n || n' || X || Y)$ 
10   $Q_S \leftarrow Q_S \cup \{(id, L_{is} || n || n' || X || Y)\}$ 
11   $\tau' \leftarrow M.Tag(k_t, L_{im} || n || n' || id)$ 
```

14 ...

RunResp1($id, sk, st, peerpk, m$):

```
15 ...
16  $\sigma \leftarrow S.Sign(sk, L_{rs} || n || n' || X || Y)$ 
17  $Q_S \leftarrow Q_S \cup \{(id, L_{rs} || n || n' || X || Y)\}$ 
18 ...
19  $\tau \leftarrow M.Tag(k_t, L_{rm} || n || n' || id)$ 
```

21 ...

RunResp2($id, sk, st, peerpk, m$):

```
22 ...
23 if  $S.Vrfy(peerpk[peerid], L_{is} || n || n' || X || Y, \sigma')$ 
   and  $M.Vrfy(k_t, L_{im} || n || n' || peerid, \tau')$  then
24   if  $revltk_{peerid} = \infty$  and
   ( $peerid, L_{rs} || n || n' || X || Y$ )  $\notin Q_S$  then
25      $bad_S \leftarrow true$   $\boxed{abort}$ 
```

```
28    $st'.status \leftarrow accepted$ 
29    $st'.skey \leftarrow k_s$ 
30 else  $st'.status \leftarrow rejected$ 
31 return  $st', m'$ 
```

$G_{10}, \boxed{G_{11}}$

RunInit2($id, sk, st, peerpk, m$):

```
1 ...
2 if  $S.Vrfy(peerpk[peerid], L_{rs} || n || n' || X || Y, \sigma)$ 
   and  $M.Vrfy(k_t, L_{rm} || n || n' || peerid)$  then
3   if  $revltk_{peerid} = \infty$  and
   ( $peerid, L_{rs} || n || n' || X || Y$ )  $\notin Q_S$  then
4     abort
5   if  $S[st'.sid] \neq \perp$  and
   ( $st'.sid, L_{rm} || n || n' || peerid$ )  $\notin Q_M$  then
6      $bad_M \leftarrow true$   $\boxed{abort}$ 
7    $st'.status \leftarrow accepted$ 
8    $st'.skey \leftarrow k_s$ 
9    $\sigma' \leftarrow S.Sign(sk, L_{is} || n || n' || X || Y)$ 
10   $Q_S \leftarrow Q_S \cup \{(id, L_{is} || n || n' || X || Y)\}$ 
11   $\tau' \leftarrow M.Tag(k_t, L_{im} || n || n' || id)$ 
12  if  $S[st'.sid] \neq \perp$  then
13     $Q_M \leftarrow Q_M \cup \{(st'.sid, L_{im} || n || n' || id)\}$ 
14 ...
```

RunResp1($id, sk, st, peerpk, m$):

```
15 ...
16  $\sigma \leftarrow S.Sign(sk, L_{rs} || n || n' || X || Y)$ 
17  $Q_S \leftarrow Q_S \cup \{(id, L_{rs} || n || n' || X || Y)\}$ 
18 ...
19  $\tau \leftarrow M.Tag(k_t, L_{rm} || n || n' || id)$ 
20 if  $S[st'.sid] \neq \perp$  then
21    $Q_M \leftarrow Q_M \cup \{(st'.sid, L_{rm} || n || n' || id)\}$ 
22 ...
```

RunResp2($id, sk, st, peerpk, m$):

```
23 ...
24 if  $S.Vrfy(peerpk[peerid], L_{is} || n || n' || X || Y, \sigma')$ 
   and  $M.Vrfy(k_t, L_{im} || n || n' || peerid, \tau')$  then
25   if  $revltk_{peerid} = \infty$  and
   ( $peerid, L_{is} || n || n' || X || Y$ )  $\notin Q_S$  then
26     abort
27   if  $S[st'.sid] \neq \perp$  and
   ( $st'.sid, (peerid, L_{im} || n || n' || peerid)$ )  $\notin Q_M$  then
28      $bad_M \leftarrow true$   $\boxed{abort}$ 
29    $st'.status \leftarrow accepted$ 
30    $st'.skey \leftarrow k_s$ 
31 else  $st'.status \leftarrow rejected$ 
32 return  $st', m'$ 
```

Figure 11: Games G_8, G_9, G_{10} , and G_{11} of the SIGMA-I proof. Changes in G_8 and G_{10} are highlighted in gray, changes in G_9 and G_{11} are highlighted in $\boxed{\text{frames}}$.

the simulation, sessions will share a user identifier i if and only if they are partnered and would share keys in Game G_{10} . These keys are furthermore unique to one initiator and one responder session only, so consistency is maintained. Furthermore, k_t cannot be exposed (by `REVLONGTERMKEY` or `REVSESSIONKEY`) to adversary \mathcal{A} , hence implicitly replacing it with the MAC game's oracles is sound, and \mathcal{B}_4 makes no `CORRUPT` queries. Except for oracle replacements, the pseudocode of \mathcal{B}_4 is identical to that of \mathcal{A} in G_{10} , so $t_{\mathcal{B}_4} \approx t$.

If `badM` is triggered, then $S[st'.sid] \neq \perp$, so $st'.sid$ corresponds to some user identifier i in the multi-user EUF-CMA game. Additionally, a tag τ for message m was verified under identity i , and $(st'.sid, m)$ was not logged in Q_M . Since \mathcal{B}_4 logs $(st'.sid, m)$ every time it calls its `Tag` oracle on the pair (i, m) , this call cannot have occurred. Then τ is a valid forgery on m , which \mathcal{B}_4 will output for user i to win the EUF-CMA game. Thus,

$$\Pr[G_{10} \Rightarrow 1] \leq \Pr[G_{11} \Rightarrow 1] + \text{Adv}_{M, \mu}^{\text{EUF-CMA}}(t_{\mathcal{B}_4}, q_S, q_S, 1, q_S, 1, 0).$$

We can now consider the final advantage of an adversary playing Game G_{11} . Adversary \mathcal{A} has a non-zero advantage if in the final oracle query `FINALIZE`(b')

1. `Sound` is false,
2. `ExplicitAuth` is false, or
3. `Fresh` is true and $b' = b$.⁵

Soundness. The flag `Sound` is set if (1) three honest sessions hold the same session identifier, or if (2) two partnered sessions hold different session keys.

For (1): No three honest sessions can share the same session identifiers, as this would require a collision in either the contained initiator or responder nonce, which is excluded by Game G_2 .

For (2): The session identifier includes both nonces n and n' and DH shares X and Y , which together determine the derived master key $mk = \text{RO}(n \| n' \| X \| Y \| Z)$ (where Z is the DH secret from X and Y) and thus the session key. Agreement on the session identifier hence implies deriving the same session key.

Hence, in Game G_{11} , `Sound` is always true.

Explicit authentication. The predicate `ExplicitAuth` requires that for any session π_u^i accepting with a non-compromised peer v , there exists a partnered session π_v^j of user v with opposite role which, if it accepts, has u set as its peer.

The session π_u^i , prior to accepting, obtained a valid signature on $\pi_u^i.sid$ and a label corresponding to a role $r \neq \pi_u^i.role$. Due to Game G_9 , this signature must have been issued by an honest session π_v^j (since v was not compromised at this point). All honest sessions sign their own `sid` and a label corresponding to their own role, so $\pi_u^i.sid = \pi_v^j.sid$ and $\pi_u^i.role = r \neq \pi_v^j.role$ are satisfied.

Furthermore, when π_v^j accepts, it must have received a valid MAC tag τ on a label identifying an opposite-role session and that session's user identity, as well as their shared nonces. Due to Game G_{11} , this MAC value must have been computed by an honest session holding the same nonces, as π_v^j has an honest partner session and therefore $S[\pi_v^j.sid] \neq \perp$. Furthermore, by Game G_2 , nonces do not collide and hence that session must have been π_u^i , thus computing the MAC on user identity u , which π_v^j accordingly sets as peer identity.

Therefore `ExplicitAuth` is always true in G_{11} . Note that we did not require that the long-term key of user u was uncorrupted, and we allow the adversary to continue interacting with sessions after compromise; hence covering key compromise impersonation attacks.

⁵If `Fresh` is false, $b = b' = 0$ happens with probability $\frac{1}{2}$, so \mathcal{A} 's advantage is 0.

Guessing the challenge bit. Finally, we have to consider \mathcal{A} 's chance of guessing the challenge bit b , which it may only learn through TEST queries such that all tested sessions are fresh (i.e., Fresh is true).

The Fresh predicate being true ensures that all tested sessions (those in T) accepted prior to their respective partner being corrupt. Then, as ExplicitAuth is true, we have that for each tested session there exists an honest session with the same sid and different roles. This session, by Fresh, was not tested or revealed. Being partnered, the first message (n, X) between these two honest sessions was not tampered with, so in the responder session, whether it was the tested session or its partner, the master and session keys are sampled uniformly at random (due to Games G_6 and G_7). Since the initiator session holds the same sid , it copied the responder's random session key (due to Game G_4). This random session key was not revealed in either of the two sessions (by Fresh), and hence from \mathcal{A} 's perspective is a uniformly random and independent value. In all TEST oracle responses, k_0 and k_1 are hence identically distributed and so G_{11} is fully independent of b . It follows that the adversary \mathcal{A} has no better than a $\frac{1}{2}$ probability of choosing b' equal to b , so

$$\Pr[G_{11} \Rightarrow 1] = \frac{1}{2},$$

which concludes the proof. □

6 The TLS 1.3 Handshake Protocol

The Transport Layer Security (TLS) protocol in version 1.3 [47] bases its key exchange design (the so-called handshake protocol) on a variant of SIGMA-I. Following the core SIGMA design, the TLS 1.3 main handshake is an ephemeral Diffie–Hellman key exchange where authentication is performed through a combination of signing and MAC-ing the (full, hashed) communication transcript.⁶ Additionally, and similar to SIGMA-I, beyond establishing the main session key (the application traffic secret), handshake traffic keys are derived and used to encrypt part of the handshake.

Except for additional protocol complexity like negotiating the cryptographic algorithms to be used, communicating further information in extensions, etc.—features which we do not capture here—, TLS 1.3 however deviates in two core cryptographic aspects from the more simplistic and abstract SIGMA(-I) design: it hashes the communication transcript when deriving keys and computing signatures and MACs, and it uses a significantly more complicated key schedule. In this section we revisit the TLS 1.3 handshake to see under which additional assumptions our tight security results for SIGMA-I can be translated to its main key exchange mode.

6.1 Protocol Description

Due to the substantial complexity of the TLS 1.3 handshake, we focus here on a slightly simplified version of the protocol which encompasses all essential cryptographic aspects but abstracts away further technical protocol details that are not relevant for our analysis. In particular, we only consider mutual authentication and security of the main application traffic keys (see [24, 26, 28] for full computational, multi-stage key exchange analyses of the different modes with varying authentication), and accordingly leave out some computations and additional messages. To ease linking back to the underlying SIGMA-I structure, we describe the protocol in the following referencing

⁶TLS 1.3 also specifies an abbreviated resumption-style handshake based on pre-shared keys; we focus on the main DH-based handshake in this work.

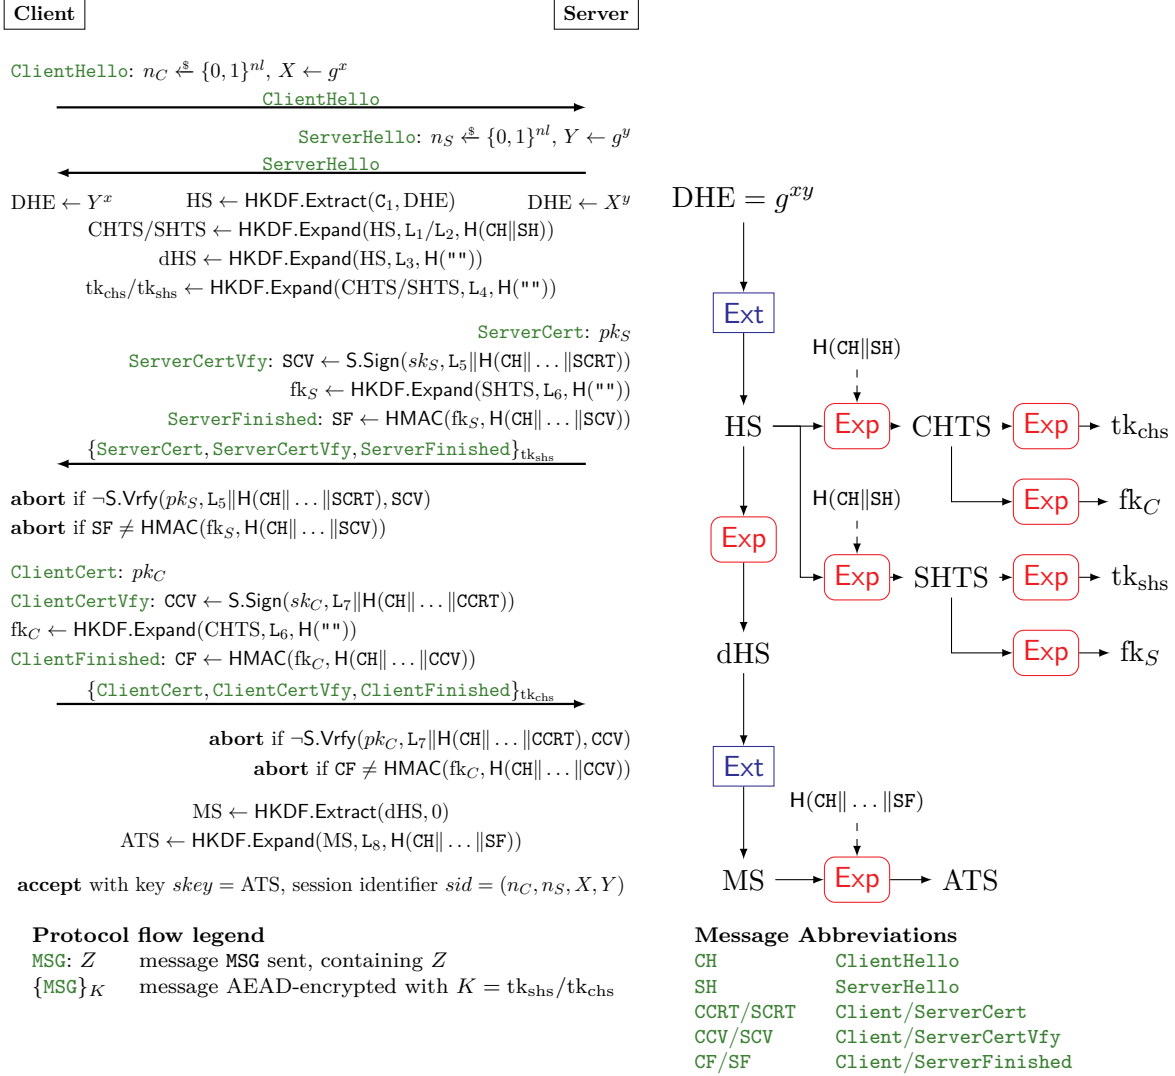


Figure 12: The simplified TLS 1.3 main Diffie–Hellman handshake protocol (left) and key schedule (right). Values L_i and C_i indicate bitstring labels, resp. constant values, (distinct per i). Boxes Ext and Exp denote HKDF extraction and expansion, respectively, with dashed inputs to Exp indicating context information (see protocol figure for detailed computations).

back to the latter (cf. Section 4). We illustrate the handshake protocol and its accompanying key schedule in Figure 12, the latter deriving keys in the extract-then-expand paradigm of the HKDF key derivation function [39].⁷

In the TLS 1.3 handshake, the client acts as initiator and the server as responder. Within Hello messages, both send nonce values n_C resp. n_S together with ephemeral Diffie–Hellman shares g^x resp. g^y . Based on these values, both parties extract a handshake secret HS from the shared DH value $\text{DHE} = g^{xy}$ using HKDF.Extract with a constant salt input.⁸ In a second step, client and

⁷We follow the standard HKDF syntax: $\text{HKDF.Extract}(XTS, SKM)$ on input salt XTS and source key material SKM outputs a pseudorandom key PRK . $\text{HKDF.Expand}(PRK, CTXinfo)$ on input a pseudorandom key PRK and context information $CTXinfo$ outputs pseudorandom key material KM .

⁸This salt input becomes relevant for pre-shared key handshakes, but in the full handshake takes the constant value $C_1 = \text{Expand}(\text{Extract}(0, 0), \text{"derived"}, \text{H}(""))$.

server derive their respective handshake traffic keys tk_{chs} , tk_{shs} and MAC keys fk_C , fk_S through two levels of `HKDF.Expand` steps from the handshake secret HS , including in the first level distinct labels and the hashed communication transcript $H(CH\|SH)$ so far as context information.

The handshake traffic keys are then used to encrypt the remaining handshake messages. First the server, then the client send their certificate (carrying their identity and public key), a signature over the hashed transcript up to including their certificate ($H(CH\|\dots\|SCRT)$, resp. $H(CH\|\dots\|CCRT)$), as well as a MAC over the (hashed) transcript up to incl. their signatures ($H(CH\|\dots\|SCV)$, resp. $H(CH\|\dots\|CCV)$). Note the similarity to SIGMA-I here: each party signs both nonces and DH values (within $CH\|SH$, modulo transcript hashing) together with a unique label, and then MACs both nonces and their own identity (the latter being part of their certificate).⁹ The application traffic secret ATS —which we treat as the session key *key*, unifying secrets of both client and server—is then derived from the master secret MS through `HKDF.Expand` with handshake context up to the `ServerFinished` message. The master secret in turn is derived through (context-less) `Expand` and `Extract` from the handshake secret HS .

6.2 Handling the TLS 1.3 Key Schedule

As mentioned before, the message flow of the TLS 1.3 handshake relatively closely follows the SIGMA-I design [37, 38] (cf. Figure 4): after exchanging nonces and DH shares (in `Hello`) from both sides, the remaining (encrypted) messages carry identities (`Certificate`), signatures over the nonces and DH shares (`CertificateVerify`), and MACs over the nonces and identities (`Finished`).

What differentiates the TLS 1.3 handshake from the basic SIGMA-I design (beyond putting more under the respective signatures and MACs, which does not negatively affect the key exchange security we are after) is the way keys are derived. While SIGMA-I immediately derives a master key through a random oracle with input both the shared DH secret as well as the session identifying nonces and DH shares, TLS 1.3 separates both in its HKDF-based extract-then-expand key schedule: The core secrets—handshake secret (HS) and master secret (MS)—are derived without further context purely from the shared DH secret $DHE = g^{xy}$ (beyond other constant inputs). Only when deriving the specific-purpose secrets—handshake traffic keys (tk_{chs} , tk_{shs}), MAC keys (fk_C , fk_S), and session key (ATS)—is context added to the key derivation, including in particular the nonces and DH shares identifying the session. To complicate matters even further, this context is hashed before entering key derivation (or signature and MAC computation), and the final session key ATS depends on more messages than just the session-identifying ones. Since our tighter security proof for the SIGMA(-I) protocol (cf. Section 5) heavily makes use of (exactly) the session identifiers being input together with DH secrets to a random oracle when programming the latter, the question arises how to treat the TLS 1.3 key schedule when aiming at a similar proof strategy.

Our approach to achieving a tighter security proof for the TLS 1.3 key schedule will be as follows: First of all, we model both `HKDF.Extract` and `HKDF.Expand` as (programmable) random oracles. We then have to carefully orchestrate the programming of intermediate secrets and session keys in a two-level approach, taking into account that inputs to deriving the session keys depend on values established through the intermediate secrets (namely, the server’s `Finished` MAC). Along the way, we separately ensure that we recognize any hashed inputs of interest that the adversary might query to the random oracle, without modeling the hash function H as a random oracle itself. By tracking intermediate programming points (especially HS and MS) in the random oracles, we recover the needed capability of linking sessions and their session identifiers and DH shares exchanged to the corresponding session keys. This finally allows us to again (efficiently) determine when and on what

⁹Instead of using distinct labels for the client and server MAC computations, TLS 1.3 employs distinct MAC keys for client and server, achieving separation between the two MAC values this way.

input to query the strong Diffie–Hellman oracle when programming challenge DH shares into the TLS 1.3 key exchange execution during the proof.

7 Tighter Security Proof for the TLS 1.3 Handshake

We now give our full key exchange security bound for the handshake protocol of TLS 1.3 in the random oracle model, with tighter reductions compared to prior analyses.

Theorem 7.1. *Let \mathcal{A} be a key exchange security adversary against the TLS 1.3 handshake protocol as specified in Figure 12 based on a hash function H , a signature scheme S , and a group \mathbb{G} of prime order p , and let the HKDF functions `Extract` and `Expand` in the protocol be modeled as (independent) random oracles RO_1 , resp. RO_2 . For any $(t, q_N, q_S, q_{RS}, q_{RL}, q_T)$ -KE-SEC-adversary against SIGMA-I making at most q_{RO} queries to the random oracle, we give algorithms $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$, and \mathcal{B}_4 in the proof, with running times $t_{\mathcal{B}_i} \approx t$ (for $i = 1, 3, 4$) and $t_{\mathcal{B}_2} \approx t + 6q_{\text{RO}} \log_2 p$ close to that of \mathcal{A} , such that*

$$\begin{aligned} \text{Adv}_{\text{TLS 1.3}}^{\text{KE-SEC}}(t, q_N, q_S, q_{RS}, q_{RL}, q_T) &\leq \frac{q_S^2}{2^{nl}} + \frac{q_S^2}{2p} + \frac{q_S^2}{2^{nl} \cdot p} + \text{Adv}_H^{\text{CR}}(t_1) \\ &\quad + 2 \cdot \text{Adv}_{\mathbb{G}}^{\text{stDH}}(t_2, q_{\text{RO}}) + \frac{q_{\text{RO}} \cdot q_S}{2^{kl-1}} + \text{Adv}_S^{\text{mu-EUF-CMA}}(t_3, q_N, q_S, q_S, q_{RL}) \\ &\quad + \text{Adv}_{\text{HMAC}}^{\text{mu-EUF-CMA}}(t_4, q_S, q_S, 1, q_S, 1, 0), \end{aligned}$$

where $nl = 256$ is the nonce length in TLS 1.3, kl is the output length of $\text{RO}_2 = \text{HKDF.Expand}$, \mathbb{G} is the used Diffie–Hellman group of prime order p , and $q_S \cdot q_{\text{RO}} \leq 2^{kl-2}$.¹⁰

Proof. We prove our bound by making an incremental series of changes to the key exchange security game and limiting the amount that each change affects the success probability of \mathcal{A} .

Game 0. The initial game, Game G_0 is the key exchange security game for TLS played by \mathcal{A} should be equivalent to the key exchange security game, using the implicit `KGen`, `Activate`, and `Run` routines defined by the TLS protocol specification on the left side of Figure 12. (In this game, `HKDF.Extract` and `HKDF.Expand` are modeled by random oracles RO_1 and RO_2 respectively.) By definition,

$$\Pr[G_0 \Rightarrow 1] = \Pr[G_{\text{TLS}, \mathcal{A}}^{\text{KE-SEC}} \Rightarrow 1].$$

Game 1. In game G_0 , we start logging the nonces and group elements chosen by honest sessions. Whenever two honest sessions choose the same nonces or group elements, we set a flag `badC`. Whenever an honest responder session chooses a nonce and group element that have already been received by another session, we set a flag `badO`. We also make both random oracles RO_1 and RO_2 lazily sampled using internal tables H_1 and H_2 . These changes only affect the values of the game’s internal state, and the view of the adversary remains the same as in G_0 , so

$$\Pr[G_1 \Rightarrow 1] = \Pr[G_0 \Rightarrow 1].$$

Game 2. Starting with G_2 , we abort whenever two honest sessions sample the same nonce or group element and whenever an honest responder samples a nonce and group element that are

¹⁰To simplify the factor on $\text{Adv}_{\mathbb{G}}^{\text{stDH}}$ to 2, we assume that $q_S \cdot q_{\text{RO}} \leq 2^{kl-2}$, which is true for any reasonable real-world parameters. See the proof for the exact bound.

already in use. Since this happens only after one of the flags bad_C and bad_O is set, by the identical-until-bad lemma,

$$\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1] \leq \Pr[\text{bad}_C \leftarrow \text{true or } \text{bad}_O \leftarrow \text{true in } G_1].$$

One nonce and one group element is chosen in each RunInit1 call and each RunResp1 call, so at most one nonce and one group element is chosen for each query the adversary makes to its SEND oracle. Let the adversary make no more than q_{SEND} such queries. We use the birthday bound to limit the probability of a collision in either the set of honest sessions' nonces or the set of honest sessions' DH shares. Every time a responder session chooses a nonce and group element, there are at most q_S values have already been chosen, so by the union bound bad_O is set with probability at most $\frac{q_S^2}{2^{nl} \cdot p}$. Therefore

$$\Pr[G_1 \Rightarrow 1] - \Pr[G_2 \Rightarrow 1] \leq \frac{q_S^2}{2^{nl}} + \frac{q_S^2}{2p} + \frac{q_S^2}{2^{nl} \cdot p}.$$

Game 3. Next, we must ensure that partial transcripts between honest sessions do not collide under the hash function H . This is a step unique to the TLS proof, which hashes all of its context with a collision-resistant hash function before it is input into key-derivation. In G_3 , honest sessions will log all of their hash outputs in a look-up table T : whenever an honest session computes $d = H(s)$ for some string s , it sets $T[d] \leftarrow s$ if $T[d]$ has not already been defined. If $T[d]$ is not empty, then some prior honest session has computed $d = H(s')$ for some string s' . The session will set a flag bad_H if $s' \neq s$, noting that a collision has occurred. We also remove the now superfluous bad_C flag. These administrative changes do not affect the view of the adversary, so

$$\Pr[G_3 \Rightarrow 1] = \Pr[G_2 \Rightarrow 1].$$

Game 4. In Game G_4 , we abort whenever hashes computed by honest sessions collide (i.e. the bad_H flag is set). By the identical-until-bad lemma,

$$\Pr[G_3 \Rightarrow 1] - \Pr[G_4 \Rightarrow 1] \leq \Pr[\text{bad}_H \leftarrow \text{true in } G_3].$$

We bound the probability that bad_H is set via a reduction \mathcal{B}_1 to the collision-resistance security of H . The reduction simulates G_3 honestly for the adversary \mathcal{A} . If the flag bad_H is set, then the reduction has obtained strings s , s' , and d such that $s' \neq s$, and $H(s) = H(s') = d$. Then \mathcal{B}_1 outputs (s, s') and wins the collision-resistance game, so $\text{Adv}_{H, \mathcal{B}_1}^{\text{CR}} \geq \Pr[\text{bad}_H \leftarrow \text{true in } G_3]$. The runtime of \mathcal{B} approximately equals the runtime of \mathcal{A} in G_3 . It follows that

$$\Pr[G_3 \Rightarrow 1] - \Pr[G_4 \Rightarrow 1] \leq \text{Adv}_H^{\text{CR}}(\text{time}_{\mathcal{B}_1}).$$

Game 5. In Game G_5 , we remove the superfluous bad_H flag and make additional internal changes to the behavior of honest sessions. As in the SIGMA-I proof, all honest initiator sessions now log the first message they send in a set Sent , and honest responder sessions use this set to check whether their first received message came from an honest session without tampering. If so, we say the responder session has an "honest origin partner." In the SIGMA-I protocol, partnering between honest sessions was sufficient to ensure agreement on the derived master key and all subsequently computed keys, since partners are guaranteed to hold the same nonces and group elements. In

TLS 1.3, partnering also ensures agreement on the handshake traffic secrets SHTS and CHTS, but it does not ensure agreement on the session key ATS. Therefore the responder only logs the handshake traffic keys fk_S, fk_C, tk_{shs} , and tk_{chs} in a look-up table S under its session identifier. In addition to the session identifier, the application traffic secret ATS depends on the server’s identity SCRT, signature SCV, and MAC tag SF. These values are not necessarily shared by partner sessions in Game G_5 , so two partnered sessions may derive different values of ATS. The responder session therefore logs its session key ATS in a second look-up table S' indexed by all of the dependencies of the session key: $sid, SCRT, SCV$, and SF . All of this is just bookkeeping, so

$$\Pr[G_5 \Rightarrow 1] = \Pr[G_4 \Rightarrow 1].$$

Game 6. Going forward from Game G_6 , honest initiators copy their key material from tables S and S' if it is consistent for them to do so. In the case where the adversary has tampered with the values of SCRT, SCV, or SF, the partner’s session key depends on the untampered values and should not be copied. Therefore honest initiators always copy encryption and MAC keys from the table S if they have an honest partner session, but they only copy ATS when the SCRT, SCV, and SF messages they received match the ones sent by their partner. The initiator session can check whether tampering occurred using the table S' , which will contain a session key ATS at index $sid||SCRT||SCV||SF$ if and only if the honest partner session computed and sent SCRT, SCV, and SF.

We argue that all copied keys are consistent with the keys that would be derived in G_5 . Recall that partnered sessions agree on the nonces and the DH shares X and Y as components of sid , so they also agree on the shared DH secret Z associated with the pair (X, Y) . Partnered sessions therefore agree on the handshake secret HS, which is derived from Z without context, and on the handshake traffic secrets, which are derived with the session identifier as context. Thus partnered sessions agree on the values of the handshake traffic keys fk_S, fk_C, tk_{shs} , and tk_{chs} which are derived from the handshake traffic secrets. For the adversary it is hence unobservable if honest sessions compute the handshake traffic keys themselves, or copy the keys from their partners. By agreeing on the handshake secret HS, partnered sessions will also agree on the master secret MS, which is derived from HS without context. The if SCRT, SCV, and SF are left untampered, both sessions will derive the session key as $RO_2(\text{MS}, L_8, H(sid||SCRT||SCV||SF))$. Hence it is again unobservable whether an honest initiator derives ATS itself or copies ATS from an honest partner which agrees on the values of SCRT, SCV, SF; consequently

$$\Pr[G_6 \Rightarrow 1] = \Pr[G_5 \Rightarrow 1].$$

Game 7. In Game G_7 , all responders sample ATS, SHTS and CHTS randomly (unless their values have already been fixed by queries to random oracle RO_2 on the corresponding input), then retroactively programs random oracle RO_2 by setting its internal table H_2 on the appropriate input. Partnered initiator sessions which have not copied ATS (i.e., those who received tampered SCRT, SCV, and SF) also sample ATS randomly and program RO_2 when necessary. We choose to program ATS, SHTS, and CHTS, as opposed to only mk in the SIGMA-I proof, because these three keys are derived with context. Most importantly, the DH shares X and Y indirectly enter the key derivation for these keys, which will be critical for the reduction in the next step. This simply moves the lazy sampling process from RO_2 to RunResp1 and RunInit2 for certain queries, which is unobservable to the adversary; therefore

$$\Pr[G_7 \Rightarrow 1] = \Pr[G_6 \Rightarrow 1].$$

Game 8. The step between G_7 and G_8 is most technically involved step of this proof, and it is also the most significantly altered from the corresponding step in the proof of SIGMA-I. In G_8 , partnered initiators and responder sessions with honest origin partners will stop maintaining the consistency of their keys ATS , SHTS , and CHTS with the random oracle RO_2 . Specifically, responders with honest origin partners sample ATS , SHTS , and CHTS uniformly at random even if RO_2 has already been queried on the string HS, L, d for the appropriate label and hash, and they do not retroactively program RO_2 . Partnered initiator sessions which have not copied ATS from their partner also sample ATS uniformly without checking or programming RO_2 . These keys are therefore completely random, and they will be inconsistent with any random oracle queries made before or after the keys are sampled. In order to detect this inconsistency, the adversary must make a query to RO_2 that would, in G_7 , return one of the unprogrammed keys. Which queries are these? They are the queries that an honest responder session with honest origin partner would use to derive SHTS , CHTS , and ATS , and the queries that an honest partnered initiator which received a tampered message would use to derive ATS . Formally, let $\text{sid} = (n, n', X, Y)$ be the session ID held by some honest responder session with honest origin partner, and let SCRT , SCV , SF be the identity, signature, and MAC tag sent by this session. Let DHE be the DH secret corresponding to the pair (X, Y) . Then the adversary \mathcal{A} can detect inconsistency in game G_8 if at any point during the game \mathcal{A} queries RO_2 on the tuple

$$(\text{RO}_1(\text{C}_1, \text{DHE}), \text{L}, \text{H}(\text{sid})),$$

with $\text{L} \in \{\text{L}_1, \text{L}_2\}$, or the tuple

$$(\text{MS}, \text{L}_8, \text{H}(\text{sid} \parallel \text{SCRT} \parallel \text{SCV} \parallel \text{SF})),$$

where for some HS , dHS , we have that $\text{HS} = \text{RO}_1(\text{C}_1, \text{DHE})$, that $\text{dHS} = \text{RO}_2(\text{HS}, \text{L}_3, \text{H}(""))$, and that $\text{MS} = \text{RO}_1(\text{dHS}, 0)$. Otherwise, let sid be the session ID held by an honest partnered initiator session, and let SCRT , SCV , and SF be the identity, signature, and MAC tag received by this session. Assume that at least one of these values was not sent by the honest partner. Then the adversary \mathcal{A} can detect inconsistency in game G_8 if at any point it queries RO_2 on the tuple

$$(\text{MS}, \text{L}_8, \text{H}(\text{sid} \parallel \text{SCRT} \parallel \text{SCV} \parallel \text{SF})),$$

and for some HS , dHS , we have that $\text{HS} = \text{RO}_1(\text{C}_1, \text{DHE})$, that $\text{dHS} = \text{RO}_2(\text{HS}, \text{L}_3, \text{H}(""))$, and that $\text{MS} = \text{RO}_1(\text{dHS}, 0)$. Let event F denote the event that the adversary \mathcal{A} makes at least one of the above queries. If event F does not occur, then ATS , SHTS , and CHTS are chosen uniformly at random in both G_7 and G_8 , hence

$$\Pr[G_7 \Rightarrow 1] - \Pr[G_8 \Rightarrow 1] \leq \Pr[F \text{ occurs in } G_7].$$

We bound the probability of event F via a reduction \mathcal{B}_2 to the strong Diffie-Hellman assumption in group \mathbb{G} . The reduction will make no more queries to its stDH oracle than \mathcal{A} makes to its RO_2 oracles.

Given its strong DH challenge $(A = g^a, B = g^b)$ and having access to the strong Diffie Hellman oracle stDH_a , \mathcal{B}_2 simulates G_7 for an adversary \mathcal{A} in the following manner: In all honest initiator sessions, \mathcal{B}_2 samples r uniformly at random from \mathbb{Z}_q and sets the session's DH share $X \leftarrow A \cdot g^r$. In all honest responder sessions with honest origin partner, \mathcal{B}_2 samples r' uniformly from \mathbb{Z}_q and sets the session's DH share $Y \leftarrow B \cdot g^{r'}$. Both of these DH shares are still distributed uniformly over \mathbb{Z}_q as long as q is prime and A and B are not the identity. To extract g^{ab} when event F occurs, the reduction \mathcal{B}_2 will follow the same general strategy as the reduction \mathcal{B}_1 in the proof of SIGMA-I, with four major points of divergence. We address these points first, before giving a full description of \mathcal{B}_2 .

1. Since \mathcal{B}_2 no longer knows x or y such that $X = g^x$ or $Y = g^y$, it cannot compute the Diffie-Hellman secret DHE or the derived handshake secret HS, so it samples HS randomly for honest responder sessions with honest origin partners and for honest partnered initiator sessions. The adversary can only tell that HS was not correctly computed if it notices that SHTS, CHTS, or dHS are derived from an incorrect value of HS. The former two cases require the adversary to make a query that triggers event F . In the latter case, dHS is not revealed to the adversary through any oracle, so the adversary must notice that ATS, which is derived indirectly from dHS via the master secret, is derived from an incorrect value of HS. This also requires \mathcal{A} to make a query that triggers event F . Therefore, until event F occurs, this change is unobservable to the adversary.
2. In the TLS protocol, the context string, including the Diffie-Hellman shares X and Y , is hashed with H before it enters key derivation, so \mathcal{B}_2 cannot directly associate a query to RO_2 with the honest session(s) whose session ID is being used. The reduction addresses this by having each honest responder with honest origin partner and each honest partnered initiator, log the hash of its context in a reverse look-up table R . (The context does not include the handshake or master secrets.) Then in the RO_2 oracle, \mathcal{B}_2 can use R to efficiently check whether the hash d of a query is used to derive a handshake or application traffic key.
3. Due to TLS's complex key schedule, no one random oracle query contains both a pair of Diffie-Hellman shares and the DH secret associated with that pair. Instead, \mathcal{B}_2 will augment the RO_1 and RO_2 oracles to log in a reverse look-up table T the DH secret associated with each of the intermediate values HS, dHS, and MS. The DH secret for $\text{dHS} = \text{RO}_2(\text{HS}, L_3, H(""))$ simply be copied from $T[\text{HS}]$, and the DH secret for $\text{MS} = \text{RO}_1(\text{dHS}, 0)$ will be copied from $T[\text{dHS}]$. For each query to RO_2 with secret s , the reduction can efficiently check using T whether s was derived from some DH secret via RO_1 .
4. The TLS key schedule uses multiple random oracle queries (if we model HKDF.Extract and HKDF.Expand as random oracles) whereas the SIGMA-I protocol uses only one. If \mathcal{A} can guess the intermediate value $\text{HS} = \text{RO}_1(\mathcal{C}_1, \text{DHE})$, where DHE is the DH secret associated to some pair of embedded shares (X, Y) chosen by honest sessions, then it can trigger event F without ever submitting DHE to an oracle. In this case, \mathcal{A} can trigger event F , but \mathcal{B}_2 cannot win the Strong DH game. However, if $\text{RO}_1(\mathcal{C}_1, \text{DHE})$ is never queried, then it is uniformly random, and the probability that \mathcal{A} guesses correctly is bounded by $\frac{q_{\text{RO}_2} \cdot q_{\text{SEND}}}{2^{kt}}$ by the birthday bound.

To compute the correct handshake and application traffic keys, \mathcal{B}_2 needs to be able to correctly program CHTS, SHTS, and ATS. When these keys are chosen by an honest responder with honest origin partner or a partnered initiator, \mathcal{B}_2 uses its strong DH oracle to check whether RO_2 has already received the query that the adversary needs to make to generate these keys. If the query has already been made, \mathcal{B}_2 can look up the DH secret using T and win the game. Otherwise, \mathcal{B}_2 hashes the session's context and logs it in R , so that future RO_2 queries can identify this session for retroactive programming. It also logs the session's randomness in a look-up table Q , to be used if event F is triggered relative to this session by a future RO_2 query.

Like in the SIGMA-I proof, \mathcal{B}_2 must be able to correctly compute handshake and application traffic keys for unpartnered initiator sessions. Because all initiator sessions have embedded DH shares, \mathcal{B}_2 cannot compute the DH secret DHE for these sessions. However, it can use its StrongDH oracle to check whether the adversary has queried such a secret and copy the expected keys to preserve consistency in this case. If no query has been made, the keys are selected randomly and the initiator session stores its context, randomness, and keys in R . In future queries to the RO_2

oracle, \mathcal{B}_2 will use R to efficiently check whether a query should output one of the initiator session's keys. If so, it retroactively programs the oracle using the keys from R .

Therefore, if event F occurs, reduction \mathcal{B}_2 wins the strong Diffie-Hellman game except with probability $\frac{q_{\text{RO}} \cdot q_{\text{S}}}{2^{kl}}$, so $\text{Adv}_{\mathbb{G}}^{\text{stDH}}(\text{time}_{\mathcal{B}_2}, q_{\text{RO}}) \geq (1 - \frac{q_{\text{RO}} \cdot q_{\text{S}}}{2^{kl}}) \cdot \Pr[F]$. Then $\Pr[F] \leq \frac{2^{kl}}{2^{kl} - q_{\text{RO}} \cdot q_{\text{S}}} \cdot \text{Adv}_{\mathbb{G}}^{\text{stDH}}(\text{time}_{\mathcal{B}_2}, q_{\text{RO}})$. Otherwise, the reduction simulates G_7 perfectly except with probability $\frac{q_{\text{RO}} \cdot q_{\text{SEND}}}{2^{kl}}$.

$$\begin{aligned} \Pr[G_7 \Rightarrow 1] &= \Pr[G_8 \Rightarrow 1] + \Pr[F] + (1 - \Pr[F]) \cdot \frac{q_{\text{RO}} \cdot q_{\text{S}}}{2^{kl}} \\ &\leq \Pr[G_8 \Rightarrow 1] + \frac{2^{kl} + q_{\text{RO}} \cdot q_{\text{S}}}{2^{kl} - q_{\text{RO}} \cdot q_{\text{S}}} \cdot \text{Adv}_{\mathbb{G}}^{\text{stDH}}(\text{time}_{\mathcal{B}_2}, q_{\text{RO}}) + \frac{q_{\text{RO}} \cdot q_{\text{S}}}{2^{kl}} \\ &\leq \Pr[G_8 \Rightarrow 1] + 2 \cdot \text{Adv}_{\mathbb{G}}^{\text{stDH}}(\text{time}_{\mathcal{B}_2}, q_{\text{RO}}) + \frac{q_{\text{RO}} \cdot q_{\text{S}}}{2^{kl}}, \end{aligned}$$

where the last simplification step assumes that $q_{\text{S}} \cdot q_{\text{RO}} \leq 2^{kl-2}$, which is true for any reasonable real-world parameters.

Game 9. In Game G_9 , honest responders with honest origin partners sample $\text{fk}_S, \text{fk}_C, \text{tk}_{\text{chs}}$ and tk_{shs} uniformly at random, so these keys are no longer consistent with RO_2 . The adversary can distinguish this change if and only if it queries RO_2 on a string $\text{SHTS}, \text{L}, \text{H}(\text{" "})$, or $\text{CHTS}, \text{L}, \text{H}(\text{" "})$, where $\text{L} \in \{\text{L}_4, \text{L}_6\}$, and SHTS and CHTS are chosen by an honest responder sessions with honest origin partner. Call this event E . In these sessions, SHTS and CHTS are chosen uniformly at random by G_8 , and they are never revealed by any oracle. Therefore the probability of event E is at most $\frac{q_{\text{RO}} \cdot q_{\text{S}}}{2^{kl}}$ by the birthday bound, hence

$$\Pr[G_8] \leq \Pr[G_9] + \frac{q_{\text{RO}} \cdot q_{\text{S}}}{2^{kl}}.$$

Note that this step in the SIGMA-I proof introduced a multi-user PRF security bound due to final keys being derived through a PRF, not the random oracle. Modeling HKDF.Expand as random oracle RO_2 , we here instead incur a birthday bound under the random oracle instead of a multi-user PRF security bound for HKDF.Expand .

The remaining game hops are identical to those in the proof of SIGMA-I, so we discuss them only briefly.

Game 10. In Game G_{10} , we log all messages signed by an honest session in a look-up table Q_S , and we set a flag bad_S whenever a partnered session verifies a signature with an uncorrupted public key on a message that was not in Q_S . This is just administrative, so

$$\Pr[G_{10} \Rightarrow 1] = \Pr[G_9 \Rightarrow 1].$$

Game 11. In Game G_{11} , we abort if the flag bad_S is set. In this case, an honest partnered session received a signature which was not computed by an honest session, and which was verified by an uncorrupted public key. We can give a straightforward reduction \mathcal{B}_3 to the multi-user EUF-CMA security of the signature scheme that wins whenever bad_S is set and has runtime approximately equal to that of \mathcal{A} in G_{10} . By the identical-until-bad lemma,

$$\Pr[G_{10} \Rightarrow 1] - \Pr[G_{11} \Rightarrow 1] \leq \text{Adv}_{\mathbb{S}}^{\text{mu-EUF-CMA}}(\text{time}_{\mathcal{B}_3}, q_{\text{N}}, q_{\text{S}}, q_{\text{S}}, q_{\text{RL}}).$$

Interestingly and in contrast to the SIGMA-I proof, soundness is still not guaranteed after this game hop, because we do not require the signature scheme to be strongly unforgeable. Therefore the adversary may be able to produce a new signature on a message that had been signed by an honest session, allowing it to tamper with SCV without setting the bad_S flag.

Game 12. In Game G_{12} we log all messages for which an honest session computed a MAC tag in a look-up table Q_M . We remove the bad_S flag and instead set a flag bad_M if an honest partnered session verifies a MAC on a message that is not in Q_M . Again, this is only bookkeeping and does not impact the view of \mathcal{A} , hence

$$\Pr[G_{12} \Rightarrow 1] = \Pr[G_{11} \Rightarrow 1].$$

Game 13. Finally, in Game G_{13} , we abort if an honest session with an honest partner verifies a MAC tag on a message which was not tagged by any honest session; i.e if the bad_M flag is set. We can give a simple reduction \mathcal{B}_4 to multi-user MAC security. The reduction \mathcal{B}_4 assigns a pair of indices $i, i + 1$ to each session identifier held by an honest session with honest origin partner. When an honest session with honest origin partner needs to compute a server MAC tag, \mathcal{B}_4 finds the pair $(i, i + 1)$ using the session identifier and calls its **Tag** oracle with user identity i . When the session needs a client MAC tag \mathcal{B}_4 calls **Tag** with user identity $i + 1$. The reduction calls its **Tag** oracle no more than twice for every query \mathcal{A} makes to **SEND** (once to generate a tag, and once to verify a tag). Since by Game G_9 all honest sessions with honest origin partners sample their MAC keys fk_S and fk_C uniformly at random, the keys implicitly generated by the MAC security game are consistent with the expected operation of Game G_{13} . When the flag bad_M is set, a partnered session has received a valid tag on a message which was never logged in Q_M . The reduction can look up the pair $(i, i + 1)$ using the session identifier of whichever session set bad_M . Since \mathcal{B}_4 logs every message for which it calls its **Tag** oracle, this is a valid forgery for either identity i or identity $i + 1$, and \mathcal{B}_4 will win. Then

$$\Pr[G_{12} \Rightarrow 1] - \Pr[G_{13} \Rightarrow 1] \leq \text{Adv}_M^{\text{mu-EUF-CMA}}(\text{time}_{\mathcal{B}_4}, q_S, q_S, 1, q_S, 1, 0).$$

The runtime of \mathcal{B}_4 is about that of \mathcal{A} in G_{12} .

We can now finally argue that the advantage of \mathcal{A} in G_{13} is zero. The adversary \mathcal{A} would win game G_{13} with probability better than $\frac{1}{2}$ in one of three ways:

1. Sound is false,
2. **ExplicitAuth** is false, or
3. **Fresh** is true and $b' = b$.

Soundness. The flag **Sound** is set if (1) three honest sessions hold the same session identifier, or if (2) two partnered sessions accept with different session keys. By Game G_2 , each session identifier is held by at most one session of each role. There are only two roles so case (1) never occurs. If two partnered sessions π_1 and π_2 accept, the initiator session π_1 verified a MAC tag τ on the message $m = n||n'||X||Y||\text{SCRT}||\text{SCV}$. Because τ was verified by an honest partnered session, by Game G_{13} , this message was tagged by an honest session. Honest sessions only tag strings including their own nonce, and by Game G_2 , the only honest session with nonce n' is π_2 . Then π_2 must have tagged the message m , so π_1 and π_2 agree on both τ and m . Since the DH shares X and Y are components of m , π_1 and π_2 also agree on the DH secret DHE associated with the pair (X, Y) . Consequently, π_1 and π_2 will agree on any value derived deterministically from m , τ , and DHE, including the session key **ATS**. Then the flag **Sound** is always true in G_{13} .

Explicit authentication. The flag `ExplicitAuth` is set if there exists a session π_u^i that accepts with uncorrupted peer identity v , and either (1) no honest session π_v^j is partnered with π_u^i , or (2) a session π_v^j is partnered with π_u^i but accepts with peer identity $w \neq u$. To have accepted with peer identity v , the session π_u^i must have received and verified a signature σ using the public key of identity v on a message m containing the session identifier of π_u^i . As v was uncorrupted at the time that π_u^i accepted, by Game G_{11} , the message m must have been signed by some honest session π_v^j . As honest sessions only sign messages containing their own session identifiers, $\pi_v^j.sid = \pi_u^i.sid$, so π_v^j and π_u^i are partnered. If case (2) occurs, π_v^j must have accepted a MAC tag τ on message m' containing its session ID and the identity w of its peer. We know that π_v^j is a partnered session, so by G_{13} , m' was tagged by some honest session. Honest sessions tag only messages containing their own session identifiers, so by G_2 , the message m' must have been tagged by either π_u^i or π_v^j . In SIGMA-I, the messages tagged by these two sessions are differentiated by their labels. Here, they are differentiated by their length: one role signs a message including values `SF`, `CCRT`, and `CCV`, while the other signs a message which does not contain these values. For this reason π_v^j will not verify the tag on a message it signed itself. Therefore m' must have been tagged by π_u^i , so m' contains the identity u . This contradicts the assumption that $w \neq u$, so case (2) never occurs, and the flag `ExplicitAuth` is always false in G_{13} .

Guessing the challenge bit. Now the adversary can only win with advantage better than zero is by guessing the correct value of b when the `Fresh` flag is set to true. This requirement ensures that all tested sessions accepted with uncorrupted peer identities. Since `ExplicitAuth` is true, each tested session must therefore have an honest session with which it is partnered, and by `Sound`, this session holds the same session key. Then by G_6 , each tested initiator session copies the session key of its partner. By G_8 each tested responder session, and each responder session partnered with a tested initiator session chooses its session key uniformly at random. By `Fresh`, the partners of tested sessions were not tested or revealed. Then the session keys of all tested sessions are sampled uniformly and never revealed to the adversary by any oracle. Therefore the key returned by each `TEST` query is uniformly random and independent of the bit b . The adversary’s view is independent of the bit b , so it will win G_{13} with probability $\frac{1}{2}$, and consequently its advantage is 0.

Collecting the bounds across all game hops gives the theorem statement. \square

8 Evaluation

Tighter security results in terms of loss factors are practically meaningful only if they materialize in better concrete advantage bounds when taking the underlying assumptions into account. In our case, this amounts to the question: How does the overall concrete security of the SIGMA/SIGMA-I and the TLS 1.3 key exchange protocols improve based on our tighter security proofs?

Parameter selection. In order to evaluate our and prior bounds practically, we need to make concrete choices for each of the parameters entering the bounds. Let us explain the choices we made in our evaluation:

Runtime $t \in \{2^{40}, 2^{60}, 2^{80}\}$. We parameterize the adversary’s runtime between well within computational reach (2^{40}) and large-scale attackers (2^{80}).

Number of users $\#U = q_{NW} \in \{2^{20}, 2^{30}\}$. We consider the number of users a global-scale adversary may interact with to be in the order of active public-key certificates on the Internet,

reported at 130–160 million¹¹ ($\approx 2^{27}$).

Number of sessions $\#S \approx q_S \in \{2^{35}, 2^{45}, 2^{55}\}$. Chrome¹² and Firefox¹³ report that 78–97% of all web page accesses through these browsers are encrypted, with an active daily base of about 2 billion ($\approx 2^{30}$) users. We consider adversaries may easily see 2^{35} sessions and a global-scale attacker may have access to 2^{55} sessions over an extended timespan. Note that the number of send queries essentially corresponds to the number of sessions.

Number of RO queries $\#RO = q_{RO} = \frac{t}{2^{10}}$. We fix this bound at a 2^{10} -fraction of the overall runtime accounting for all adversarial steps.

Group order $p \in \{2^{254}, 2^{384}, 2^{518}, 2^{252}, 2^{446}\}$. We consider all five elliptic curves standardized for use with TLS 1.3 (bit-security level b in parentheses): `secp256r1` (128), `secp384r1` (192), `secp521r1` (256), `x25519` (128), and `x448` (224). We focus on elliptic curve groups only, as they provide high efficiency and the best known algorithms for solving discrete-log and DH problems are generic, allowing us to apply GGM bounds for the involved DDH and strong DH assumptions.

Signature schemes. In order to unify the underlying hardness assumptions, we consider the ECDSA/EdDSA signature schemes standardized for use with TLS 1.3, based on the five elliptic curves above, treating their single-user unforgeability as equally hard as the corresponding discrete logarithm.

Symmetric schemes and key/nonce/output lengths $kl = nl = ol$. Since our focus is mostly on evaluating ECDH parameters, we idealize the symmetric primitives (PRF, MAC, and hash function) in the random oracle model. We consider two choices: primarily (esp. in Table 1), we consider symmetric lengths doubling the elliptic curve security level b , i.e., $kl = nl = ol = 2 \cdot b$. Since however four of the five cipher suites standardized for TLS 1.3 fix 256-bit symmetric parameters, we additionally consider this variant (in Appendix A.3).

Reveal and Test queries q_{RS}, q_{RL}, q_T . Using a generic reduction to single-user signature unforgeability, the number of `REVLONGTERMKEY`, `REVSESSIONKEY`, and `TEST` queries do not affect the bounds; we hence do not place any constraints on them.

Fully-quantitative CK/DFGS bounds for SIGMA/TLS 1.3. For our evaluation, we need to reconstruct fully-quantitative security bounds from the more abstract prior security proofs for SIGMA by Canetti-Krawczyk [19] and for TLS 1.3 by Dowling et al. [24, 26]. We report them in Appendix A.1 and A.2 for reference. In terms of their reduction to underlying DH problems, the CK SIGMA bound reduces to the DDH problem with a loss of $\#U \cdot \#S$, whereas the DFGS TLS 1.3 bound reduces to the strong DH problem with a loss of $(\#S)^2$.

Numerical advantage bounds for CK, DFGS, and ours. We report the numerical advantages for SIGMA and TLS 1.3 based on prior (CK [19], DFGS [24, 26]) and our bounds when ranging over the full parameter space detailed above in Tables 2 and 3 (on pages 49 and 50 in the appendix). Table 1 summarizes the key data points for 128-bit and 192-bit security levels. Our bounds improve on the prior ones across all parameters, with security-level improvements of up to 82 bits for SIGMA and 85 bits for TLS 1.3.

¹¹<https://letsencrypt.org/stats/>, <https://trends.builtwith.com/ssl/traffic/Entire-Internet>

¹²<https://transparencyreport.google.com/https/>

¹³<https://telemetry.mozilla.org/>




Figure 13: Elliptic curve group order (y axis) required to achieve 128-bit (top) and 192-bit (bottom) AKE security for SIGMA and TLS 1.3 based on the CK [19] SIGMA resp. Dowling et al. [24, 26] TLS 1.3 bounds in comparison to our bound (ours giving the same result for SIGMA and TLS 1.3), for a varying number of sessions $\#S$ (x axis), both axis in log-scale. For each security and bound, we plot a smaller-resource “(sm)” setting (with runtime $t = 2^{60}$, number of users $\#U = 2^{20}$, and number of random oracle queries $\#RO = 2^{70}$) and a larger-resource “(lg)” setting (with $t = 2^{80}$, $\#U = 2^{30}$, and $\#RO = 2^{70}$). The group orders of NIST elliptic curves `secp256r1` ($p = 2^{254}$) and `secp384r1` ($p = 2^{384}$) are shown as horizontal lines for context.

Group size requirements based on CK, DFGS, and our bound. Finally, let us take a slightly different perspective on what the prior and our bounds tell us: Figure 13 illustrates the group size required to achieve 128-bit resp. 192-bit AKE security for SIGMA and TLS 1.3 based on the different bounds, dependent on a varying number of sessions $\#S$. (Note that, for the CK SIGMA and our SIGMA and TLS 1.3 bound, the group-size requirement is dominated by the signature scheme (due to a $\#U$ generic loss to single-user unforgeability), whereas for the DFGS TLS 1.3 bound, it is dominated by the $(\#S)^2$ -loss reduction to strong DH.) Prior bounds require the use of larger, less efficient curves to achieve 128-bit security already in the presence of 2^{35} sessions. For large-scale attackers, they furthermore exceed the 192-bit security margin at about 2^{55} sessions. We highlight that, in contrast, with our bounds a curve with 128-bit, resp. 192-bit, security is sufficient to guarantee the same security level for SIGMA and TLS 1.3, for both small- and large-scale adversaries and for very conservative bounds on the number of sessions.

Acknowledgments

We thank Mihir Bellare for insightful discussions and helpful comments.

References

- [1] M. Abdalla, M. Bellare, and P. Rogaway. The oracle Diffie-Hellman assumptions and an analysis of DHIES. In D. Naccache, editor, *CT-RSA 2001*, volume 2020 of *LNCS*, pages 143–158. Springer, Heidelberg, Apr. 2001. 4, 5
- [2] M. Abdalla, F. Benhamouda, and P. MacKenzie. Security of the J-PAKE password-authenticated key exchange protocol. In *2015 IEEE Symposium on Security and Privacy*, pages 571–587. IEEE Computer Society Press, May 2015. 7
- [3] M. Abdalla, P.-A. Fouque, and D. Pointcheval. Password-based authenticated key exchange in the three-party setting. In S. Vaudenay, editor, *PKC 2005*, volume 3386 of *LNCS*, pages 65–84. Springer, Heidelberg, Jan. 2005. 7
- [4] C. Bader, D. Hofheinz, T. Jager, E. Kiltz, and Y. Li. Tightly-secure authenticated key exchange. In Y. Dodis and J. B. Nielsen, editors, *TCC 2015, Part I*, volume 9014 of *LNCS*, pages 629–658. Springer, Heidelberg, Mar. 2015. 5, 6, 12
- [5] C. Bader, T. Jager, Y. Li, and S. Schäge. On the impossibility of tight cryptographic reductions. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 273–304. Springer, Heidelberg, May 2016. 6, 13
- [6] M. Bellare, D. J. Bernstein, and S. Tessaro. Hash-function based PRFs: AMAC and its multi-user security. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 566–595. Springer, Heidelberg, May 2016. 6, 12, 14, 18
- [7] M. Bellare, R. Canetti, and H. Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *37th FOCS*, pages 514–523. IEEE Computer Society Press, Oct. 1996. 12
- [8] M. Bellare and W. Dai. The multi-base discrete logarithm problem: Non-rewinding proofs and improved reduction tightness for identification and signatures. Cryptology ePrint Archive, Report 2020/416, 2020. <https://eprint.iacr.org/2020/416>. 51
- [9] M. Bellare, O. Goldreich, and A. Mityagin. The power of verification queries in message authentication and authenticated encryption. Cryptology ePrint Archive, Report 2004/309, 2004. <http://eprint.iacr.org/2004/309>. 14
- [10] M. Bellare, D. Pointcheval, and P. Rogaway. Authenticated key exchange secure against dictionary attacks. In B. Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 139–155. Springer, Heidelberg, May 2000. 3
- [11] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In D. E. Denning, R. Pyle, R. Ganesan, R. S. Sandhu, and V. Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, Nov. 1993. 5, 7
- [12] M. Bellare and P. Rogaway. Entity authentication and key distribution. In D. R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 232–249. Springer, Heidelberg, Aug. 1994. 3, 5, 7
- [13] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006. 5, 8, 18

- [14] M. Bellare and B. Tackmann. The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 247–276. Springer, Heidelberg, Aug. 2016. 7
- [15] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang. High-speed high-security signatures. In B. Preneel and T. Takagi, editors, *CHES 2011*, volume 6917 of *LNCS*, pages 124–142. Springer, Heidelberg, Sept. / Oct. 2011. 6
- [16] D. Boneh. The decision Diffie-Hellman problem. In *Third Algorithmic Number Theory Symposium (ANTS)*, volume 1423 of *LNCS*. Springer, Heidelberg, 1998. Invited paper. 5, 11
- [17] J. Brendel, M. Fischlin, F. Günther, and C. Janson. PRF-ODH: Relations, instantiations, and impossibility results. In J. Katz and H. Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 651–681. Springer, Heidelberg, Aug. 2017. 6, 7, 47
- [18] R. Canetti and H. Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In B. Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, Heidelberg, May 2001. 3, 6, 11, 46
- [19] R. Canetti and H. Krawczyk. Security analysis of IKE’s signature-based key-exchange protocol. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 143–161. Springer, Heidelberg, Aug. 2002. <http://eprint.iacr.org/2002/120/>. 3, 4, 5, 6, 7, 8, 40, 41, 46, 49, 50
- [20] K. Cohn-Gordon, C. Cremers, and L. Garratt. On post-compromise security. In *2016 Computer Security Foundations Symposium*, pages 164–178. IEEE, 2016. 11
- [21] K. Cohn-Gordon, C. Cremers, K. Gjøsteen, H. Jacobsen, and T. Jager. Highly efficient key exchange protocols with optimal tightness. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 767–797. Springer, Heidelberg, Aug. 2019. 5, 6, 7, 9, 11, 18, 26
- [22] K. Cohn-Gordon, C. Cremers, K. Gjøsteen, H. Jacobsen, and T. Jager. Highly efficient key exchange protocols with optimal tightness – enabling real-world deployments with theoretically sound parameters. Cryptology ePrint Archive, Report 2019/737, 2019. <https://eprint.iacr.org/2019/737>. 5
- [23] D. Diemert and T. Jager. On the tight security of TLS 1.3: Theoretically-sound cryptographic parameters for real-world deployments. *Journal of Cryptology*, 2020. To appear. Available as Cryptology ePrint Archive, Report 2020/726. <https://eprint.iacr.org/2020/726>. 7
- [24] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In I. Ray, N. Li, and C. Kruegel, editors, *ACM CCS 2015*, pages 1197–1210. ACM Press, Oct. 2015. 3, 4, 6, 7, 29, 40, 41, 47, 49, 50
- [25] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. Cryptology ePrint Archive, Report 2015/914, 2015. <http://eprint.iacr.org/2015/914>. 47
- [26] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 draft-10 full and pre-shared key handshake protocol. Cryptology ePrint Archive, Report 2016/081, 2016. <http://eprint.iacr.org/2016/081>. 3, 4, 6, 7, 29, 40, 41, 47

- [27] M. Fischlin and F. Günther. Multi-stage key exchange and the case of Google’s QUIC protocol. In G.-J. Ahn, M. Yung, and N. Li, editors, *ACM CCS 2014*, pages 1193–1204. ACM Press, Nov. 2014. 7, 47
- [28] M. Fischlin and F. Günther. Replay attacks on zero round-trip time: The case of the TLS 1.3 handshake candidates. In *2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017*, pages 60–75, Paris, France, Apr. 26–28, 2017. IEEE. 3, 7, 29
- [29] K. Gjøsteen and T. Jager. Practical and tightly-secure digital signatures and authenticated key exchange. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 95–125. Springer, Heidelberg, Aug. 2018. 5
- [30] S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, Apr. 1988. 12
- [31] F. Günther. *Modeling Advanced Security Aspects of Key Exchange and Secure Channel Protocols*. PhD thesis, Technische Universität Darmstadt, Darmstadt, Germany, 2018. <http://tuprints.ulb.tu-darmstadt.de/7162/>. 7
- [32] D. Harkins and D. Carrel. The Internet Key Exchange (IKE). IETF RFC 2409 (Proposed Standard), 1998. 3, 15
- [33] V. T. Hoang, S. Tessaro, and A. Thiruvengadam. The multi-user security of GCM, revisited: Tight bounds for nonce randomization. In D. Lie, M. Mannan, M. Backes, and X. Wang, editors, *ACM CCS 2018*, pages 1429–1440. ACM Press, Oct. 2018. 7
- [34] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DHE in the standard model. In R. Safavi-Naini and R. Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 273–293. Springer, Heidelberg, Aug. 2012. 6, 7
- [35] C. Kaufman (Ed.). Internet Key Exchange (IKEv2) Protocol. RFC 4306 (Proposed Standard), Dec. 2005. Obsoleted by RFC 5996, updated by RFC 5282. 15
- [36] S. Kent and R. Atkinson. Security Architecture for the Internet Protocol. RFC 2401 (Proposed Standard), Nov. 1998. Obsoleted by RFC 4301, updated by RFC 3168. 3, 15
- [37] H. Krawczyk. SIGMA: The “SIGn-and-MAC” approach to authenticated Diffie-Hellman and its use in the IKE protocols. In D. Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 400–425. Springer, Heidelberg, Aug. 2003. 3, 5, 8, 15, 31
- [38] H. Krawczyk. SIGMA: the ‘SIGn-and-MAC’ approach to authenticated Diffie-Hellman and its use in the IKE protocols, 2003. Full version. <https://webee.technion.ac.il/~hugo/sigma-pdf.pdf>. 15, 31
- [39] H. Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 631–648. Springer, Heidelberg, Aug. 2010. 7, 30
- [40] H. Krawczyk and H. Wee. The OPTLS protocol and TLS 1.3. In *2016 IEEE European Symposium on Security and Privacy*, pages 81–96. IEEE, Mar. 2016. 3
- [41] B. A. LaMacchia, K. Lauter, and A. Mityagin. Stronger security of authenticated key exchange. In W. Susilo, J. K. Liu, and Y. Mu, editors, *ProvSec 2007*, volume 4784 of *LNCS*, pages 1–16. Springer, Heidelberg, Nov. 2007. 6, 11

- [42] A. Langley, M. Hamburg, and S. Turner. Elliptic Curves for Security. RFC 7748 (Informational), Jan. 2016. 4
- [43] U. M. Maurer. Abstract models of computation in cryptography (invited paper). In N. P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Heidelberg, Dec. 2005. 5, 11
- [44] A. Menezes and N. Smart. Security of signature schemes in a multi-user setting. *Designs, Codes and Cryptography*, 33(3):261–274, Nov. 2004. 12
- [45] National Institute of Standards and Technology. FIPS PUB 186-4: Digital Signature Standard (DSS), 2013. 4, 6
- [46] T. Okamoto and D. Pointcheval. The gap-problems: A new class of problems for the security of cryptographic schemes. In K. Kim, editor, *PKC 2001*, volume 1992 of *LNCS*, pages 104–118. Springer, Heidelberg, Feb. 2001. 11
- [47] E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Proposed Standard), Aug. 2018. 3, 6, 8, 15, 29, 47
- [48] V. Shoup. Lower bounds for discrete logarithms and related problems. In W. Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Heidelberg, May 1997. 5, 11, 53

A Evaluation Details

A.1 Fully-quantitative CK SIGMA Bound

Recall our security bound for SIGMA/SIGMA-I from Theorem 5.1:

$$\begin{aligned} & \text{Adv}_{\text{SIGMA-I}}^{\text{KE-SEC}}(t, q_N, q_S, q_{RS}, q_{RL}, q_T) \\ & \leq \frac{q_S^2}{2^{nl+1}} + \frac{q_S^2}{2p} + \frac{q_S^2}{2^{nl} \cdot p} + \text{Adv}_{\mathbb{G}}^{\text{stDH}}(t_{\mathcal{B}_1}, q_{RO}) + \text{Adv}_{\text{PRF}}^{\text{mu-PRF}}(t_{\mathcal{B}_2}, q_S, 3 \cdot q_S) \\ & \quad + \text{Adv}_{\mathbb{S}}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_3}, q_N, q_S, q_S, q_{RL}) + \text{Adv}_{\mathbb{M}}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_4}, q_S, q_S, 1, q_S, 1, 0), \end{aligned}$$

Comparing this bound to the original security proof for SIGMA by Canetti and Krawczyk [19] faces two complications: First, the original proof does not include any concrete-security bounds and merely refers to the DDH assumption (cf. Section 3.1, Definition 3.1) and “standard security notions” for digital signatures, MACs, and PRFs (i.e., single-user EUF-CMA and PRF security). We accordingly have to reconstruct the concrete security bound from their proof, our result given below. Second, the original result is in an extension of the Canetti–Krawczyk (CK) security model for key exchange [18], which is stronger than ours in allowing state-reveal attacks. (Technically, the models are incomparable due to specifying a single (CK) vs. multiple (ours) TEST queries and our model capturing key-compromise impersonation attacks, but we expect the CK-model results could be likewise extended.) Further, the proof in [19] assumes out-of-band unique session identifiers, whereas protocols in practice have to establish those from, e.g., nonces (introducing a corresponding collision bound as in our analysis). We are therefore inherently constrained to compare qualitatively different security properties here.

Let us informally consider a game-based definition of the CK model [18] in the same style as our AKE model (cf. Definition 2.1), capturing the same oracles plus an additional state-reveal oracle, with q_{RST} denoting the number of queries to this oracle, and session identifiers derived from the execution. Translating the SIGMA-I security proof from [19, Theorem 6 (full version)], we obtained the following concrete security bound:

$$\begin{aligned} & \text{Adv}_{\text{SIGMA-I}}^{\text{CK}}(t, q_N, q_S, q_{RS}, q_{RL}, q_{\text{RST}}, q_T) \\ & \leq \frac{q_S^2}{2^{nl+1}} + \text{Adv}_{\mathbb{S}}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_1}, q_N, q_S, q_S, q_{RL}) \quad // \text{sid collision \& property P1} \\ & \quad + q_N \cdot q_S \cdot \left(\text{Adv}_{\mathbb{G}}^{\text{DDH}}(t_{\mathcal{B}_2}) + \text{Adv}_{\text{PRF}}^{\text{mu-PRF}}(t_{\mathcal{B}_5}, 1, 3) \right) \quad // \text{property P2 \dots} \\ & \quad + (q_N + 1) \cdot \text{Adv}_{\mathbb{S}}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_3}, 1, q_S, q_S, 0) + \text{Adv}_{\mathbb{M}}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_4}, 1, 2, 2, 2, 2, 0) \end{aligned}$$

where the number of test queries is restricted to $q_T = 1$ and \mathcal{B}_i (for $i = 1, \dots, 5$) are the described reductions for property P1 and P2 of [19, Theorem 6 (full version)] all running in time $t_{\mathcal{B}_i} \approx t$. For simplicity, we present the above bound in terms of “multi-user” PRF, signature, and MAC advantages for a single user $q_{\text{Nw}} = 1$, which are equivalent to the corresponding single-user advantages (cf. Section 3).

A.2 Fully-quantitative DFGS TLS 1.3 Bound

Recall our bound for TLS 1.3 from Theorem 7.1:

$$\begin{aligned} \text{Adv}_{\text{TLS 1.3}}^{\text{KE-SEC}}(t, q_N, q_S, q_{RS}, q_{RL}, q_T) &\leq \frac{q_S^2}{2^{nl}} + \frac{q_S^2}{2^p} + \frac{q_S^2}{2^{nl} \cdot p} + \text{Adv}_{\text{H}}^{\text{CR}}(t_1) \\ &\quad + 2 \cdot \text{Adv}_{\mathbb{G}}^{\text{stDH}}(t_2, q_{RO}) + \frac{q_{RO} \cdot q_S}{2^{kl-1}} + \text{Adv}_{\text{S}}^{\text{mu-EUF-CMA}}(t_3, q_N, q_S, q_S, q_{RL}) \\ &\quad + \text{Adv}_{\text{HMAC}}^{\text{mu-EUF-CMA}}(t_4, q_S, q_S, 1, q_S, 1, 0), \end{aligned}$$

We compare with the bound of [24, 25, 26], adapted to the final standard [47]. Note that this bound is established in a multi-stage key exchange model [27], here we focus only on the main application key derivation, as in our proof. The DFGS bound needs instantiation through the random oracle only in one step (the PRF-ODH assumption on HKDF.Extract) while other PRF step remain in the standard model. Our proof instead models both HKDF.Extract and HKDF.Expand as random oracles.

$$\begin{aligned} \text{Adv}_{\text{TLS 1.3}}^{\text{DFGS}}(t, q_N, q_S, q_{RS}, q_{RL}, q_T) &\leq \frac{q_S^2}{2^{nl} \cdot p} + q_S \cdot \left(\text{Adv}_{\text{H}}^{\text{CR}}(t_{\mathcal{B}_1}) + q_N \cdot \text{Adv}_{\text{S}}^{\text{mu-EUF-CMA}}(t_{\mathcal{B}_2}, 1, q_S, q_S, 0) \right. \\ &\quad + q_S \cdot \left(\text{Adv}_{\text{HKDF.Extract}, \mathbb{G}}^{\text{dual-snPRF-ODH}}(t_{\mathcal{B}_3}) + \text{Adv}_{\text{HKDF.Expand}}^{\text{mu-PRF}}(t_{\mathcal{B}_4}, 1, 3, 3, 0) \right. \\ &\quad + 2 \cdot \text{Adv}_{\text{HKDF.Expand}}^{\text{mu-PRF}}(t_{\mathcal{B}_5}, 1, 2, 2, 0) + \text{Adv}_{\text{HKDF.Extract}}^{\text{mu-PRF}}(t_{\mathcal{B}_6}, 1, 1, 1, 0) \\ &\quad \left. \left. + \text{Adv}_{\text{HKDF.Expand}}^{\text{mu-PRF}}(t_{\mathcal{B}_7}, 1, 1, 1, 0) \right) \right), \end{aligned}$$

Let us further unpack the PRF-ODH bound. Following Brendel et al. [17], it can be reduced to the strong Diffie–Hellman assumption modeling HKDF.Extract as a random oracle.¹⁴ In this reduction, the single DH oracle query is checked against each random oracle query via the strong-DH oracle, hence establishing the following bound:

$$\text{Adv}_{\text{RO}, \mathbb{G}}^{\text{dual-snPRF-ODH}}(t_{\mathcal{B}_3}, q_{RO}) \leq \text{Adv}_{\mathbb{G}}^{\text{stDH}}(t_{\mathcal{B}_3}, q_{RO}).$$

A.3 Numerical bounds for SIGMA and TLS 1.3

In the following Tables 2 and 3 we evaluate prior AKE bounds against our bounds for SIGMA and TLS 1.3 with a wide variety of parameters. Throughout, we assume that an adversary with running time t makes no more than $t \cdot 2^{-10}$ queries to its random oracles. We target the bit-security of whatever curve we use; this means that for b bits of security we want an advantage of $t/2^b$. If a bound does not achieve this goal, we color it red. We consider runtimes between 2^{40} and 2^{80} , a total number of users between to vary between 2^{20} and 2^{30} , and a total number of sessions between 2^{35} and 2^{55} (cf. Section 8 for a discussion of these parameter choices). We evaluate these parameters in relation to all of the elliptic curve groups standardized for use with TLS 1.3.

In Table 2, we scale the symmetric (nonce, key, tag) length with the curve’s security level b , setting the symmetric length to $2 \cdot b$ bits. In Table 3, we instead fix the symmetric lengths to 256 bits. Considering the latter setting is particularly relevant because all but one TLS 1.3 cipher

¹⁴The same paper suggests that a standard-model instantiation of the PRF-ODH assumption via an algebraic black-box reduction to common cryptographic problems is implausible.

suite (TLS_AES_256_GCM_SHA384) allow only 256-bit symmetric lengths. In this case, for larger groups, MAC and PRF advantages dominate the strong Diffie–Hellman advantage, and our bounds in Table 3 hence do not achieve higher levels of security beyond 192 bits; for those, symmetric lengths will need to scale as in Table 2. Notably though, with symmetric parameters at 256 bits, the prior CK SIGMA and DFGS TLS 1.3 bounds are significantly worse, especially failing to meet the 192-bit security level with curve `secp384r1` which our bounds satisfy for most parameters.

G₀

INITIALIZE():

- 1 $p \leftarrow |\mathbb{G}|$; $E \xleftarrow{\$} \text{Bijections}(\mathbb{Z}_p, \mathbb{G})$
- 2 $\mathcal{K} \leftarrow E(0)$; $g \leftarrow E(1)$
- 3 $x, y \xleftarrow{\$} \mathbb{Z}_p^*$; $X \leftarrow E(x)$; $Y \leftarrow E(y)$
- 4 $GL \leftarrow \{\mathcal{K}, g, x, y\}$
- 5 return (\mathcal{K}, g, x, y)

OP(A, B, sgn):

- 6 If $A \notin GL$ or $B \notin GL$ then return \perp
- 7 $c \leftarrow E^{-1}(A) \text{ sgn } E^{-1}(B) \pmod p$
- 8 $C \leftarrow E(c)$; $GL \leftarrow GL \cup \{C\}$
- 9 Return C

stDH(A, B):

- 10 If $A \notin GL$ or $B \notin GL$ then return \perp
- 11 $z \leftarrow x \cdot E^{-1}(A) \pmod p$
- 12 $Z \leftarrow E(z)$
- 13 Return $[Z = B]$

FINALIZE(Z):

- 14 If $Z \notin GL$ then return **false**
- 15 $z \leftarrow x \cdot y \pmod p$; return $[Z = E(z)]$

Figure 14: Game G₀ of the stDH proof.

B Proof of the Strong Diffie–Hellman GGM Bound

We begin by giving a code-based game for the strong Diffie–Hellman problem in the generic group model. First, we establish some preliminaries, using the setting and notation of [8]. Let \mathbb{G} be an arbitrary set of strings with prime order p , and let $E : \mathbb{Z}_p \rightarrow \mathbb{G}$ be a bijection, called the encoding function. For any two strings $A, B \in \mathbb{G}$, we define the operation $A \text{ OP}_E B = E(E^{-1}(A) + E^{-1}(B) \pmod p)$. The set \mathbb{G} is a group with respect to this operation, and it is isomorphic to \mathbb{Z}_p . Therefore, \mathbb{G} has the identity $E(0)$, and it is generated by $E(1)$.

In the generic group model, we wish for the adversary to compute group operations only through an oracle OP . We accomplish this by picking the encoding function E at random and keeping it secret; then providing oracle access to OP_E through OP . In this model, we can give a sequence of games bounding the advantage of any adversary \mathcal{A} that makes t queries to the OP oracle and q queries to the stDH oracle.

Game 0. This first game formalizes the strong Diffie–Hellman problem in the generic group model. Note that for any $a \in \mathbb{Z}_p$, a is the discrete logarithm of the group element $E(a)$.

It follows that

$$\text{Adv}_{\mathbb{G}}^{\text{stDH}}(t, q_{\text{SDH}}) = \Pr[\text{G}_0 \Rightarrow 1].$$

Game 1. In Game G₁, we change the internal notation of the game. First, for clarity, we assume the adversary queries its oracles only on valid inputs (elements of GL). Instead of representing each element of \mathbb{G} with an element of \mathbb{Z}_p , we use a vector over \mathbb{Z}_p^3 . We define the basis vectors $\vec{e}_1 := (1, 0, 0)$, $\vec{e}_2 := (0, 1, 0)$, and $\vec{e}_3 := (0, 0, 1)$. We map \mathbb{Z}_p^3 to \mathbb{Z}_p by taking the inner product with the vector $(1, x, y)$. (Effectively, we are representing each element of \mathbb{Z}_p as a linear combination modulo p of 1, x , and y .) We cache the map from \mathbb{Z}_p^3 to \mathbb{G} induced by this transformation in a table TV and its inverse map in a table TI . From this point forward, we will assume without loss of generality that all queries to the OP and stDH oracles are valid (meaning their inputs are valid group elements in GL).

Although one element of \mathbb{G} may now have multiple representations, the bilinearity of the inner product ensures that the view of the adversary is not changed, and $\Pr[\text{G}_1] = \Pr[\text{G}_0]$.

G₁

INITIALIZE():

- 1 $p \leftarrow |\mathbb{G}|$; $E \xleftarrow{\$} \text{Bijections}(\mathbb{Z}_p, \mathbb{G})$
- 2 $k \leftarrow 0$; $\mathcal{K} \leftarrow \text{VE}(\vec{0})$; $g \leftarrow \text{VE}(e_1)$
- 3 $x, y \xleftarrow{\$} \mathbb{Z}_p^*$; $\vec{x} \leftarrow \mathcal{K}, x, y$
- 4 $X \leftarrow \text{VE}(e_2)$; $Y \leftarrow \text{VE}(e_3)$
- 5 return (\mathcal{K}, g, x, y)

OP(A, B, sgn):

- 6 $\vec{c} \leftarrow \text{VE}^{-1}(A) \text{sgn } \text{VE}^{-1}(B) \pmod p$
- 7 $C \leftarrow \text{VE}(\vec{c})$; return C

VE(\vec{t}):

- 1 If $TV[\vec{t}] \neq \perp$ then return $TV[\vec{t}]$
- 2 $k \leftarrow k + 1$; $\vec{t}_k \leftarrow \vec{t}$
- 3 $v \leftarrow \langle \vec{t}, \vec{x} \rangle$; $C \leftarrow E(v)$; $GL \leftarrow GL \cup \{C\}$
- 4 $TV[\vec{t}] \leftarrow C$; $TI[C] \leftarrow \vec{t}$
- 5 Return $TV[\vec{t}]$

stDH(A, B):

- 8 $\vec{a} \leftarrow \text{VE}^{-1}(A)$; $\vec{b} \leftarrow \text{VE}^{-1}(B)$
- 9 Return $[\text{VE}(x\vec{a}) = B]$

FINALIZE(Z):

- 10 Return $[\text{VE}(xe_3) = Z]$

VE⁻¹(C):

- 1 Return $TI[C]$

Figure 15: Game G₁ of the stDH proof.

Game 2. Next, we replace the random encoding function E with a lazily sampled encoding represented by table TV for the forward direction and TI for the backward direction. Because we want our encoding to be one-to-one, we sample from the set $\mathbb{G} \setminus GL$. This assigns a unique element of \mathbb{G} to each vector \vec{t} . However, as we've noted, each integer in \mathbb{Z}_p has multiple representations in \mathbb{Z}_p^3 . If two representations of the same integer are submitted to the encoding algorithm VE , we set a **bad** flag and program the encoding table to maintain consistency.

We also change the format of the check in the stDH oracle. Since $\text{VE}(x\vec{a}) = B = \text{VE}(\vec{b})$ if and only if $\langle x\vec{a}, \vec{x} \rangle = \langle \vec{b}, \vec{x} \rangle$, we return **true** if the latter condition holds and **false** otherwise. These two conditions are equivalent, so $\Pr[\text{G}_2] = \Pr[\text{G}_1]$.

Game 3. In this game, we stop programming the encoding table after the bad flag is set. Let F_1 denote the event that G_3 sets the **bad** flag at any point. By the fundamental lemma of game playing, $\Pr[\text{G}_2] \leq \Pr[\overline{F_1}] + \Pr[F_1]$.

Game 4. Game G_4 removes the now-redundant bad flag, but the FINALIZE oracle now returns **true** if at any point in game G_3 the bad flag would have been set (i.e. if event F_1 occurs). Otherwise, all oracles behave exactly as they did in G_3 . It follows that $\Pr[\text{G}_3 \text{ and } \overline{F_1}] + \Pr[F_1] \leq \Pr[\text{G}_4]$.

Additionally, in the stDH oracle, we separate out checking for trivial queries: if the adversary computed $A = g^a$ and $B = X^a$ for an integer a of their choosing. If this is so, then $\vec{a} = ae_1$ and $\vec{b} = ae_2$, so $\langle x\vec{a}, \vec{x} \rangle = xa = \langle \vec{b}, \vec{x} \rangle$, so may return **true**. If the query is nontrivial but should still return true according to our previous condition, we set a **bad₂** flag. This does not change the oracle's response to any query, so the above bound still holds.

Game 5. In Game G_5 , we no longer return **true** in the stDH oracle after the **bad₂** flag is set. This makes the second check redundant and has the effect that the stDH oracle's behavior is no longer dependent on the value of either x or \vec{x} . Let event F_2 denote the event that G_5 sets the **bad₂** flag. By the fundamental lemma of game playing, $\Pr[\text{G}_4] \leq \Pr[\text{G}_5 \text{ and } \overline{F_2}] + \Pr[F_2]$.

G₂, G₃

stDH(A, B):

- 1 $\vec{a} \leftarrow \text{VE}^{-1}(A); \vec{b} \leftarrow \text{VE}^{-1}(B)$
- 2 If $\langle x\vec{a}, \vec{x} \rangle = \langle \vec{b}, \vec{x} \rangle$ then return **true**
- 3 Return **false**

VE(\vec{t}):

- 1 If $TV[\vec{t}] \neq \perp$ then return $TV[\vec{t}]$
- 2 $C \leftarrow \mathbb{G} \setminus GL$
- 3 If $(\exists \vec{s} : TV[\vec{s}] \neq \perp \text{ and } \langle \vec{t}, \vec{x} \rangle = \langle \vec{s}, \vec{x} \rangle)$
- 4 then **bad** \leftarrow **true**; $C \leftarrow TV[\vec{s}]$
- 5 $k \leftarrow k + 1; \vec{t}_k \leftarrow \vec{t}$
- 6 $GL \leftarrow GL \cup \{C\}$
- 7 $TV[\vec{t}] \leftarrow C; TI[C] \leftarrow \vec{t}$
- 8 Return $TV[\vec{t}]$

G₄, G₅

stDH(A, B):

- 1 $\vec{a} \leftarrow \text{VE}^{-1}(A); \vec{b} \leftarrow \text{VE}^{-1}(B); \vec{a} \leftarrow \vec{a}[1]$
- 2 If $\vec{a} = a\vec{e}_1$ and $\vec{b} = a\vec{e}_2$ then return **true**
- 3 If $\langle x\vec{a}, \vec{x} \rangle = \langle \vec{b}, \vec{x} \rangle$ then **bad**₂ \leftarrow **true**; **return true**
- 4 Return **false**

VE(\vec{t}):

- 1 If $TV[\vec{t}] \neq \perp$ then return $TV[\vec{t}]$
- 2 $C \leftarrow \mathbb{G} \setminus GL$
- 3 $k \leftarrow k + 1; \vec{t}_k \leftarrow \vec{t}$
- 4 $GL \leftarrow GL \cup \{C\}$
- 5 $TV[\vec{t}] \leftarrow C; TI[C] \leftarrow \vec{t}$
- 6 Return $TV[\vec{t}]$

FINALIZE(Z):

- 5 If $\exists i, j : 1 \leq i < j \leq k$ and $\langle \vec{t}_i - \vec{t}_j, \vec{x} \rangle = 0$
- 6 then return **true**
- 7 Return $([\text{VE}(xe_3) = Z])$

Figure 16: Left: Games G₂ (changes highlighted in gray) and G₃ (changes highlighted in frames) of the strong Diffie–Hellman proof. Right: Games G₄ and G₅.

Game 6. In Game G₆, we remove the redundant check and **bad** flag from the **stDH** oracle, and in the **FINALIZE** oracle we return **true** whenever the **bad**₂ flag would have been set in G₅. Otherwise all oracles behave precisely as they did in G₅. It follows that $\Pr[G_5 \text{ and } \overline{F_2}] + \Pr[F_2] \leq \Pr[G_6]$. We also move the initialization of variables x , y , and \vec{x} from **INITIALIZE** to **FINALIZE**. Since these variables are not used by any oracle but **FINALIZE**, this does not change the view of the adversary.

At this point, we can collect the bounds from each gamehop to see that

$$\text{Adv}_{\mathbb{G}}^{\text{stDH}}(t, q_{\text{sDH}}) \leq \Pr[G_6].$$

Therefore we analyze the advantage of an adversary in game G₆. Since the adversary made at most t queries to **OP**, each of which incremented k once, $k \leq t + 3$. We can separately analyze each condition. We know that x and y are chosen independently of all t_i . For each index i , let F_i be the bivariate linear polynomial over \mathbb{Z}_p whose coefficients are given by the vector \vec{t}_i . Then $\langle \vec{t}_i - \vec{t}_j \rangle = 0$ only if $(1, x, y)$ is a root of $F_i - F_j$. Using Lemma 1 of [48], the probability of this event is at most $1/p$.

For the second condition; we see that $\langle x\vec{t}_i - \vec{t}_j \rangle = 0$ only if $(1, x, y)$ is a root of $XF_i - F_j$, which is a bivariate quadratic polynomial over \mathbb{Z}_p . Again Using Lemma 1, this occurs with probability at most $2/p$.

If neither event occurs, then the adversary wins only if $[\text{VE}(xe_3) = Z]$. Because the second condition failed, we know that (xe_3) is not an entry in table TV . Therefore $\text{VE}(xe_3)$ is sampled uniformly at random, and it equals Z with probability at most $1/p$. Then by the union bound, $\Pr[G_6] \leq (3(t + 3)^2 + 1)/p$. Collecting the bounds gives the theorem statement for all $t > 3$. \square

G₆

INITIALIZE():

- 1 $p \leftarrow |\mathbb{G}|$;
- 2 $k \leftarrow 0$; $\mathcal{K} \leftarrow \text{VE}(\vec{0})$; $g \leftarrow \text{VE}(\vec{e}_1)$
- 3 $X \leftarrow \text{VE}(\vec{e}_2)$; $Y \leftarrow \text{VE}(\vec{e}_3)$
- 4 **return** (\mathcal{K}, g, x, y)

FINALIZE(Z):

- 5 $x, y \xleftarrow{\$} \mathbb{Z}_p^*$; $\vec{x} \leftarrow (\mathcal{K}, x, y)$
- 6 **If** $\exists i, j : 1 \leq i < j \leq k$ and $\langle \vec{t}_i - \vec{t}_j, \vec{x} \rangle = 0$
- 7 **then return true**
- 8 **If** $\exists i, j : 1 \leq i < j \leq k$
- and $\langle x\vec{t}_i - \vec{t}_j, \vec{x} \rangle = 0$ or $\langle x\vec{t}_j - \vec{t}_i, \vec{x} \rangle = 0$
- 9 **then return true**
- 10 **Return** $([\text{VE}(x\vec{e}_3) = Z])$

stDH(A, B):

- 11 $\vec{a} \leftarrow \text{VE}^{-1}(A)$; $\vec{b} \leftarrow \text{VE}^{-1}(B)$; $a \leftarrow \vec{a}[1]$
- 12 **If** $(\vec{a} = a\vec{e}_1$ and $\vec{b} = a\vec{e}_2)$ **then return true**
- 13 **Return** 0

Figure 17: Game G₆ (changes highlighted in gray) of the strong Diffie–Hellman proof.