

An abridged version of this paper is under submission. This is the full version.

Downgrade Resilience in Key-Exchange Protocols

KARTHIKEYAN BHARGAVAN*

CHRISTINA BRZUSKA†

CÉDRIC FOURNET, MARKULF KOHLWEISS, SANTIAGO ZANELLA-BÉGUELIN‡

MATTHEW GREEN§

January 2016

Abstract

Key-exchange protocols such as TLS, SSH, IPsec, and ZRTP are highly configurable, with typical deployments supporting multiple protocol versions, cryptographic algorithms and parameters. In the first messages of the protocol, the peers negotiate one specific combination: the *protocol mode*, based on their local configurations. With few notable exceptions, most cryptographic analyses of configurable protocols consider a single mode at a time. In contrast, downgrade attacks, where a network adversary forces peers to use a mode weaker than the one they would normally negotiate, are a recurrent problem in practice.

How to support configurability while at the same time guaranteeing the preferred mode is negotiated? We set to answer this question by designing a formal framework to study downgrade resilience and its relation to other security properties of key-exchange protocols. First, we study the causes of downgrade attacks by dissecting and classifying known and novel attacks against widely used protocols. Second, we survey what is known about the downgrade resilience of existing standards. Third, we combine these findings to define downgrade security, and analyze the conditions under which several protocols achieve it. Finally, we discuss patterns that guarantee downgrade security by design, and explain how to use them to strengthen the security of existing protocols, including a newly proposed draft of TLS 1.3.

*Inria Paris-Rocquencourt. E-mail: karthikeyan.bhargavan@inria.fr

†Hamburg University of Technology. E-mail: brzuska@tuhh.de

‡Microsoft Research. E-mail: {fournet,markulf,santiago}@microsoft.com

§Johns Hopkins Univeristy. E-mail: mgreen@cs.jhu.de

Contents

1	Introduction	3
1.1	Motivating example	3
1.2	Overview of our approach	4
1.3	Summary of our results	6
1.4	Outline of the paper	6
2	Modeling Multi-Mode Key-Exchanges	7
2.1	Unique identifiers and partnering	7
2.2	Multi-mode authentication	8
2.3	Key-indistinguishability and privacy	8
2.4	Instantiating our model for SIGMA-N	9
3	Defining Downgrade Resilience	9
3.1	Downgrade resilience of SIGMA-N	11
3.2	Downgrade resilience and multi-mode security	11
3.3	Downgrade secure sub-protocols	12
3.4	Downgrade security by whitelisting	13
4	Secure SHell	14
4.1	SSHv2 is partnering and downgrade secure	15
4.1.1	Agile hash functions and signatures	15
4.1.2	Partnering security	16
4.1.3	Downgrade security	17
5	Internet Key Exchange	18
5.1	IKEv1 does not prevent downgrades	18
5.2	IKEv2 does not prevent downgrades	22
5.2.1	IKEv2 with signatures	23
5.2.2	IKEv2 with EAP client authentication	24
5.3	Version downgrades from IKEv2 to IKEv1	25
6	Z Real-Time Protocol	25
6.1	ZRTP does not prevent downgrades	26
7	Transport Layer Security	27
7.1	Negotiation in TLS 1.2	28
7.2	TLS 1.2 does not prevent downgrades	29
7.3	On downgrade protection in Draft 10 of TLS 1.3	31
8	Related Work	36
9	Conclusion and Future Work	37

1 Introduction

Popular protocols such as TLS, SSH and IPsec as used in practice do not fit a simple textbook definition of a key-exchange protocol, where the state machine, cryptographic algorithms, parameters and message formats are all fixed in advance. Rather, these modern protocols feature cryptographic agility, which provides for configurable selection of multiple protocol and cipher modes, so that the key exchange actually executed between two peers depends on a negotiation phase embedded in the exchange.

Agility has proven important in securing real-world protocol implementations. For example, in the wake of recent vulnerability disclosures in TLS [25, 2–4, 41, 26, 10], network operators reacted by updating client and server configurations to disable weak algorithms and protocol versions. Moreover, experience shows that when sufficient agility is not present within a single protocol, application developers construct their own *ad hoc* negotiation mechanisms, for example, by sequentially attempting connections with different versions of a protocol and “falling back” to the best one supported [40].

Unfortunately, support for algorithm agility opens up opportunities for *downgrade attacks*, where an active network adversary interferes with the negotiation, causing honest peers to complete a key exchange, albeit using a mode that is weaker than the one they would have used on their own. Such attacks have been identified in a number of protocols, most famously in the early versions of the SSL protocol [44] and even in recent versions of TLS [2, 40].

Surprisingly, there has been relatively little formal work around the security of negotiation in modern cryptographic protocols. Several recent works formally prove the security of different aspects of TLS and SSH. Some [28, 33] only model a single mode at a time. Some [12, 13] do model negotiation of *weak* algorithms, but do not guarantee negotiation of the preferred mode. Some others [9, 23] consider only interactions where both parties have secure configurations. For this reasons, all of these works overlook certain downgrade attacks that occur when one party supports an insecure mode.

This is concerning because negotiation has proven to be fertile ground for attacks, e.g. [10, 44, 2], and because recent Internet-wide scans have revealed the prevalence of hosts supporting insecure protocol modes [43, 2]. In this setting, it is insufficient to restrict our attention to situations where both parties support secure configurations.

In this work we aim to address this situation by systematically investigating the problem of *downgrade resilience* in cryptographic protocols.

1.1 Motivating example

We begin with a simple motivating example: we adapt the SIGMA protocol of Krawczyk [32] by adding a naïve extension intended to negotiate Diffie-Hellman groups: In the first message, A proposes a list of groups it supports; in the second message, B indicates which of these groups should be used in the exchange. The modified protocol appears in Fig. 1. The goal of the protocol is to compute session keys (k_m, k_s) .

Under normal circumstances, the protocol succeeds in correctly selecting a group. However, consider a scenario where both participants support both strong and weak groups. B 's signature authenticates the chosen group, but not A 's proposal. This leads to a downgrade attack (see Fig. 2) similar to the Logjam attack on TLS [2], where an attacker can break the session keys at leisure and compromise the connection.

Protocol designers have adopted a number of techniques to prevent such downgrade attacks. Based on a review of deployed protocols, we identify three common patterns. In the first, exemplified by SSH, protocol designers assume the existence of strong signing keys shared between the two parties, and use these keys to authenticate all negotiation messages, either at the time they are transmitted, or after the fact.

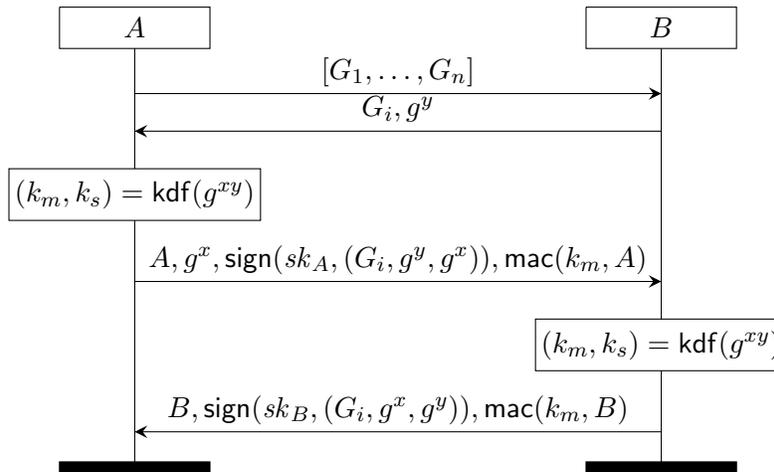


Figure 1: SIGMA-N: Basic SIGMA [32] with group negotiation

In a second pattern, exemplified by TLS, designers transmit unauthenticated protocol messages, perform a key exchange, and then subsequently use the resulting shared secrets to retroactively authenticate the negotiation messages. The final approach relies on *whitelisting* certain modes, and is best exemplified by Google’s TLS False Start proposal [34], which is being codified as part of TLS 1.3 [22].

Each approach has various advantages and disadvantages. The devil is often in the details: each protocol is sensitive to the precise nature of the implementation, e.g. the inputs of authentication functions, or the specifics of what a *valid* mode is for whitelisting. As a concrete example, modern versions of TLS-DHE fail to sign the identity of the ciphersuite chosen by a server, leading to cross-protocol attacks [36, 2]. Similarly, TLS False Start relies solely on ciphersuite identifier (rather than more detailed information such as key strength) in its selection of which modes to whitelist, which converts the online attack of Adrian et al. [2] into an offline one.

1.2 Overview of our approach

We give a definition and a theorem for downgrade resilience that model the following intuitive and desirable property for deployed key-exchange protocols:

To prevent an attack on a particular protocol mode, it is sufficient to deactivate the configurations that lead to its negotiation.

Our work builds on the definitions of Bhargavan et al. [13], used to model security in MITLS, a reference implementation of the TLS standard. A fundamental difference between these definitions and previous work is that they attempt to model entire deployed protocols. This requires a definition of security cognizant of the fact that some aspects (modes) of the protocol may be insecure. To deal with this, the definitions of [13] incorporate predicates determining modes that are expected to provide security guarantees, e.g., key indistinguishability. This approach allows to define security when secure modes are chosen, yet *tolerates* the existence of insecure modes.

One limitation of these definitions is that they do not take into account *how* modes are chosen. In a protocol secure under the MITLS framework, two parties under adversarial influence may arrive at an

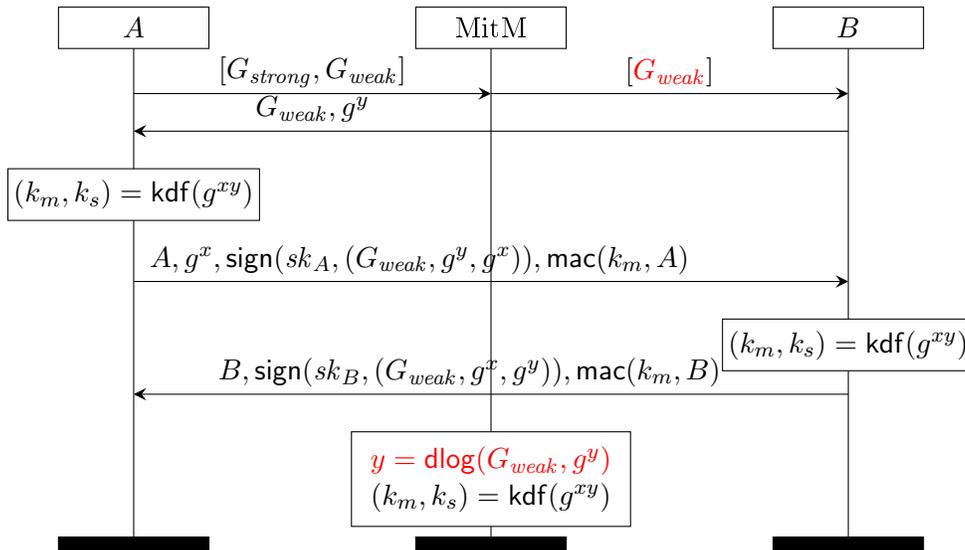


Figure 2: Man-in-the-Middle downgrade attack on SIGMA-N

insecure mode even when otherwise they would use a secure mode. In theory each party can detect and react to the negotiation of an insecure mode, e.g., by terminating the protocol execution. Nonetheless, this does not guarantee that the preferred common mode is selected. Our solution is to incorporate *downgrade resilience* in our security definitions, to ensure that an adversary cannot force the selection of another mode than the preferred one.

We consider protocols between an Initiator and a Responder. These two parties each have their own local static *configurations*, expressing their preferences and their intent to negotiate a shared protocol *mode*. To define *downgrade resilience* formally, we introduce a downgrade protection predicate DP that operates on pairs of configurations (analogous to MITLS predicates on modes), and that identifies pairs of configurations from which we expect downgrade resilience. We also introduce a function $Nego$ that maps two opposite-role configurations to the protocol mode that should be negotiated in the absence of active adversaries. Intuitively, our definition says that a protocol is *downgrade secure* if two peers starting from configurations satisfying DP can only negotiate the mode determined by $Nego$, even in the presence of an active adversary.

By way of example, a specific instantiation of $Nego$ for the TLS protocol might determine that two TLS peer configurations would normally result in the negotiation of TLS 1.2 in combination with a ciphersuite such as `DHE-RSA-AES256-GCM-SHA384` with a 2048-bit Diffie-Hellman modulus. However, if a server supports an insecure mode, such as a `DHE-EXPORT` ciphersuite, an adversary might force the pair to downgrade to this mode [2]. This shows that without additional countermeasures, TLS 1.2 does not meet our definition. On the other hand, protocols with only one possible mode are obviously secure. The challenge we address in this paper is to consider agile protocols that support multiple modes (e.g., ciphersuites, versions).

To apply our definition to real-world protocols, we adopt the following approach. Rather than analyzing a protocol in its entirety, we first extract a core negotiation sub-protocol, which captures the main downgrade-protection mechanisms of the larger protocol. We next prove that this sub-protocol is *complete* for downgrade security, in the sense that an adversary that succeeds in downgrading the full protocol will also succeed in downgrading the sub-protocol. This technique of *lifting* security from the sub-protocol to

the main protocol was previously employed by Bergsma et al. [9] to prove multi-ciphersuite security.

In our analysis we restrict ourselves to the manual extraction of sub-protocols that only cover specific families of modes, e.g., signature-based modes or pre-shared key modes, while some of our attacks are cross family attacks. Proving the absence of cross-family attacks requires either to consider more complex sub-protocols that encompass several families, or to study families independently and prove a composition theorem similar to that in Bergsma et al. [9]. Our work is a stepping stone in this direction, and our results are readily applicable in situations where peer configurations are from the same family.

In developing our definitional approach we did not use any automated protocol analysis tools. We could have written machine-checked proofs to show that our sub-protocols correctly abstract full protocols, but doing so requires to write formal descriptions of protocols in excruciating detail, a significant amount of work on its own. Proving properties such as downgrade security is at the limit of what current tools can handle [7]. However costly, we think this is a worthwhile endeavor for standards, and we plan to extend our formalization of TLS in F^* to verify our results on downgrade security.

1.3 Summary of our results

Our primary contribution is a novel downgrade security definition for key-exchange protocols. We devise a methodology to analyze the downgrade security of a complex protocol by abstracting away irrelevant details and studying only the core negotiation sub-protocol.

We demonstrate the relevance of our definition and the applicability of our methodology by analyzing the downgrade security of several exemplary real-world protocols, namely TLS, SSH, IPsec and ZRTP. We do so by taking in their standard specifications and extracting appropriate core negotiation sub-protocols. Our analysis identifies known and novel attacks as well as sufficient conditions under which these protocols achieve downgrade security. These conditions inform users of these protocols as to how to restrict host configurations to best avoid downgrade attacks.

The following are concrete novel contributions:

- We describe two new attacks on protocol standards: 2 on IKEv2 and 2 on ZRTP.
- We confirm the conclusion evidenced by recent attacks: TLS versions up to 1.2 are not generally downgrade secure.
- We prove a downgrade security theorem for SSHv2 with `publickey` client authentication that is stronger than previous results. This stems from both peers signing all the messages that determine the protocol mode.
- We show that although TLS 1.3 Draft 10 [22] includes a mandatory server-side message for signing the handshake transcript, this does not prevent downgrades to earlier versions of TLS or non-preferred groups. Informed by this analysis, we define and prove two new downgrade protection mechanisms. The concrete countermeasures, designed jointly with the core TLS 1.3 working group, have been included in Draft 11.

1.4 Outline of the paper

The remainder of this paper proceeds as follows. In §2 we introduce the terminology used throughout and we provide a primer on security definitions for key exchange protocols. We also introduce downgrade resilience, formally defined in §3. In §4 through §7 we apply these definitions to analyze the security of SSH, IPsec IKE, ZRTP and TLS. We survey related work in §8.

2 Modeling Multi-Mode Key-Exchanges

The two main security aspects of popular key-exchange models [13, 9, 24] are entity authentication and key-indistinguishability [8]. Our focus lies on considering multi-mode protocols and incorporating the negotiation of the mode into the security model.

A key exchange protocol Π is a two-party protocol with an initiator role I and a responder role R (sometimes called client and server). The adversary interacts with multiple sessions of the protocol. Each session π maintains variables in a local state and makes assignments to them before sending or after receiving a message. We write $\pi.x$ for the value of variable x in session π . We will consider the following variables:

- $\pi.cfg$ initial configuration (including the role);
- $\pi.uid$ unique identifier of the session;
- $\pi.mode$ negotiated *mode* (including long-term identities);
- $\pi.key$ session key;
- $\pi.complete$ flag set when the session completes successfully.

Each session assigns a value to each variable only once, typically in the order given above. The configuration variable $\pi.cfg$ is assigned when a session is created and contains other variables, including one for the session role. We use $\pi.role$ as shorthand for $\pi.cfg.role$ and let $\bar{I} = R, \bar{R} = I$.

An adversary interacts with sessions via queries to oracles. A query $\pi \leftarrow \text{Init}(cfg)$ initializes a session. Recall that cfg determines *role* and furthermore, in the setting where we have symmetric or public keys, cfg will contain handles to those keys. A query $m_{out} \leftarrow \text{Send}(\pi, m_{in})$ sends a message m_{in} to session π , which processes it to update its local state and output an ongoing message m_{out} . A query $k \leftarrow \text{Reveal}(\pi)$ reveals the session key of π , i.e., returns the value of $\pi.key$. There are several variants for handling long-term keys and other authentication mechanisms as well as corruption settings, and each of those requires different variables and oracles. As these settings are mostly standard and orthogonal to our definition, we leave those details deliberately unspecified for now and get back to them in Section 4. Note that our definitions only become complete once we add the specifics of long-term keys or other authentication mechanisms.

2.1 Unique identifiers and partnering

The goal of a key exchange protocol is to *match* two sessions of two different parties so that they compute the same key and agree on the algorithms and authentication setting. We say that two sessions match if and only if they derive the same session key [31]. For defining downgrade security, we rely on the weaker notion of *partnering*, based on unique identifiers—at *most* two sessions may assign the same value to the session variable *uid*.

Definition 1 (Partnering). *Sessions π and π' are partnered if $\pi'.role = \overline{\pi.role}$ (they have opposite roles) and $\pi'.uid = \pi.uid$. A session π is unpartnered when there is no such π' .*

For example, in TLS, a suitable *uid* is the pair of nonces sent by the client and server in their *Hello* messages. To guarantee partnering upon completion, a protocol needs to protect the messages that influence the *uid* against man-in-the-middle attacks. As some configurations—in particular those where entity authentication is optional—do not protect against man-in-the-middle attacks, our definition depends on a predicate *PS* that indicates configurations that provide **Partnering Security**. Typically, these configurations demand peer authentication.

Definition 2 (Partnering security). *The advantage $\text{Adv}_{\Pi, PS}^{\text{partnering}}(\mathcal{A})$ of adversary \mathcal{A} against the partnering security of Π is the probability that, when \mathcal{A} interacts with protocol Π , there is an unpartnered session π such that $\pi.\text{complete} = \text{true}$ and $PS(\pi.\text{cfg})$ holds.*

2.2 Multi-mode authentication

We now define authentication for protocols in which long-term identifiers (e.g. public keys or pre-shared key identifiers) of peers are themselves negotiated. This is also known as the *post-specified peer setting* [18] and the type of authentication (e.g. mutual or bilateral) is determined as part of the negotiation [24]. We incorporate entity identifiers eid_r for $r \in \{I, R\}$ and authentication type, together with the negotiated cryptographic algorithms in the *mode* variable. We write eid_r as shorthand for $\text{mode}.\text{eid}_r$. As algorithms can be weak, keys can be compromised, and authentication can be unilateral, whether participants get guarantees depends crucially on the outcome of negotiation.

While the predicate PS for partnering is defined over configurations fixed upon creation of a session, our authentication definition depends on a predicate $\text{Auth}(\text{mode}, r)$, which holds when *mode* is expected to authenticate role \bar{r} .

Authentication classically guarantees agreement on the variables of authenticated peers.

Definition 3 (Agreement). *A session π agrees with π' on x when $\pi.x = \pi'.x$ or $\pi'.x = \perp$. For agreement on a set X we require that π agrees with π' on all $x \in X$.*

Definition 4 (Multi-mode authentication). *A session π completes maliciously for X when $\pi.\text{complete} = \text{true}$ but there is no partnered session π' matching π that agrees with π on X .*

The advantage $\text{Adv}_{\Pi, \text{Auth}, X}^{\text{mm-auth}}(\mathcal{A})$ of an adversary \mathcal{A} against the multi-mode authentication security with agreement on X of protocol Π is the probability that, when \mathcal{A} interacts with protocol Π , a session π completes maliciously for X and $\text{Auth}(\pi.\text{mode}, \pi.\text{role})$ holds.

Let $r = \pi.\text{role}$. Note that $\text{Auth}(\pi.\text{mode}, r)$ typically includes the requirement that the long term key $\pi.\text{eid}_{\bar{r}}$ of the peer is honest. If, as in SIGMA-N, the mode is secure against *key-compromise impersonation* attacks [29] then $\pi.\text{eid}_{\bar{r}}$ need not be honest. In addition, the predicate Auth models concurrent mixed-mode authentication. A protocol *mode* provides mutual-authentication if $\text{Auth}(\pi.\text{mode}, r)$ holds regardless of r . It provides server-only authentication if only $\text{Auth}(\pi.\text{mode}, I)$ holds, i.e., only clients get guarantees.

Observe that the authentication mode is itself negotiated. The same long-term keys eid_r routinely appear in different modes and protocols may assign the same *key* in different modes. Agreement on *mode* and other variables may be critical for higher-level protocols; *mode* may include record algorithms and using the same keys with different algorithms may lead to agile security problems. In any case it contains the entity identifiers that should be in agreement to avoid identity confusion attacks [21]. As we will see, protocols need to have sufficient downgrade resilience to guarantee that the preferred authentication mode is negotiated.

2.3 Key-indistinguishability and privacy

Classical definitions of key indistinguishability are parameterized by a freshness predicate Fresh that determines the sessions with uncompromised keys. Key indistinguishability requires that for fresh sessions, an adversary cannot tell apart the real session key from a random one.

For SIGMA-N, a suitable Fresh predicate holds for π when the group in $\pi.\text{mode}$ is strong, \mathcal{A} neither queried $\text{Reveal}(\pi)$ nor $\text{Reveal}(\pi')$ for a matching session π' , and $\pi.\text{eid}_{\bar{r}}$ is honest.

Identity protection and *deniability* are other orthogonal security requirements of key-exchange protocols. We do not formally capture them in this work, but note that many design decisions in real-world key-exchange protocols are motivated by user privacy in addition to security.

2.4 Instantiating our model for SIGMA-N

Consider the SIGMA-N protocol of Fig. 1. The configurations should include sufficient detail to determine the negotiated mode. We thus include the acceptable *groups* and a function *PK* from identities to peer public keys. The latter would normally be implemented by looking up the public key of the peer in a certificate store. We thus have variables

$$\begin{aligned} \text{cfg} &\triangleq \begin{cases} (I, A, pk_A, PK, groups) & \text{for initiator } I \\ (R, B, pk_B, PK, groups) & \text{for responder } R \end{cases} \\ \text{uid} &\triangleq (g^x, g^y) \\ \text{mode} &\triangleq (G_i, pk_A, pk_B). \end{aligned}$$

SIGMA-N inherits the security properties of SIGMA [32]. It thus provides authentication and key-indistinguishability for *Auth* and *Fresh* predicates that hold only for modes with strong groups and honest keys. Since g^x and g^y are signed, it also provides partnering security, even if the group is weak.

3 Defining Downgrade Resilience

Downgrade resilience is motivated by protocols such as SIGMA-N that despite satisfying all of the definitions above remain vulnerable to practical attacks. We model the desired outcome of negotiation using a function *Nego* that maps two configurations with opposite roles to the protocol mode negotiated (if any) in the absence of active adversaries. Formally, if a session π talking to a session π' completes, it must be the case that $\pi.\text{mode} = \text{Nego}_r(\pi.\text{cfg}, \pi'.\text{cfg})$, where Nego_r is an abbreviation defined by case:

$$\text{Nego}_r(\text{cfg}_r, \text{cfg}_{\bar{r}}) \triangleq \begin{cases} \text{Nego}(\text{cfg}_r, \text{cfg}_{\bar{r}}) & \text{when } r = I \\ \text{Nego}(\text{cfg}_{\bar{r}}, \text{cfg}_r) & \text{when } r = R. \end{cases}$$

Definition 5 (Negotiation correctness). *The protocol negotiation is correct if, whenever a session π with role r and configuration cfg_r completes, there exists a peer configuration $\text{cfg}_{\bar{r}}$ such that $\pi.\text{mode} = \text{Nego}_r(\text{cfg}_r, \text{cfg}_{\bar{r}})$.*

Although we expect this basic property to hold unconditionally, many implementation errors may break it. For instance, the FREAK attack stems from TLS clients that do not offer export ciphersuites but still accept export-grade RSA keys. An implementation of SIGMA-N in which an initiator accepts groups it did not propose would also fail to satisfy negotiation correctness.

Downgrade security complements negotiation correctness. Informally, a protocol is *downgrade secure* when two sessions of opposite roles with the same unique identifier *uid* always negotiate the mode *prescribed* by their configurations. Hence, downgrade security concerns situations in which *one participant can save the other*, even if the latter supports broken cryptography. However, we have to assume that at least some of the mechanisms of the protocol (e.g., its signature modes) are strong enough. Conversely, if both participants enable (among others) a mode that is entirely insecure, then there is no cryptographically sound way to prevent an attacker from downgrading their connection.

Our definition is parameterized by a downgrade-protection predicate *DP* on pairs of configurations.

$DP(cfg_r, cfg_{\bar{r}})$ indicates the pairs of configurations from which we expect downgrade protection; it is not necessarily symmetric. By convention, cfg_r is the local configuration, $cfg_{\bar{r}}$ is the peer configuration, and when $DP(cfg_r, cfg_{\bar{r}})$ holds, we expect that the local session is protected.

Definition 6 (Downgrade security). *A session π is downgraded when $\pi.complete = \mathbf{true}$ and there is a partnered session π' such that $DP(\pi.cfg, \pi'.cfg)$, and $\pi.mode \neq \mathit{Nego}_{\pi.role}(\pi.cfg, \pi'.cfg)$.*

The advantage $\mathbf{Adv}_{\Pi, DP, X}^{\text{downgrade}}(\mathcal{A})$ of \mathcal{A} against downgrade security with agreement on X is the probability that, when \mathcal{A} terminates after interacting with Π , there exists a session π that either is downgraded or disagrees with a partnered session π' on X . We write $\mathbf{Adv}_{\Pi, DP}^{\text{downgrade}}(\mathcal{A})$ when $X = \{\}$.

First observe that protocols that support a single mode are trivially downgrade secure with agreement on *mode*. Note also that only partnered sessions get downgrade protection guarantees, so our definition is meaningful only for protocols for which partnering security holds. For role r , if $DP(cfg_r, cfg_{\bar{r}})$ holds for any peer configuration $cfg_{\bar{r}}$, we should also have $PS(cfg_r)$; we write this concisely as $DP \subseteq_r PS$, and observe that this property holds in our case studies. Downgrade security is a complementary, but intuitively stronger property.

Agreement on *mode* (or some of its parts) is desirable but not essential for downgrade protection. Conversely, for configurations cfg_r and $cfg_{\bar{r}}$ for which both $DP(cfg_r, cfg_{\bar{r}})$ and $DP(cfg_{\bar{r}}, cfg_r)$ hold, we do have downgrade protection with agreement on *mode*.

The DP predicate for downgrade protection plays a role similar to *Auth* for authentication, but it depends only on static configurations and on the honesty of long-term credentials. This reflects that downgrade protection should depend only on the inputs to the negotiation, and not the negotiation itself, which may be under the influence of an adversary.

Our formal configurations are session-specific, and do not necessarily coincide with concrete configuration in real-world protocol deployments. In particular, each configuration contains credentials only for the intended peer (e.g. cached certificates, key fingerprints). As an example, our configurations for TLS include the authentication settings of the session: the client's configuration expresses its intent to communicate with a particular server, who may nevertheless support multiple negotiable certificates, e.g., those negotiated using the server name indication extension [15].

Ideally, $DP(\pi.cfg, \cdot)$ would hold regardless of the second configuration. Anticipating on our results, this is the case for SSH, where DP is defined as follows: the configuration of π must require authentication of its peer, all peer keys accepted by π must be honest, and all signature algorithms must be (agile) strong. However, this is not the case e.g. for TLS 1.2 clients, which do not get downgrade protection with servers that support weak Diffie-Hellman groups.

Generalizing downgrade security For simplicity, our definitions above assume there is at most one correct mode reachable from configurations $\pi.cfg$ and $\pi'.cfg$, and consider adversaries that lead sessions to pick any other mode.

More generally, we may let *Nego* return a set of equally-acceptable modes and tolerate attacks that influence which of these modes is picked by partnered sessions. In that case, of course, we would still insist that the two sessions agree on the negotiated mode.

As an example, we may interpret configurations as sets of acceptable modes, and let *Nego* compute their intersection. As long as the protocol is negotiation-correct and guarantees agreement on *mode*, it would be downgrade secure for this generalized definition.

3.1 Downgrade resilience of SIGMA-N

Recall that SIGMA-N configurations are tuples $(r, ID, pk_{ID}, PK, groups)$ where PK is a function mapping identities to public keys. The negotiation function describes the correct mode upon completion. Given a function $nego$ that selects the preferred common group, $Nego$ is defined as

$$Nego(cfg_I, cfg_R) \triangleq \begin{cases} (nego(cfg_I.groups, cfg_R.groups), cfg_I.pk_A, cfg_R.pk_B) & \text{if } cfg_I.pk_A = cfg_R.PK(A) \\ \perp & \text{otherwise .} \end{cases}$$

Note that SIGMA-N does not guarantee agreement on *mode* because the responder R completes first without receiving any confirmation from I . Thus, $Nego$ is defined even when $cfg_R.pk_B \neq cfg_I.PK(B)$.

For such a $Nego$ function, DP can hold only for pairs of configurations with at most one group in common. However, observe that if DP guarantees the honesty of peer public keys, a participant only accepts groups whitelisted by its partnered peer. Thus, SIGMA-N is downgrade secure if we generalize $Nego$ to tolerate the negotiation of any mode that uses a group in the intersection of the groups whitelisted by both participants, i.e.

$$Nego(cfg_I, cfg_R) \triangleq \begin{cases} \{(G, cfg_I.pk_A, cfg_R.pk_B) \mid G \in cfg_I.groups \cap cfg_R.groups\} & \text{if } cfg_I.pk_A = cfg_R.PK(A) \\ \emptyset & \text{otherwise .} \end{cases}$$

3.2 Downgrade resilience and multi-mode security

Protocol analysts often consider protocols restricted to specific modes and configurations. For instance it is common practice to analyze individual protocol modes in isolation. Similarly we can restrict the initial configurations of a protocol to those that provide downgrade protection. Consider sets of configurations C_I and C_R picked by initiators and responders respectively. We consider restricted protocols in which sessions abort whenever they are initialized with a configuration outside of the set $C_I \cup C_R$.

Definition 7 (Protected configurations). *Let DP be a downgrade protection predicate. A pair of sets of configurations $(C_r, C_{\bar{r}})$ gives downgrade protection to role r if $C_r \times C_{\bar{r}} \subseteq DP$.*

The following theorem expresses that when downgrade security holds, only the security of modes that can be negotiated in the absence of an adversary matters. That is, if peers support insecure modes, but with such a low priority that they never negotiate them on their own, then these modes do not affect security in the presence of an adversary.

Theorem 1 (Downgrade resilience and multi-mode security). *Let Π be a protocol, $(C_r, C_{\bar{r}})$ sets of configurations, DP a downgrade protection predicate, and $\mathcal{N} = \{Nego_r(cfg_r, cfg_{\bar{r}}) \mid cfg_r, cfg_{\bar{r}} \in C_r \times C_{\bar{r}}\}$ the modes negotiable without adversary influence. If $DP \subseteq_r PS$ and*

- $(C_r, C_{\bar{r}})$ gives downgrade protection to r ,
- Π is multi-mode authentication secure for $Auth, X$,
- Π is partnering secure for PS , and
- Π is downgrade secure for DP ,

then the protocol Π restricted to configurations in $C_r \cup C_{\bar{r}}$ is multi-mode authentication secure for a more lax Auth' predicate that deems all modes outside of \mathcal{N} as “good”, i.e. $\text{Auth}'(m, \text{role}) \triangleq \text{Auth}(m, \text{role}) \vee (m \notin \mathcal{N} \wedge \text{role} = r)$. Concretely, given an adversary \mathcal{A} against authentication for Auth', X , we have

$$\mathbf{Adv}_{\Pi', \text{Auth}', X}^{\text{auth}}(\mathcal{A}) \leq \mathbf{Adv}_{\Pi, PS}^{\text{partnering}}(\mathcal{A}) + \mathbf{Adv}_{\Pi, DP}^{\text{downgrade}}(\mathcal{A}) + \mathbf{Adv}_{\Pi, \text{Auth}, X}^{\text{auth}}(\mathcal{A}),$$

where Π' is Π restricted to configurations $C_r \cup C_{\bar{r}}$.

Proof sketch. Consider the multi-mode authentication experiment \mathbf{G}_0 for Π' . Let S hold when at some point through \mathbf{G}_0 a session π completes maliciously on X and $\text{Auth}(\pi.\text{mode}, \pi.\text{role})$ holds (i.e., \mathcal{A} succeeds in breaking authentication iff S holds at the end of \mathbf{G}_0).

Game \mathbf{G}_1 behaves as \mathbf{G}_0 except it aborts just before a session π of role r would complete without being partnered. Because of the restriction in Π' and the hypothesis that $DP \subseteq_r PS$, it must be the case that $PS(\pi.\text{cfg})$. Thus, any time \mathbf{G}_1 aborts, \mathcal{A} succeeds in breaking the partnering security of Π' , and thus that of Π . Hence, the difference in the probability of S between \mathbf{G}_0 and \mathbf{G}_1 is at most $\mathbf{Adv}_{\Pi, PS}^{\text{partnering}}(\mathcal{A})$.

Game \mathbf{G}_2 behaves as \mathbf{G}_1 except it aborts just before a session π of role r would complete and there is a partnered session π' such that $\pi.\text{mode} \neq \text{Nego}_r(\pi.\text{cfg}, \pi'.\text{cfg})$. The difference in the probability of S between \mathbf{G}_1 and \mathbf{G}_2 is at most $\mathbf{Adv}_{\Pi, DP}^{\text{downgrade}}(\mathcal{A})$ since any time \mathbf{G}_2 aborts but \mathbf{G}_1 does not, \mathcal{A} succeeds in breaking the downgrade security of Π' , and thus that of Π .

By definition of \mathcal{N} , \mathbf{G}_2 never completes with a session of role r assigning a mode outside of \mathcal{N} . Consequently, the probability of S in this game is at most

$$\mathbf{Adv}_{\Pi', \text{Auth}', X}^{\text{auth}}(\mathcal{A}) \leq \mathbf{Adv}_{\Pi, \text{Auth}, X}^{\text{auth}}(\mathcal{A}). \quad \square$$

Interestingly, partnering security is similar to the aliveness requirement in some (single-mode) security definitions which Krawczyk [32] does not consider as fundamental for key-exchange security. Our second game transformation however only works if a partnered session π' with the same *wid* exists. Otherwise an abort in \mathbf{G}_2 cannot be translated into a downgrade security attack.

For key-indistinguishability and a freshness predicate $\text{Fresh}(\pi)$ that requires for matching π' with $\pi'.\text{mode} \neq \perp$ that $\pi.\text{mode} = \pi'.\text{mode} = \text{Nego}_r(\pi.\text{cfg}, \pi'.\text{cfg})$, we have an analogous theorem for

$$\text{Fresh}'(\pi) \triangleq \text{Fresh}(\pi) \vee (\pi.\text{mode} \notin \mathcal{N} \wedge \pi.\text{role} = r).$$

In the context of TLS, Bhargavan et al. [13] observe that this is a sufficient condition for the security of session keys that are released before the handshake completes.

3.3 Downgrade secure sub-protocols

We are interested in minimal core sub-protocols that guarantee downgrade security. To justify our use of sub-protocols in further sections as a sound abstraction of the full protocol we use simulation. Our sub-protocols can take additional input as part of **Init** and **Send** queries to allow for an accurate simulation of the full protocol. This is akin to the sub-protocols of Bergsma et al. [9] which allow for additional signing oracles (restricted to not break security of the sub-protocol).

For simplicity, the following definition leaves out details about handling of long-term keys and corruption models. When filling in the details for a particular setting, we require the simulation to be accurate with respect to e.g. corruption, so that it issues exactly the same corruption queries as in the full protocol.

Definition 8 (Sub-protocol). *A protocol $\tilde{\Pi}$ is a sub-protocol of Π for X if we have an efficient simulator \mathcal{S} with only oracle access to $\tilde{\Pi}$ and exposing the same oracles as Π , such that $\mathcal{S} \circ \tilde{\Pi}$ is indistinguishable from Π for an information-theoretic distinguisher with access to all session variables in X .*

Formally, we model a protocol (and a simulator) as a collection of oracles sharing state, each oracle being a probabilistic algorithm. We model access to session variables using oracles that just return the value of the corresponding variable. The composition $\mathcal{S} \circ \tilde{\Pi}$ of a simulator \mathcal{S} and a sub-protocol $\tilde{\Pi}$ is well-defined when $\tilde{\Pi}$ includes all algorithms called by the oracles of \mathcal{S} . The composition itself is a new collection of algorithms, one for each oracle of \mathcal{S} . Operationally, the oracles of $\mathcal{S} \circ \tilde{\Pi}$ behave as the algorithmic composition of the oracles of \mathcal{S} and $\tilde{\Pi}$. Similarly, we model an adversary \mathcal{A} as a single probabilistic algorithm with access to oracles, and the composition $\mathcal{A} \circ \mathcal{S}$ (resp. $\mathcal{A} \circ \Pi$) behaves as the algorithmic composition of this algorithm with the oracles of \mathcal{S} (resp. Π).

As the next theorem shows, simulation allows to lift security properties satisfied by a sub-protocol to the full protocol.

Theorem 2 (Downgrade security lifting). *Let $\tilde{\Pi}$ be a sub-protocol of a protocol Π for session variables $\{cfg, uid, mode, key, complete\} \cup X$, and DP a downgrade protection predicate. Let \mathcal{S} be a simulator for $\tilde{\Pi}$ as in Definition 8. Then, for any adversary \mathcal{A} against the DP -downgrade security of Π with agreement on X , $\mathcal{A} \circ \mathcal{S}$ is an adversary against the downgrade security of $\tilde{\Pi}$ with agreement on X , and*

$$\mathbf{Adv}_{\tilde{\Pi}, DP, X}^{\text{downgrade}}(\mathcal{A} \circ \mathcal{S}) = \mathbf{Adv}_{\Pi, DP, X}^{\text{downgrade}}(\mathcal{A}) .$$

Proof sketch. If \mathcal{A} is successful when interacting with Π through the protocol oracles, then during the downgrade security experiment there must be a session π partnered with a session π' such that $DP(\pi.cfg, \pi'.cfg)$ holds and either $\pi.mode \neq \text{Nego}_{\pi.role}(\pi.cfg, \pi'.cfg)$ or π and π' disagree on X . Let \mathbf{E} denote this event. Note that the probability of \mathbf{E} in the experiment $\mathcal{A} \circ \Pi$ is exactly $\mathbf{Adv}_{\Pi, DP, X}^{\text{downgrade}}(\mathcal{A})$.

Now, since the simulation $\mathcal{S} \circ \tilde{\Pi}$ is accurate with respect to all variables this event depends on, and $\mathcal{S} \circ \tilde{\Pi}$ is indistinguishable from Π for \mathcal{A} , the probability of \mathbf{E} occurring in the experiment $\mathcal{A} \circ (\mathcal{S} \circ \tilde{\Pi})$ is the same as in the experiment $\mathcal{A} \circ \Pi$.

Because the composition operator \circ is such that $\mathcal{A} \circ (\mathcal{S} \circ \tilde{\Pi}) = (\mathcal{A} \circ \mathcal{S}) \circ \tilde{\Pi}$, we conclude by construing the composition of \mathcal{A} and \mathcal{S} as an adversary against the downgrade security of $\tilde{\Pi}$ with agreement on X . \square

An analogous theorem holds for partnering security.

3.4 Downgrade security by whitelisting

Consider a protocol that is negotiation correct and guarantees multi-mode authentication with agreement on all variables that influence the computation of $mode$, then we get downgrade protection for

$$DP(cfg, \cdot) \triangleq \forall cfg'. \text{Auth}(\text{Nego}_{cfg.role}(cfg, cfg'), cfg.role) .$$

That is, all negotiable modes from downgrade secure configurations must provide authentication security. This generalizes the *Negotiation-authentication* theorem of [23].

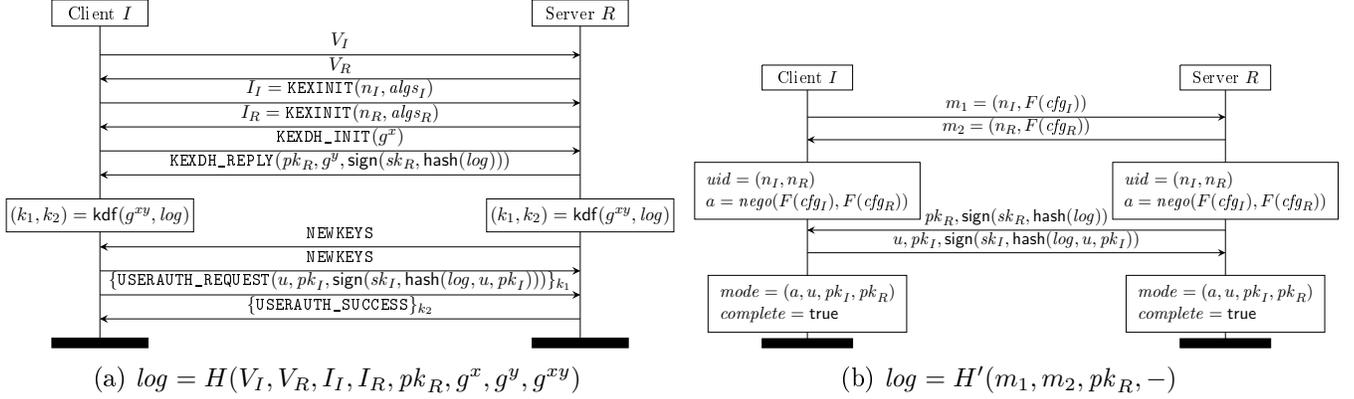


Figure 3: SSHv2 mutually-authenticated key exchange: (a) full protocol and (b) sub-protocol *SSH-sub*.

4 Secure SHell

Figure 3a models a run of the SSHv2 [46] protocol with a client that authenticates using the `publickey` method [45]. We analyze the downgrade security of this protocol using the sub-protocol shown on Figure 3b. The functions H, H' in these figures stand for the composition of a fixed injective formatting function and a negotiated hash function. Note that there are potential downgrade attacks in SSHv2 from `publickey` authentication to other mechanisms like `password` but the protocol we consider does not model the negotiation of the authentication mechanism. We stress that our analysis only applies assuming servers are configured to require public key authentication.

We next describe the notation used in Figure 3 and the way the protocol works. We then prove that *SSH-sub* is indeed a sub-protocol of the full protocol for the variables they have in common.

Client and server configurations include lists $algs$ of key exchange, server signature, encryption and MAC algorithms ordered by preference. We let $F(cfg) = cfg.algs$. Each party computes the negotiated ciphersuite independently, following the rules in the protocol specification [46, Sect. 7.1], which we encode in a *nego* function. Roughly, these rules dictate that the first algorithm for each category in cfg_I that is also in cfg_R be selected. Each session locally assigns $nego(F(cfg_I), F(cfg_R))$ to a . In addition, a client configuration cfg_I includes a user name and a service name u , a function PK_I mapping a pair (a, u) to a public key, and a function PKs_R mapping a value a to a set of acceptable server public keys. Conversely, a server configuration cfg_R includes a function PK_R mapping a value a to a public key, and a function PKs_I mapping a pair (a, u) to a set of acceptable client public keys. For instance, in OpenSSH the keys $cfg_I.PKs_R$ of acceptable server public keys are taken from the clients `known_hosts` file, whereas the keys $cfg_R.PKs_I$ of acceptable client public keys are taken from the `.ssh/authorized_keys` file in the home directory of the user on the server.

In terms of the template in Section 2, the sub-protocol uses the following session variables:

$$\begin{aligned}
 cfg &\triangleq \begin{cases} (I, algs, u, PK_I, PKs_R) & \text{for } I \\ (R, algs, PK_R, PKs_I) & \text{for } R \end{cases} \\
 uid &\triangleq (n_I, n_R) \\
 mode &\triangleq (a, u, pk_I, pk_R).
 \end{aligned}$$

Client and server exchange nonces and their algorithmic preferences $F(cfg_I), F(cfg_R)$. The server then

selects a compatible signature key pair (pk_R, sk_R) and signs a hash log that includes the first two exchanged messages. When receiving this message, the client checks that pk_R is an acceptable server key in its local configuration, computes log locally and verifies the server signature. If the signature verifies, it selects a key pair (pk_I, sk_I) in its configuration for authenticating and sends back to the server a signature over log , u , and pk_I . When receiving this message, the server checks that pk_I is an acceptable client key in $cfg_R.PKs_I(a, u)$. Each party completes the session upon successfully verifying the peer signature, otherwise aborts. Formally, a client aborts if $pk_R \notin cfg_I.PKs_R(a)$; otherwise it assigns the following value to $mode$:

$$(nego(F(cfg_I), F(cfg_R)), cfg_I.u, cfg_I.PK_I(a, u), pk_R)$$

The server’s behavior is specified analogously.

We augment the **Send** oracles of each a session in the sub-protocol with extra parameters that allow to fill in the blank $(-)$ used to compute log . This allows a simulator to compute signatures on the same values as the full protocol, as needed to consistently answer **Send** queries. Consequently, we allow an adversary against the downgrade security of the sub-protocol to fill in $-$ parameters arbitrarily.

We complete our security model with oracles $pk \leftarrow \text{KeyGen}$ for key generation, $sk \leftarrow \text{Corrupt}(pk)$ for adaptive corruption, and $\text{Coerce}(pk)$ for adversarial key registration. A public key pk is honest if it was generated by a query to oracle **KeyGen** but not corrupted by a **Corrupt** query.

Theorem 3 (Simulation). *SSH-sub (Fig. 3b) is a sub-protocol of SSH (Fig. 3a) for their common variables.*

Proof sketch. The sub-protocol is oblivious of the Diffie-Hellman exchange in the full protocol, so the simulator generates fresh Diffie-Hellman shares of his own for each session. When needed, the simulator \mathcal{S} forwards queries to *SSH-sub* after applying message parsing and formatting functions. To simulate signatures of honest sessions \mathcal{S} uses the Diffie-Hellman shares it has computed and the messages it has received to fill in the value of the extra parameter $-$ of oracles of the sub-protocol. Note that the adversary knows the secret exponents of an instance’s Diffie-Hellman shares and so it can always compute the encryption keys k_1, k_2 needed to simulate the last two messages of the full protocol. \square

4.1 SSHv2 is partnering and downgrade secure

A remarkable property of the downgrade protection sub-protocol of mutually-authenticated SSHv2 is that, because both client and server sign (a hash of) the inputs to the *nego* function, downgrade protection security relies only on the honesty of the signature keys, the collision resistance of the hash algorithm, and the strength of the signature algorithms. Notably, it does not rely on the key exchange algorithm being strong or contributive, not even on it providing high entropy inputs to H . This means that we can prove this protocol secure for a predicate DP that only constrains the signature and hash algorithms of cfg_r , and requires honesty of peer public keys in $cfg_r.PKs_{\bar{r}}$, but has no requirements on $cfg_{\bar{r}}$.

We prove the partnering and downgrade security of SSHv2 with **publickey** client authentication under the agile security assumptions on hash functions and signatures that we present next.

4.1.1 Agile hash functions and signatures

As protocol participants may negotiate different hash functions we need to capture collisions across hash functions.

Definition 9 (Agile collision resistance). *Let h^* be a hash function, and \mathcal{H} a set of hash functions. Consider the game:*

- $h, v, v' \leftarrow \mathcal{A}()$
- Return $h^*(v) = h(v') \wedge v \neq v'$

The collision resistance advantage of \mathcal{A} , $\text{Adv}_{h^*, \mathcal{P}}^{\text{CR}}(\mathcal{A})$ is the probability that the game returns true.

If the ranges of hash functions are disjoint, agile collision resistance reduces to ordinary collision resistance. Bhargavan et al. [13] also define existential unforgeability under chosen-message attacks (EUF-CMA) for agile hash-then-sign signatures. We here consider such signatures as primitives although typical constructions can be proved secure in the random oracle model.

Definition 10 (Agile EUF-CMA security). Consider an agile signature scheme $s = (\text{keygen}, \text{sign}, \text{verify})$. Let p^* be an agility parameter, and \mathcal{P} a set of parameters. Consider the forgery game:

- Let $pk, sk \leftarrow \text{keygen}()$
- Set $M := \{\}$ and run $m, \sigma \leftarrow \mathcal{A}^{\text{Sign}}(pk)$
- Return $m \notin M \wedge \text{verify}(pk, p^*, m, \sigma)$

where $\text{Sign}(p, m)$ returns \perp if $p \notin \mathcal{P}$ and otherwise sets $M := M \cup \{m\}$ before returning $\text{sign}(sk, p, m)$.

The advantage $\text{Adv}_{s, p^*, \mathcal{P}}^{\text{EUF-CMA}}(\mathcal{A})$ of \mathcal{A} in forging a signature for s is the probability that the forgery game returns true.

Since we proved that *SSH-sub* soundly abstracts negotiation in the full protocol, any downgrade attack on the full protocol can be turned into a downgrade attack on the sub-protocol. By virtue of Theorem 2 it suffices to prove that *SSH-sub* is downgrade secure. The same reasoning applies to partnering security.

4.1.2 Partnering security

To state and prove partnering security, we use the following notations.

Let $\text{Nego}(cfg_I, cfg_R) \triangleq (a, cfg_I.u, cfg_I.PK_I(a, u), pk_R)$, where $a = \text{nego}(F(cfg_I), F(cfg_R))$ if $pk_R = cfg_R.PK_R(a)$, and \perp otherwise.

Let $\mathcal{M}^* \triangleq \{\text{Nego}_{cfg.role}(cfg, cfg') \mid PS(cfg)\}$ be the modes negotiated between any pair of configurations for which the first guarantees partnering security and \mathcal{M} be the set of all supported modes.

Let $\mathcal{P}_s \triangleq \{p \mid s, p = \text{mode.sig} \wedge \text{mode} \in \mathcal{M}\}$ be the agility parameters for the signature scheme s of a peer, \mathcal{H} be the set of all hash algorithms supported, and $\mathcal{H}^* \triangleq \{\text{mode.hash} \mid \text{mode} \in \mathcal{M}^*\}$ be the hash algorithms used by partnering secure modes.

Theorem 4 (Partnering security of *SSH-sub*). Let PS be such that $PS(cfg)$ implies that all public keys in the range of $cfg.PK_{s_{cfg.role}}$ are honest. Given an adversary \mathcal{A} against the partnering security of *SSH-sub*, we construct adversaries $\mathcal{B}_{s,p,i}$ and \mathcal{B}_h running in about the same time as \mathcal{A} such that $\text{Adv}_{SSH-sub, PS}^{\text{partnering}}(\mathcal{A})$ is at most

$$\sum_{h \in \mathcal{H}^*} \text{Adv}_{h, \mathcal{H}}^{\text{CR}}(\mathcal{B}_h) + \sum_{(s,p) \in \text{sig}(\mathcal{M}^*)} \sum_{i=1}^{n_s} \text{Adv}_{s,p, \mathcal{P}_s}^{\text{EUF-CMA}}(\mathcal{B}_{s,p,i}),$$

where n_s is the number of keys generated for scheme s .

Proof sketch. Let G_0 be the original partnering security experiment. Let S hold whenever at some point during the experiment, a session π for which $PS(\pi.cfg)$ holds completes without being partnered. Let G_1 behave like G_0 , except it aborts whenever two sessions hash messages (m_1, m_2) with different (n_I, n_R) to the same value. Using \mathcal{A} , we can construct for each $h \in \mathcal{H}^*$ an adversary \mathcal{B}_h such that $|\Pr_{G_0}[S] - \Pr_{G_1}[S]| \leq \sum_{h \in \mathcal{H}^*} \mathbf{Adv}_{h, \mathcal{H}}^{CR}(\mathcal{B}_h)$.

Game G_2 behaves as G_1 , except that sessions with honest peer keys and a peer signature scheme $(s, p) \in \overline{sig}(\mathcal{M}^*)$ abort without completing whenever they verify a signature on log that was not signed by another session.

To bound the difference in the probability of S between these games we use a lemma of Bhargavan et al. [14, Lemma 4]. It shows via a hybrid argument how to construct adversaries $\mathcal{B}_{s,p,i}$ that bound the probability of a forgery for an agile hash-and-sign scheme $(s, p) \in \overline{sig}(\mathcal{M}^*)$ in an agile signature library that generates n_s honest keys for signing scheme s , thus

$$|\Pr_{G_1}[S] - \Pr_{G_2}[S]| \leq \sum_{(s,p) \in \overline{sig}(\mathcal{M}^*)} \sum_{i=1}^{n_s} \mathbf{Adv}_{s,p,\mathcal{P}_s}^{EUF-CMA}(\mathcal{B}_{s,p,i}).$$

We conclude by confirming that all sessions hash (n_I, n_R) into log and verify a signature on log before completion. For those sessions for which $PS(cfg)$ holds, $pk_{cfg.role} \in cfg.PK_{s_{cfg.role}}$ is honest, and we have that another session in possession of the signing key must have assigned the same $uid = (n_I, n_R)$. Finally, since the signatures computed by initiators and responders cannot be confounded, the other session must have the opposite role. \square

4.1.3 Downgrade security

To prove downgrade security, we define $Nego$, \mathcal{M} , \mathcal{P}_s , and \mathcal{H} as before, but re-define \mathcal{M}^* , \mathcal{H}^* to use DP instead of PS , i.e. $\mathcal{M}^* \triangleq \{Nego_{cfg.role}(cfg, cfg') \mid DS(cfg, cfg')\}$.

Theorem 5 (Downgrade security of *SSH* sub-protocol). *Let DP be such that $DP(cfg, \cdot)$ implies that all public keys in the range of $cfg.PK_{s_{cfg.role}}$ are honest. Given an adversary \mathcal{A} against the downgrade security of the sub-protocol, we construct adversaries $\mathcal{B}_{s,p,i}$ and \mathcal{B}_h running in about the same time as \mathcal{A} such that $\mathbf{Adv}_{SSH-sub, DP}^{downgrade}(\mathcal{A})$ is at most*

$$\frac{n^2}{2^{|uid|/2}} + \sum_{h \in \mathcal{H}^*} \mathbf{Adv}_{h, \mathcal{H}}^{CR}(\mathcal{B}_h) + \sum_{(s,p) \in \overline{sig}(\mathcal{M}^*)} \sum_{i=1}^{n_s} \mathbf{Adv}_{s,p,\mathcal{P}_s}^{EUF-CMA}(\mathcal{B}_{s,p,i}),$$

where n is the number of sessions, n_s the number of keys generated for scheme s , and $|uid|$ the size of unique identifiers.

Proof sketch. Let S be the event that there exists a downgraded session π and let G_0 be the original downgrade security experiment.

Game G_1 aborts whenever two sessions of the same role assign the same uid . Let n be the total number of sessions. The length of the randomness in client and server nonces is $|uid|/2$ bits. The probability that n such random values give rise to a collision is approximately $n^2 2^{-|uid|/2-1}$. As G_0 and G_1 are equivalent up to collisions, we have $|\Pr_{G_0}[S] - \Pr_{G_1}[S]| \leq n^2 2^{-|uid|/2-1}$.

Game \mathbf{G}_2 behaves as \mathbf{G}_1 , except that it aborts whenever two sessions hash messages $(m_1, m_2, -, pk_R, -)$ to the same log .

Using \mathcal{A} , we can construct for each $h \in \mathcal{H}^*$ an adversary \mathcal{B}_h such that $|\Pr_{\mathbf{G}_0}[S] - \Pr_{\mathbf{G}_1}[S]| \leq \sum_{h \in \mathcal{H}^*} \mathbf{Adv}_{h, \mathcal{H}}^{\text{CR}}(\mathcal{B}_h)$.

Observe that in \mathbf{G}_2 , every pair of nonces $(n_I, n_R) = uid$ in the message pairs (m_1, m_2) is signed at most once per role.

Game \mathbf{G}_3 is the same as Game \mathbf{G}_2 , except that sessions with honest peer keys and $(s, p) \in \overline{sig}(\mathcal{M})$ abort without completing whenever they verify a signature on log or (log, u, pk_I) that was not signed by another session.

We again bound the difference in the success probabilities using the lemma of Bhargavan et al. [14, Lemma 4] and have

$$|\Pr_{\mathbf{G}_2}[S] - \Pr_{\mathbf{G}_3}[S]| \leq \sum_{(s,p) \in \overline{sig}(\mathcal{M}^*)} \sum_{i=1}^{n_s} \mathbf{Adv}_{s,p,\mathcal{P}_s}^{\text{EUF-CMA}}(\mathcal{B}_{s,p,i}).$$

Observe that, since the signatures computed by initiators and responders cannot be confounded, in Game \mathbf{G}_3 we never verify a message that was not signed by its unique matching peer.

We now consider initiators and responders separately.

For the initiator we have that message log was signed by its partner, which computed it from the responders key pk_R and messages m_1 and m_2 containing the contribution of the local and the remote configuration respectively. Note that the only missing component of $mode$ is u, pk_I which, conditioned on the responder completing are determined by the initiators configuration.

For the responder, message (log, u, pk_I) authenticates the two public keys and all inputs to $nego$.

We thus excluded the event S from ever occurring in \mathbf{G}_3 . \square

5 Internet Key Exchange

The Internet Key Exchange (IKE) protocol is the key exchange component of the IPsec suite of protocols. Two versions of the protocol are commonly deployed: IKEv1 [27] and IKEv2 [30]. Both variants are inspired by the SIGMA protocol [32] recalled in the introduction, and are believed to inherit its authentication and key-indistinguishability guarantees. Next, we study their downgrade protection sub-protocols.

5.1 IKEv1 does not prevent downgrades

We first consider the DHE-PSK modes of IKEv1, whose first three messages are depicted in Figure 4a. The protocol presumes that both parties can select the pre-shared key (psk) to use from the negotiated security association SA_R and identifiers ID_I and ID_R ; it then confirms that the two parties agree, using a MAC based on psk . The two parties also exchange Diffie-Hellman shares and use them to derive session keys and protect application data but, in ‘aggressive’ modes, their authentication and downgrade-protection relies solely on the pre-shared key.

The corresponding downgrade protection sub-protocol is depicted in Figure 4b. The initiator begins by extracting a list of supported security associations $[SA_1, \dots, SA_n]$ from its configuration, presumably ordered by preference, formats them (using the function F), and sends them along with a nonce (n_I) to the responder. Each security association specifies a Diffie-Hellman group (for the key exchange); an encryption scheme and a hash algorithm (for protecting messages); and a peer authentication method. The responder chooses one of these associations (SA_R), based on its own configuration, and responds with its own nonce.

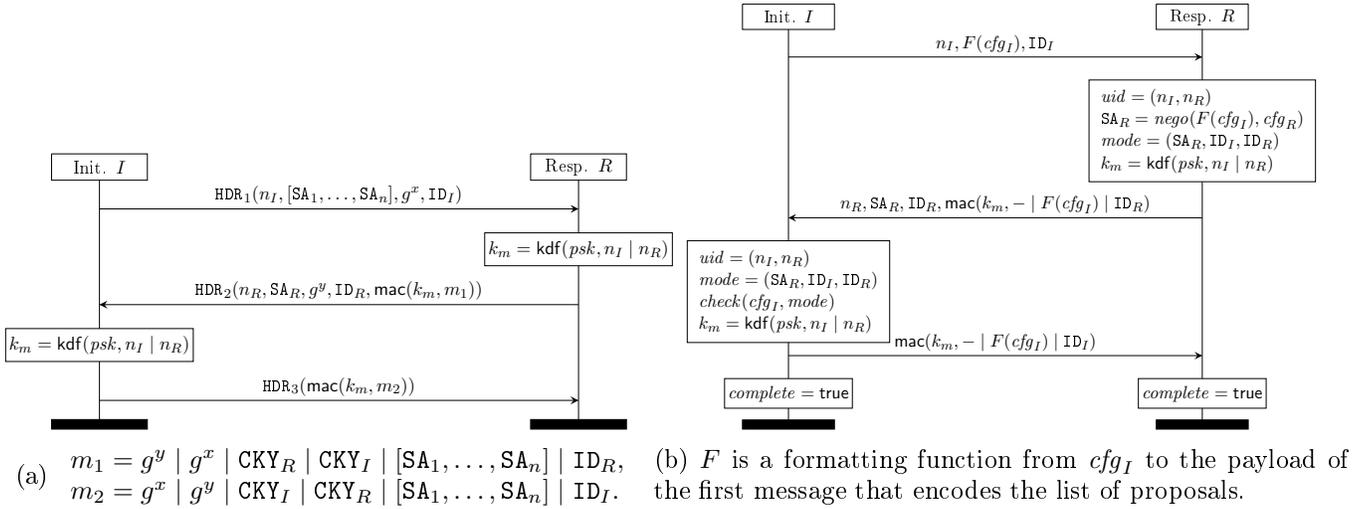


Figure 4: IKEv1 aggressive DHE-PSK protocol (a) first messages (b) downgrade protection sub-protocol.

The initiator checks that this choice is compatible with its proposals, which completes the negotiation. To authenticate one another, to provide key confirmation, and to prevent downgrades, the initiator and responder exchange MACs, optionally signed when using certificates for authentication. For simplicity, Figures 4a and 4b depict the use of just a pre-shared key for authentication. The MACs are computed with a key derived from the pre-shared key and the nonces, over some important parts of the protocol transcript: the key shares, the 8 byte ISAKMP cookies taken from the headers, the client’s offered security associations and the sender’s identity.

Surprisingly, the MAC does not cover the negotiated security association (SA_R), and this omission leads to a downgrade attack. A man-in-the-middle can simply modify the second message to replace the server’s chosen SA_R with a different SA'_R compatible with the initiator’s proposals. If this new SA'_R uses an encryption algorithm that the attacker can break (e.g. DES or NULL), then the attacker can break the confidentiality of the first messages sent by the initiator. (Similarly, the first MAC includes ID_R but not ID_I , so an attacker can modify ID_I in the first message, and yet the initiator will complete the sub-protocol without detecting the modification; this is less problematic in the full protocol because IKEv1 continues with a confirmation message from the responder.)

We instantiate our main definitions to IKEv1 to better understand this downgrade-protection failure and propose fixes. Clearly, the protocol offers no authentication guarantees unless the PSKs used by both parties are honest, so we always make that assumption in the following, which enables us to omit the choice of PSKs from the negotiation predicates. In IKEv1, the mac and kdf functions are negotiated as part of SA_R . They are effectively HMAC-MD5 or HMAC-SHA1. For simplicity, we also restrict our attention to clients and servers configured to use only HMAC-SHA1. (See §4 for an explicit handling of cryptographic agility.) We use the following notations for the sub-protocol:

- the goal is to agree on a mode $(\text{SA}_R, \text{ID}_I, \text{ID}_R)$;
- $cfg_I = (\text{ID}_I, [\text{SA}_1, \dots, \text{SA}_n])$.
- cfg_R includes ID_R and is otherwise unspecified; it would typically also include a list of SAs.

- F is a formatting function from cfg_I to the payload of the first message that encodes the list of proposals above.
- $nego$ is a partial function, used by the responder to map $F(cfg_I)$ and cfg_R to some SA_R .
- $check$ is used by the initiator to confirm that the mode is acceptable, checking for instance that SA_R matches one of the initiator's proposals $[SA_1, \dots, SA_n]$.
- $Nego(cfg_I, cfg_R)$, our specification for negotiation, is defined as $(nego(F(cfg_I), cfg_R), cfg_I.ID_I, cfg_R.ID_R)$ when $check$ succeeds, and is otherwise undefined.

The definitions of F , $nego$, and $check$ follow from the standard; their details are unimportant in our presentation.

We first prove partnering security (Definition 2) relying on the security of both kdf (modeled as a PRF keyed with psk) and mac (modeled as a MAC, relying e.g. on existential unforgeability under chosen-message attacks).

To simplify our presentation, we assume in our formal theorem statement a universe of configurations \mathcal{M}_{fix} with a fixed kdf and mac algorithms. As discussed above in practice this is HMAC-SHA1. Consequently, we can rely on standard pseudo-randomness and chosen message attack unforgeability assumptions for kdf and mac and thus simplify our theorem and proof. We hope this is useful for a reader accustomed to traditional non-agile provable security developments. This is still meaningful, as configurations may vary for instance on the choice of groups in the full protocol, moreover the proof acts as a template for a fully agile theorem analogous to Theorem 4.

We also assume that the protocol rejects runs in which $ID_I = ID_R$. This is referred to in the literature as the *self-communication* scenario, and in such settings there are well known reflection attacks on IKEv1 [37, 20]. Positive results in this setting would require an extension of our sub-protocol and assumptions about the DH groups or the ISAKMP cookies employed in the protocol.

Theorem 6 (Partnering security of *IKEv1* sub-protocol). *Let \mathcal{M}_{fix} be the universe of configurations used by protocol participants and let PS be such that all PSKs referenced by handles in cfg are honest. Given an adversary \mathcal{A} against the partnering security of *IKEv1*-sub, we construct adversaries \mathcal{B}_i and \mathcal{B}' running in about the same time as \mathcal{A} such that $\mathbf{Adv}_{IKEv1-sub, PS}^{partnering}(\mathcal{A})$ is at most*

$$\frac{n^2}{2^{|uid|/2}} + \sum_{i=1}^{n_p} \mathbf{Adv}^{PRF}(\mathcal{B}_i) + n \cdot \mathbf{Adv}^{EUF-CMA}(\mathcal{B}'),$$

where n is the number of sessions and n_p is the number different psks employed by sessions.

Proof sketch. The proof relies on the honesty of psk .

Let \mathbf{G}_0 be the original partnering security experiment. Let S hold whenever at some point during the experiment, a session π for which $PS(\pi.cfg)$ holds completes without being partnered. Game \mathbf{G}_1 aborts whenever two sessions of the same role assign the same uid .

Let n be the total number of sessions. The length of the client and server nonces is $|uid|/2$ bits. The probability that n such random values give rise to a collision is approximately $n^2 2^{-|uid|/2-1}$. As \mathbf{G}_0 and \mathbf{G}_1 are equivalent up to collisions, we have $|\Pr_{\mathbf{G}_0}[S] - \Pr_{\mathbf{G}_1}[S]| \leq n^2 2^{-|uid|/2-1}$.

Game \mathbf{G}_2 behaves like \mathbf{G}_1 , except that it keeps a map from pairs of nonces to MAC keys. Instead of defining $k_m = kdf(psk, n_I | n_R)$, game \mathbf{G}_2 looks up n_I, n_R in this map to retrieve the corresponding key

k_m when it has been defined. Otherwise, it samples a fresh random MAC key k_m and stores the entry (n_I, n_R, k_m) in the map. We bound $|\Pr_{\mathbf{G}_1}[S] - \Pr_{\mathbf{G}_2}[S]| \leq \sum_{i=1}^{n_p} \mathbf{Adv}^{\text{PRF}}(\mathcal{B}_i)$: for each of the n_p pre-shared secrets psk_k we define a hybrid, and bound the probability of \mathcal{A} distinguishing between adjacent hybrids by $\mathbf{Adv}^{\text{PRF}}(\mathcal{B}_i)$.

The keys k_m are now random and when event S holds for a session with $uid = (n_I, n_R)$, (n_I, n_R, k_m) was not accessed by any session to generate MACs. Moreover, because $ID_I \neq ID_R$ MACs cannot be reflected. Consequently, any successful adversary in \mathbf{G}_1 can be turned into a reduction \mathcal{B}' that guess the PSK handle used by the session for which S holds and returns the MAC received by that session as a forgery.

Therefore $\Pr_{\mathbf{G}_2}[S] \leq n \cdot \mathbf{Adv}^{\text{EUF-CMA}}(\mathcal{B}')$. \square

On the other hand, the protocol offers provable downgrade protection only for very restrictive configurations. For example, relying on the unambiguous formatting of ID_I and ID_R in the MACed payloads, we have downgrade protection when

1. the client (or the server) uses each PSK only for a fixed SA_R, ID_I, ID_R ; or
2. the client proposes only one SA at a time and checks that the server echoes this proposal in SA_R , and moreover SA determines ID_I .

Our analysis of the IKEv1 downgrade-protection sub-protocol suggests an obvious fix: include the mode (SA_R, ID_I, ID_R) in both MACs. We thus have

$$\begin{aligned} m_1 &= g^y \mid g^x \mid \text{CKY}_R \mid \text{CKY}_I \mid [\text{SA}_1, \dots, \text{SA}_n] \mid SA_R, ID_R, ID_I \\ m_2 &= g^x \mid g^y \mid \text{CKY}_I \mid \text{CKY}_R \mid [\text{SA}_1, \dots, \text{SA}_n] \mid SA_R, ID_I, ID_R. \end{aligned}$$

We then obtain downgrade protection under the same conditions as for partnering: that PSKs be honest and both \mathbf{kdf} and \mathbf{mac} be secure.

Theorem 7 (Downgrade security of *IKEv1* sub-protocol). *Let \mathcal{M}_{fix} be the universe of configurations used by protocol participants and let DP be such that $DP(\text{cfg}, \cdot)$ implies that all pre-shared keys referenced by handles in cfg are honest. Given an adversary \mathcal{A} against the downgrade security of the sub-protocol, we construct adversaries \mathcal{B}_i and \mathcal{B}' running in about the same time as \mathcal{A} such that $\mathbf{Adv}_{\text{IKEv1-sub}, DP}^{\text{downgrade}}(\mathcal{A})$ is at most*

$$\frac{n^2}{2^{|uid|/2}} + \sum_{i=1}^{n_p} \mathbf{Adv}^{\text{PRF}}(\mathcal{B}_i) + n \cdot \mathbf{Adv}^{\text{EUF-CMA}}(\mathcal{B}'),$$

where n is the number of sessions and n_p is the number of different pre-shared keys employed by sessions.

Proof sketch. Let \mathbf{G}_0 be the original downgrade security experiment and let S be the event that there exists a downgraded session.

Game \mathbf{G}_1 aborts whenever two sessions of the same role assign the same uid . Let n be the total number of sessions. The length of client and server nonces is $|uid|/2$ bits. The probability that n such random values give rise to a collision is approximately $n^2 2^{-|uid|/2-1}$. As \mathbf{G}_0 and \mathbf{G}_1 are equivalent up to collisions, we have $|\Pr_{\mathbf{G}_0}[S] - \Pr_{\mathbf{G}_1}[S]| \leq n^2 2^{-|uid|/2-1}$.

Game \mathbf{G}_2 behaves like \mathbf{G}_1 , except that it keeps a map from pairs of nonces to MAC keys. Instead of defining $k_m = \mathbf{kdf}(psk, n_I \mid n_R)$, game \mathbf{G}_2 looks up n_I, n_R in this map to retrieve the corresponding key k_m when it has been defined. Otherwise, it samples a fresh random MAC key k_m and stores the entry (n_I, n_R, k_m) in the map. We bound $|\Pr_{\mathbf{G}_1}[S] - \Pr_{\mathbf{G}_2}[S]| \leq \sum_{i=1}^{n_p} \mathbf{Adv}^{\text{PRF}}(\mathcal{B}_i)$: for each of the n_p pre-shared

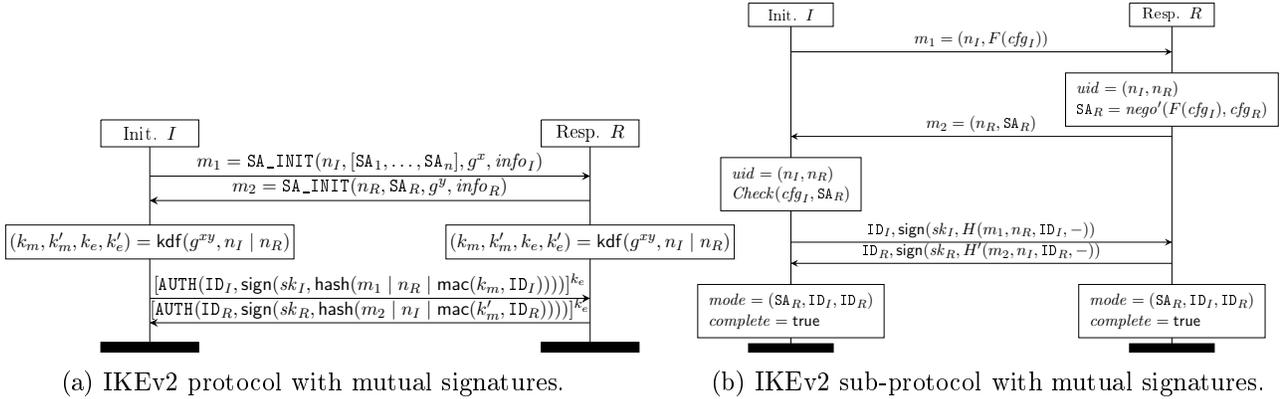


Figure 5: IKEv2 protocol and sub-protocol for signature-based authentication

secrets psk_i we define a hybrid, and bound the probability of \mathcal{A} distinguishing between adjacent hybrids by $\mathbf{Adv}^{\text{PRF}}(\mathcal{B}_i)$.

Observe that in \mathbf{G}_2 , every pair of nonces $(n_I, n_R) = uid$ appears at most once per role, and that peer instances with the same uid and configurations that contain only honest pre-shared keys share the same randomly sampled MAC key.

Game \mathbf{G}_3 is the same as Game \mathbf{G}_2 , except that sessions with honest pre-shared keys abort without completing whenever they verify a MAC on m_1 or m_2 that was not computed by another session. We bound the difference between the probability of S in these games using a reduction \mathcal{B}' to the EUF-CMA security of mac :

$$|\Pr_{\mathbf{G}_2}[S] - \Pr_{\mathbf{G}_3}[S]| \leq n \cdot \mathbf{Adv}^{\text{EUF-CMA}}(\mathcal{B}').$$

Observe that, since the MACs computed by initiators and responders cannot be confounded (because we assumed sessions abort when detecting $\text{ID}_I = \text{ID}_R$), in Game \mathbf{G}_3 we never verify a message that was not authenticated by its unique matching peer.

Note that both m_1 and m_2 contain the initiators contribution to nego and the responders choice, as well as all other components of mode . We thus can exclude the event S from ever occurring in \mathbf{G}_3 . \square

We also considered other modes of IKEv1, based on signatures instead of PSKs (much as in our introductory SIGMA example), and also when the MACs are protected using the keys derived from the Diffie-Hellman exchange. In those cases, the downgrade-protection sub-protocol is almost the same: SA_R is similarly left unauthenticated and, even if the messages are protected, there is still an attack when the client proposes a weak group, as explained in the introduction.

5.2 IKEv2 does not prevent downgrades

IKEv2 [30] is a revision of the IKEv1 protocol intended to simplify the specification and extend it to cover popular authentication methods such as EAP [1].

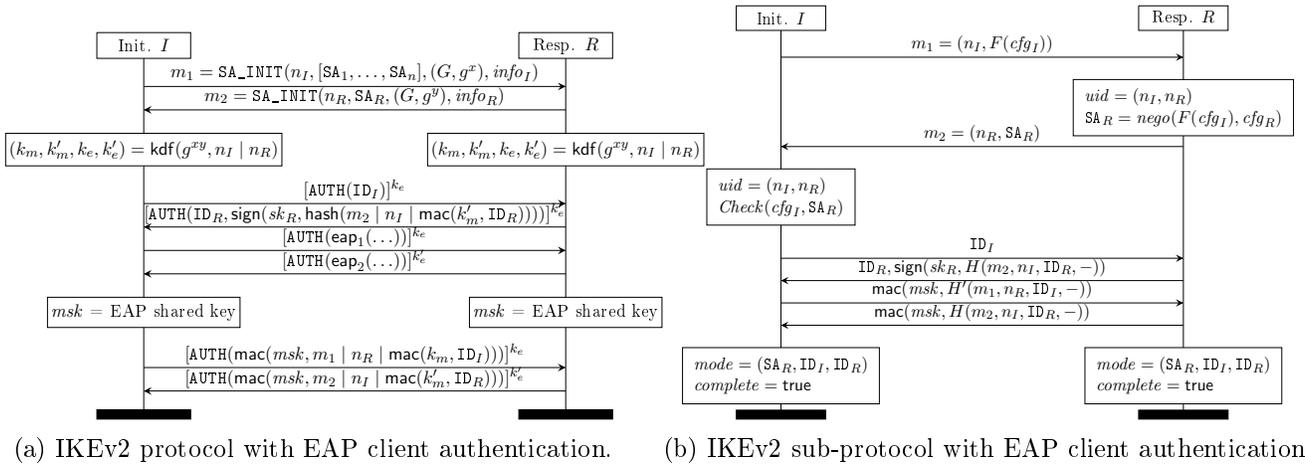


Figure 6: IKEv2 protocol and sub-protocol for EAP-based authentication

5.2.1 IKEv2 with signatures

We first consider the plain, signature-based sub-protocol (see Figure 5b). For brevity, we omit a description of the full protocol (its first messages appear in Figure 8) and we reuse the notations introduced for IKEv1. We also ignore signature agility issues, since in IKEv2 the hash algorithm for signing is not negotiated; it is chosen by the sender, who almost always picks SHA1.

As in IKEv1, the initiator begins by offering a sequence of security associations (extracted from cfg_I) and the responder chooses one of these. In the full protocol, the initiator and responder also exchange Diffie-Hellman public values and use them to derive session keys, used (in particular) to encrypt and MAC all messages after m_2 .

The client and the server then exchange signatures over MACs of their own views of the protocol (presumably to provide some deniability): their full first message, their identity, and the nonce of their peer. In particular, and in contrast with IKEv1, the server’s signature covers its chosen SA_R but not the initiator’s offered security associations.

The sub-protocol leaves important payloads unauthenticated: the peers do not sign or MAC each other’s DH public keys, and not even each other’s identities. It also ignores the fact that, in the full protocol, all messages after m_2 are encrypted and MACed using a derived key. Thus, some attacks against the sub-protocol may not occur in the full protocol.

Still, there is a downgrade attack against the full protocol as soon as the client tolerates one weak group. The attack proceeds as follows (see Fig. 7). Suppose an initiator offers two security associations, one using the 1024-bit Diffie-Hellman group 14 and another using the 768-bit group 1. The attacker tampers with the first message to delete the first association, so that the responder thinks that the initiator only supports group 1. The attacker forwards the responder’s messages, and now the initiator thinks that the responder only supports group 1. If the attacker has performed enough pre-computation on group 1 so that he can compute the discrete log of a key share, he can then compute the session and MAC keys and impersonate the responder.

In practice, executing this attack requires the MitM to send an extra `INVALID_KE` message to the client. This does not present any difficulty since this message is unauthenticated.

The attack described above is reminiscent of Logjam [2] and is arguably feasible with modern computing

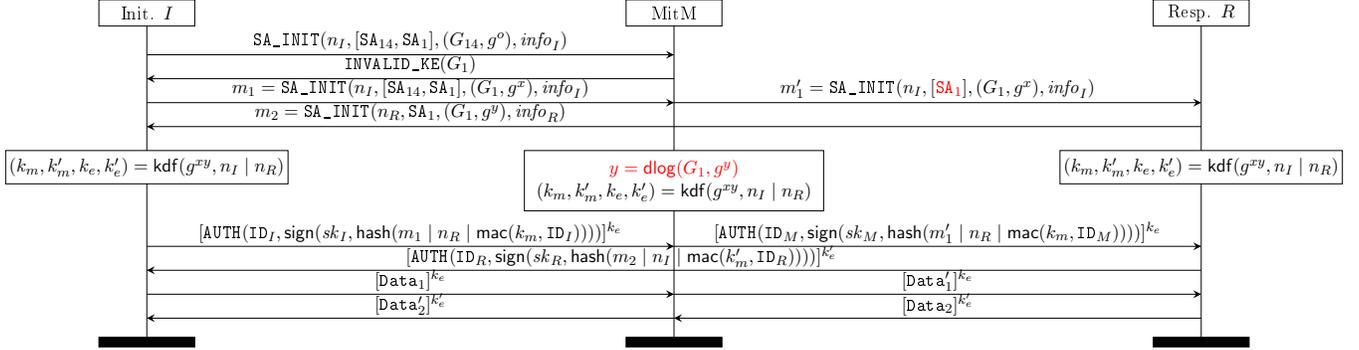


Figure 7: Man-in-the-middle downgrade on IKEv2 with mutual signatures and weak Diffie-Hellman groups

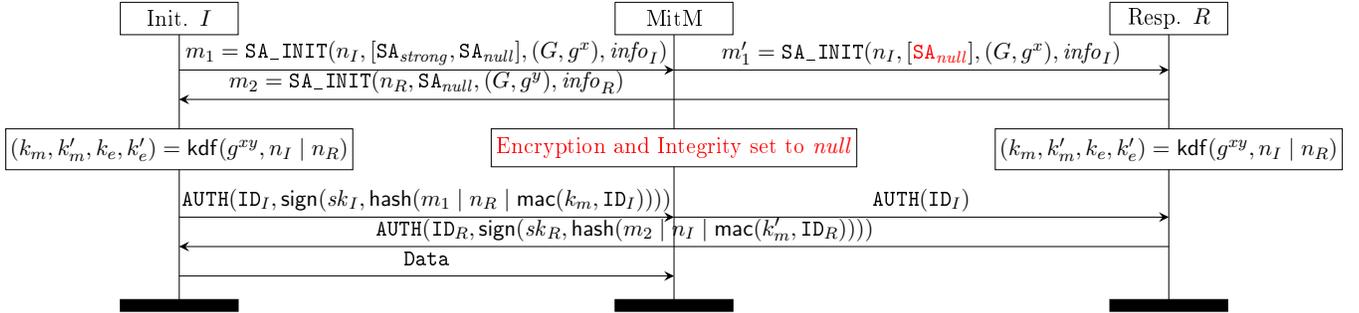


Figure 8: Man-in-the-middle cross-protocol downgrade on IKEv2 mixing signatures and EAP authentication

power, or will be in the coming years. There are other downgrade attacks with a similar impact on IKEv2: the man-in-the-middle could downgrade the security association to use weak encryption or authentication algorithms.

5.2.2 IKEv2 with EAP client authentication

We now consider the downgrade protection sub-protocol in case the initiator is authenticated using some EAP method, whereas the responder still uses a certificate and a signature (see Figure 6b).

In this variant, in the third message, the initiator sends its identity without any signature. Instead, after verifying the server's signature, it engages in an application-level 'embedded' authentication protocol that generates a shared key. Its use of EAP is asymmetric, in that EAP authenticates the initiator (ID_I) but does not re-authenticate the responder. The resulting shared key is used to MAC the initiator's view of the negotiation: the full first message, including the client's offered security associations, the responder's nonce, and a MAC over the initiator's identity with the session key.

Enabling EAP actually weakens downgrade protection: the responder (still) does not sign the initiator's proposals, and also does not sign the chosen client AUTH method (signature or EAP), and this opens the possibility of cross-authentication attacks between different AUTH methods.

For example, consider the attack in Fig. 8. Suppose the initiator disables EAP, but the responder supports it. The attacker can then replace the initiator's signature message with an EAP authentication message, forward the responder's signature, and thereby downgrade the SA used by the initiator, to use a

weak encryption algorithm, for instance. In comparison with the first attack on IKEv2 discussed above, this attack does not require breaking the Diffie-Hellman exchange to gain control of the key used to MAC the signature payloads.

This would be a powerful downgrade, and it would allow offline decryption of the initiator’s subsequent messages, but it is still difficult to implement in practice because the authentication messages are themselves encrypted-and-MACed. Hence, the attack requires that the attacker should be able to break the (downgraded) authenticated encryption mechanism in the SA.

For example, it can be mounted if the encryption and integrity algorithms are downgraded to NULL, an allowed (but not recommended) option in IKEv2. In particular, the specification says: “Though the security of negotiated Child SAs does not depend on the strength of the encryption and integrity protection negotiated in the IKE SA, implementations MUST NOT negotiate NONE as the IKE integrity protection algorithm or ENCR_NULL as the IKE encryption algorithm.” [30, Section 5]. Our attack shows that this assumption is wrong: the downgrade security of IKEv2 crucially depends on the strength of the encryption and integrity algorithms, especially when both signatures and EAP are enabled.

We also note that in case the initiator also supports EAP, any subsequent initiator authentication makes no difference since the initiator is now talking to the attacker and does not seek to re-authenticate the responder.

A simple fix for all these attacks would be for the responder to include the client’s first message and authentication mode in its signature (at the cost of losing deniability). We could then obtain downgrade protection simply by relying on the strength of the responder’s signature, irrespective of weak groups and broken encryption algorithms.

5.3 Version downgrades from IKEv2 to IKEv1

IKE does not include a version negotiation protocol. Initiators first try to connect with IKEv2 and if that fails they fall back to IKEv1. This allows a simple downgrade attack between these versions, since IKEv1 has no way of authenticating the highest supported version. The IKEv2 specification acknowledges this version downgrade possibility to IKEv1, but sets up a flag to prevent future downgrade attacks from IKEv($n > 2$) to IKEv2: “Note that IKEv1 does not follow these rules, because there is no way in v1 of noting that you are capable of speaking a higher version number. So an active attacker can trick two v2-capable nodes into speaking v1. When a v2-capable node negotiates down to v1, it should note that fact in its logs” [30, Section 2.5].

6 Z Real-Time Protocol

ZRTP [47] is a specialized protocol used to establish key material for encrypted voice-over-IP (VoIP) communications. Unlike TLS, ZRTP does not rely on public-key infrastructure or certificates for authentication. Instead, participants authenticate each other by comparing a “short authentication string” derived from the session key, also known as a SAS, via some trusted channel. For our purposes in this analysis, we assume in our model that the SAS comparison is conducted via an ideal, trusted channel that is not susceptible to tampering.

(Formally we model this using a $\text{CompareSAS}(\pi, \pi')$ oracle that makes session π compare its own *sas* to that of π' , if assigned. Session π must be in a waiting state before this oracle call and continues if the two match, otherwise aborts. We also provide a $\text{LeakSAS}(\pi)$ oracle that marks π ’s *sas* as dishonest and leaks it

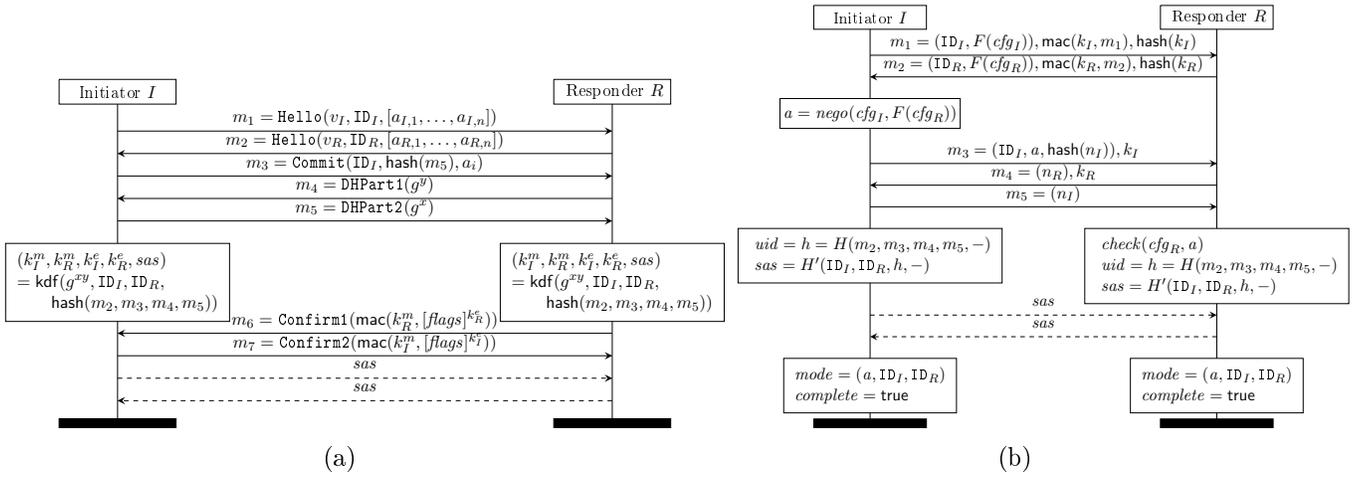


Figure 9: ZRTP (a) protocol (b) downgrade protection sub-protocol

to the adversary, and a FakeSAS(π, sas') oracle that let the adversary inject its own sas' to be compared with π 's sas .)

Because the SAS is short, the protocol offers a more limited form of protection. If the SAS length is ℓ bits, then the probability of an attacker subverting the authentication is at least $2^{-\ell}$ with each execution of the handshake. In most implementations ℓ is typically a small value, e.g. 16. The use of a short authentication string presents challenges for both key exchange and downgrade security. For example, if the SAS employed a full-length collision-resistant hash, it would suffice for the parties to exchange a hash of the full protocol transcript. However, even when constructed using a (truncated) collision-resistant hash function, the SAS is too short to provide the necessary protection, and additional measures must be taken.

6.1 ZRTP does not prevent downgrades

The ZRTP protocol is presented in Figure 9a. The downgrade protection sub-protocol is presented in Figure 9b. The ciphersuite negotiation is conducted within the first two (“Hello”) messages exchanged by the Initiator and the Responder. The chosen ciphersuite a_i is determined by selecting a ciphersuite in the intersection of the available algorithms presented by each party. Ciphersuites consist of a key exchange algorithm, a cipher and MAC algorithm for subsequent data exchange, and a SAS algorithm determining the length and format of the SAS string. Additionally, the protocol negotiates *options* such as a “trusted” PBX flag and an optional signature on the SAS.

Following the initial negotiation messages, the parties determine who will play the role of the Initiator, engage in a key exchange, and derive session keys. Transcript correctness is enforced by incorporating a hash of most of the transcript into the key derivation function, which produces both session keys and a SAS. A final mechanism tries to authenticate each of the handshake messages by computing a MAC over each message, using a key that is revealed in the subsequent message. To bind these messages together, ZRTP uses a hash chain.¹

¹Specifically, each participant computes an initial nonce H_0 and hashes it to obtain the sequence $H_3 = \text{hash}(H_2 = \text{hash}(H_1 = \text{hash}(H_0)))$. At each message in the handshake, the party reveals H_i and uses H_{i-1} as a MAC key to authenticate the current message. Verification is only possible when the next message is received. The initial value H_0 is revealed only within the encrypted confirmation message at the conclusion of the protocol.

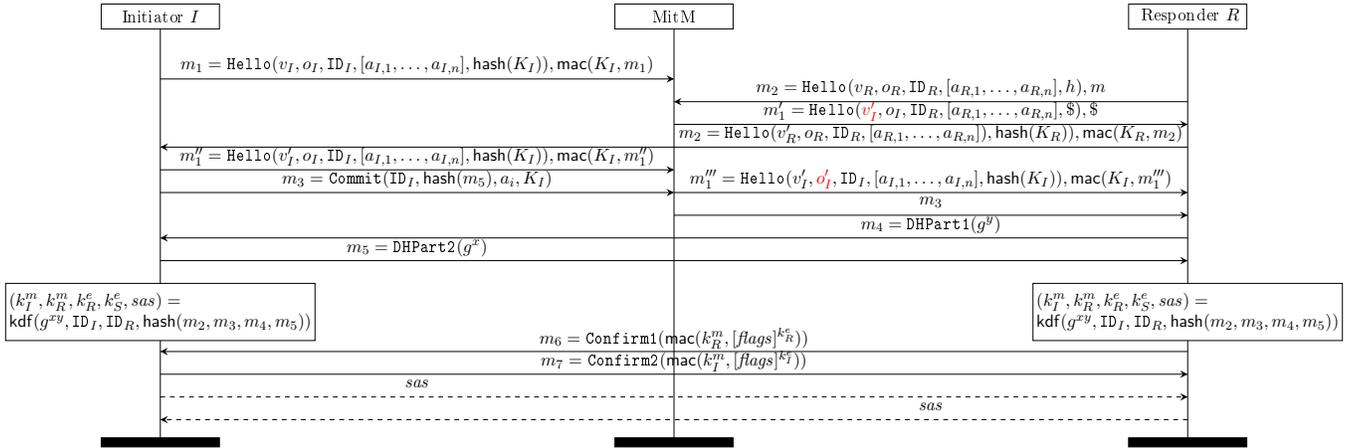


Figure 10: Man-in-the-Middle attack on ZRTP version and option negotiation. We assume that both peers prefer version $v_I = v_R$, but will support an older version $v'_I = v'_R$. The attacker additionally modifies the options flags o_I transmitted in m_1 .

Downgrading protocol versions ZRTP includes a negotiation mechanism for protocol versions and options that is not incorporated into the calculation of the shared secrets and SAS. When the parties support multiple versions of the protocol and protocol options, a MitM can substitute the protocol versions v_I, v_R to downgrade both parties to a previous version of the protocol, as illustrated in Figure 10. Moreover, since the first (Initiator Hello) message is not authenticated, the attacker can also change the options flags o_I . This second procedure requires the attacker to defeat the hash chain security mechanism. Unfortunately this may be done by capturing and delaying subsequent messages until the authentication key for earlier messages has been revealed, allowing the attacker to change messages arbitrarily. The fix for this vulnerability is straightforward: all negotiation messages should be included in the calculation of the session key and SAS.

Downgrade from DH to PSK ZRTP supports both Diffie-Hellman key exchange and a pre-shared key mode. The latter is analogous to the session resumption handshake in TLS, in that it provides an inexpensive (symmetric-key only) handshake, which operates under the assumption that the parties have previously completed a full Diffie-Hellman handshake to establish a pre-shared key. The corresponding negotiation sub-protocol is shown in Fig. 11. The limitation of this pre-shared mode is that it does not force the parties to commit to their protocol inputs before revealing them, which admits an offline attack in which a MitM may identify protocol inputs that result in a chosen SAS. A full attack is shown in Figure 12, beginning with the establishment of a shared key (via Diffie-Hellman) and restarting with the PSK mode.

7 Transport Layer Security

The Transport Layer Security protocol (TLS) is used to provide secure channels for a variety of Internet applications. It offers a number of key exchange mechanisms, authentication methods, and encryption schemes, so that users can pick and choose mechanisms best suited to their needs.

A negative consequence of this agility is the potential for downgrades. TLS clients and servers commonly support multiple protocol versions and hundreds of ciphersuites, even though some of them are known to be obsolete or even broken. For example, SSL 2 is still supported by 10% of web servers even though it has long

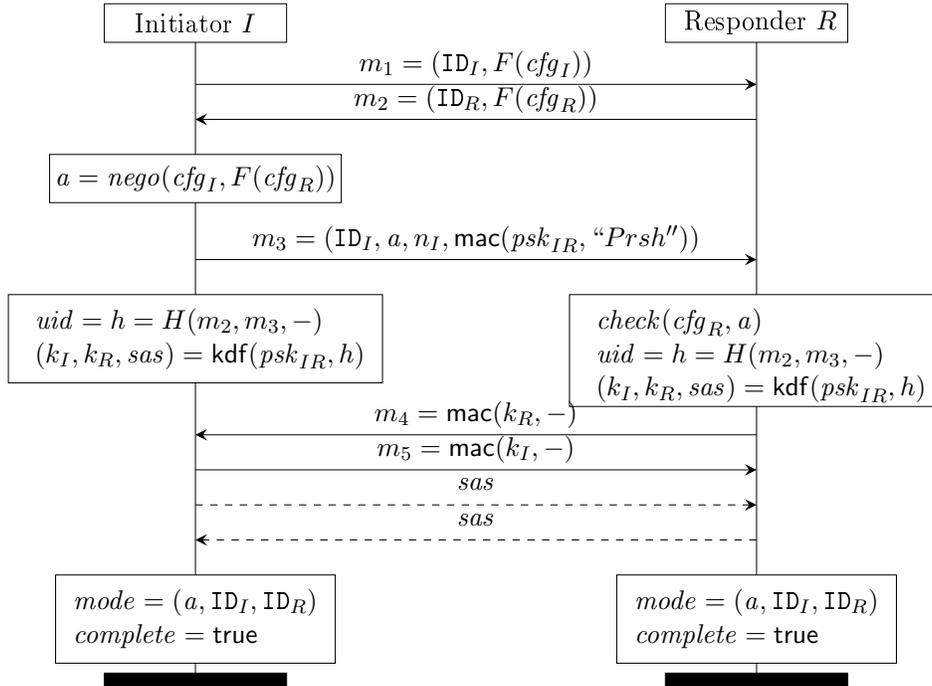


Figure 11: ZRTP with Pre-Shared Keys: negotiation sub-protocol

been known to be vulnerable to multiple attacks including, notably, a ciphersuite downgrade attack [44]. More recently, about 25% of web servers were found to still support export-grade ciphersuites that were deprecated in 2000, enabling powerful downgrade and server impersonation attacks like FREAK [10] and Logjam [2].

Since SSL 3, all versions of TLS incorporate various downgrade protection mechanisms. We will analyze the downgrade protection provided by TLS 1.2 and the proposed improvements in TLS 1.3.

7.1 Negotiation in TLS 1.2

Figure 13a depicts a mutually authenticated TLS connection incorporating a Diffie-Hellman key exchange that uses either a finite-field group (DHE) or an elliptic curve (ECDHE). Most TLS connections authenticate only the server, but the figure also depicts the optional client authentication messages.

The client I first sends a hello message (CH) with a nonce (n_I) and a list of agility parameters $[a_1, \dots, a_n]$ that include ciphersuites, compression methods, and protocol extensions. The server responds with a hello message (SH) containing its chosen parameters (a_R). At this point, the client and server know which key exchange they will execute next. In an ephemeral Diffie-Hellman key exchange (DHE/ECDHE), the server sends its public-key certificate (cert_R) and uses the private key to sign the nonces, the group (or curve) parameters (p, g) and its own Diffie-Hellman public value (g^y). The server may let the client remain anonymous, or it may require authentication (specifying the class of acceptable certificates $[\text{cert}_1, \dots, \text{cert}_m]$), in which case the client sends its own certificate (cert_I) and public value (g^x), and uses its private key to sign the full protocol transcript so far (log_1). The client and server then derive a master secret (ms) and session keys (k_1, k_2) from the nonces and shared secret (g^{xy}). To complete the key exchange, both sides com-

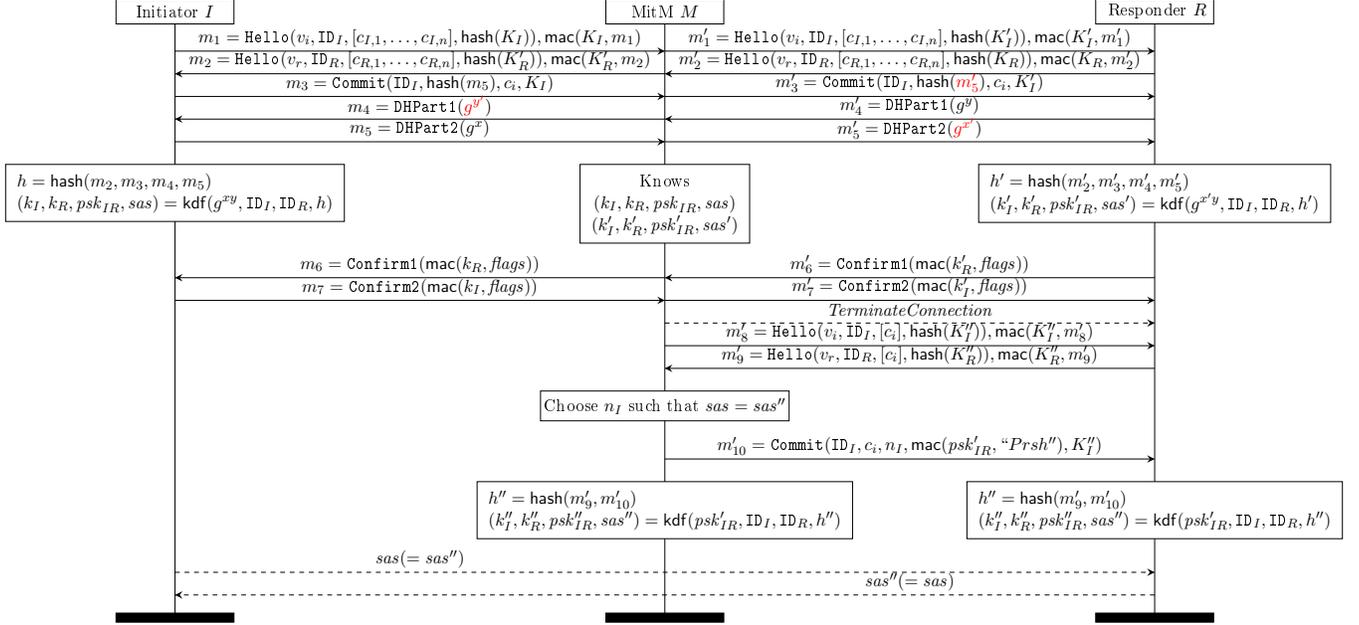


Figure 12: Man-in-the-middle cross-protocol attack on ZRTP with Diffie-Hellman exchange followed immediately by a pre-shared key exchange: M tampers with the first exchange between I and R and disconnects before they can compare their (differing) SAS values. M then runs a new pre-shared exchange with R and is able to synchronize the SAS values at I and R by finding an appropriate nonce (with 2^{16} work).

pute MACs using the master secret over the protocol transcript, and exchange them in finished messages (CFIN,SFIN). These MACs provide key confirmation as well as downgrade protection. Once exchanged, the client and server can start exchanging application data encrypted under the new session keys ($[\text{Data}]^k$).

7.2 TLS 1.2 does not prevent downgrades

The downgrade protection sub-protocol for TLS 1.2 is depicted in Fig. 13b. For simplicity, we consider only server-authenticated (EC)DHE connections, where clients are anonymous.

The client offers its entire public configuration ($F(\text{cfg}_I)$) to the server, which then computes the negotiated parameters ($mode$) that consist of the protocol version (v), the chosen parameters (a_R), the group (G_R), the server identity (pk_R), and the hash function used in the server signature (hash_1). The protocol version and the ciphersuite in a_R together determine other protocol parameters, such as the key derivation function (kdf), the authenticated encryption scheme, and the MAC and hash functions used in the finished messages (mac , hash). We note that the server may possess several identities and choose one based on the chosen ciphersuite or other protocol extensions offered by the client.

Downgrade protection primarily relies on the MACs in the finished messages, which in turn rely on the strength of the group G_R and the negotiated algorithms kdf , hash , and mac . If a client and server support a weak group, for example, then an attacker can downgrade the group and then break the master secret to forge the MACs, as in Logjam.

A second protection mechanism is the server signature, but we observe that this signature covers only the unique identifier and the group G_R , but none of the other negotiated parameters. For example, the

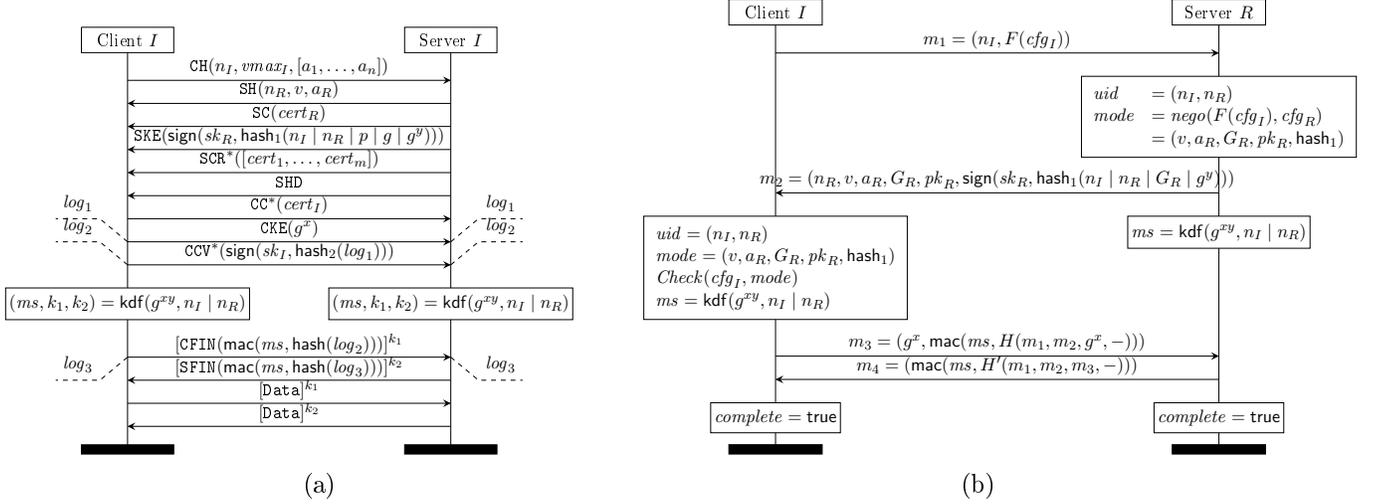


Figure 13: TLS 1.2 with (EC)DHE key exchange (a), where messages labeled with * occur only when client authentication is enabled, and (b) its downgrade protection sub-protocol

Logjam attacker tricks the server into using an export ciphersuite (DHE-EXPORT) that results in a weak Diffie-Hellman group. The client does not support DHE-EXPORT and still thinks it is using standard DHE, but the attacker can forge the MAC to hide this discrepancy. Importantly, the server signature fails to prevent this attack, because it does not include the ciphersuite. Before this attack was disclosed, many implementations of TLS clients still accepted arbitrary groups.

Furthermore, we note that the negotiated algorithms can be weak in practice. For example, TLS 1.2 supports MD5-based signatures; TLS 1.1 derives keys and transcript hashes based on combinations of MD5 and SHA1. These weak constructions also lead to downgrade and impersonation attacks [11].

Let min_r, max_r be the supported minimum and maximum protocol versions, let $algs_r = [a_1, \dots, a_m] = F(cfg_I)$ be the ciphersuites and extensions, and let $groups_r$ be the groups supported by role r . In terms of the general definition in Section 2, the downgrade protection sub-protocol uses the following session variables:

$$\begin{aligned}
 cfg &\triangleq \begin{cases} (I, min_I, max_I, algs_I, groups_I, PKs_R) & \text{for } I \\ (R, min_R, max_R, algs_R, groups_R, PK_R) & \text{for } R \end{cases} \\
 uid &\triangleq (n_I, n_R) \\
 mode &\triangleq (v, a_R, G_R, pk_R, \text{hash}_1)
 \end{aligned}$$

The negotiation function nego is executed by the server and is based on the server's configuration cfg_R and the server's partial view $F(cfg_I)$ of the client configuration. The client does not get to inspect cfg_R , but it does check that the resulting mode is consistent with its configuration.

The protocol only offers downgrade protection if the peer is authenticated with an honest key and strong signature and hash algorithms. So we will consider downgrade security from the viewpoint of a client, while assuming that all keys in PKs_R are honest and hash_1 is collision-resistant. We get partnering security from the freshness of the uid and the strength of the server signature (which includes the uid).

However, downgrade protection for the client cannot rely on just the signature, and hence requires one of the following conditions:

- the server uses its pk_R only with modes that use strong groups, key derivation algorithm kdf , hash and

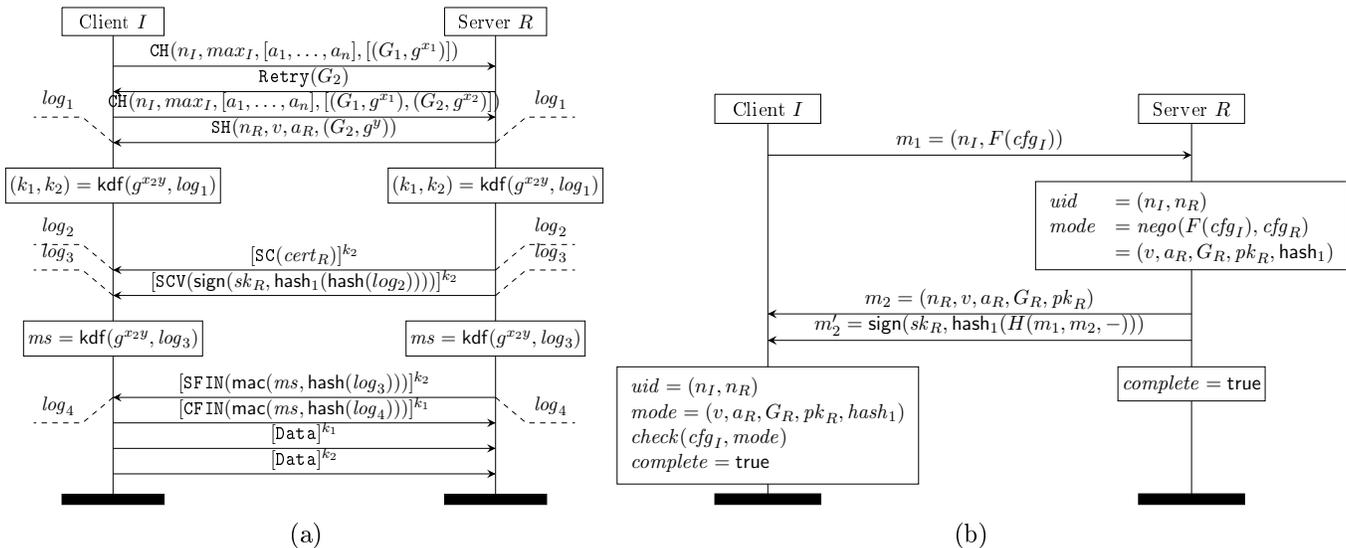


Figure 14: TLS 1.3 1-RTT mode with server-only authentication (a) and its downgrade protection sub-protocol (b)

mac algorithms and the client is aware of the servers choice and aborts whenever it sees an unexpected algorithm combination;

- the client only accepts modes with strong groups (in particular *not* the groups ‘negotiated’ by the Logjam and the ECDHE-DHE cross-protocol attacks [36]) and algorithms.

An extreme example of the first condition would be to require that the server uses a different public key for each mode; the proofs in [23] rely on this somewhat unrealistic assumption to avoid ECDHE-DHE cross-protocol attacks and the need for agile security assumptions. More pragmatically, if a client and server only support TLS 1.2 (and hence only strong hash constructions), only support strong groups and curves for (EC)DHE and all other ciphersuites that use Diffie-Hellman, then TLS clients can be protected from downgrade. Of course, we rely on the server using only honest and strong signing keys (e.g. 2048-bit RSA) with strong signature and hash algorithms (e.g. RSA-SHA256).

We also get some downgrade protection for the server when the client is authenticated, relying only on the client signature and the transcript hash algorithm `hash`. Pragmatically, TLS 1.2 servers that require client authentication and only accept strong signature and hash algorithms cannot themselves be tricked into completing a connection with a weak *mode*.

As evidenced by the Logjam attack, the TLS protocol does not satisfy downgrade security unless the *DP* predicate guarantees that the client and server configurations exclusively use strong algorithms, hence guaranteeing that all the negotiated algorithms used in the finished MACs are strong.

7.3 On downgrade protection in Draft 10 of TLS 1.3

Draft 10 of TLS 1.3, the next version of TLS, proposes a protocol that is quite different from TLS 1.2; a typical run of the 1-round-trip mode is depicted in Fig. 14a. The corresponding downgrade protection sub-protocol is in Fig. 14b.

In contrast to TLS 1.2, the client hello message includes Diffie-Hellman public values for the client’s preferred groups. The server may choose one of these groups or ask for a public value in a different group, as long as it is one supported by the client. The server sends its own public value in the server hello message, and all subsequent messages are encrypted and integrity-protected using the Diffie-Hellman shared key.

For downgrade security from the client’s viewpoint, a key difference is that server signatures in TLS 1.3 cover the full transcript, and hence they cover the full client and server hello messages. This foils most of the downgrade attacks on TLS 1.2; as long as the client only accepts strong signature and hash algorithms and honest public keys from the server, it cannot be downgraded to a weaker ciphersuite, and moreover, it yields agreement on the chosen ciphersuite.

Although Draft 10 of TLS 1.3 provides strong downgrade protection for the ciphersuite, downgrade attacks remain, in particular, because clients and servers will continue to support lower protocol versions for backward compatibility. Considering that TLS 1.2 does not provide strong downgrade protections, this unfortunately means that all the downgrade attacks on TLS 1.2 will be inherited by TLS 1.3.

There are three downgrade attacks possible on TLS 1.3 as described in Draft 10. One, an attacker downgrades the connection to TLS 1.2 or lower and mounts any of the downgrade attacks mentioned before. This will succeed as long as the attacker can forge the finished MACs. Second, an attacker uses the TLS fallback mechanism to stop TLS 1.3 connections and allows only TLS 1.2 connections to go through. Even if the endpoints implement the fallback protection mechanism [39], the attacker can use one of the downgrade attacks in TLS 1.2 to break the connection. Third, in Draft 10 of the TLS 1.3 protocol, the handshake hashes restart upon receiving a **Retry** message and hence, the attacker can downgrade the Diffie-Hellman group for some classes of negotiation functions.

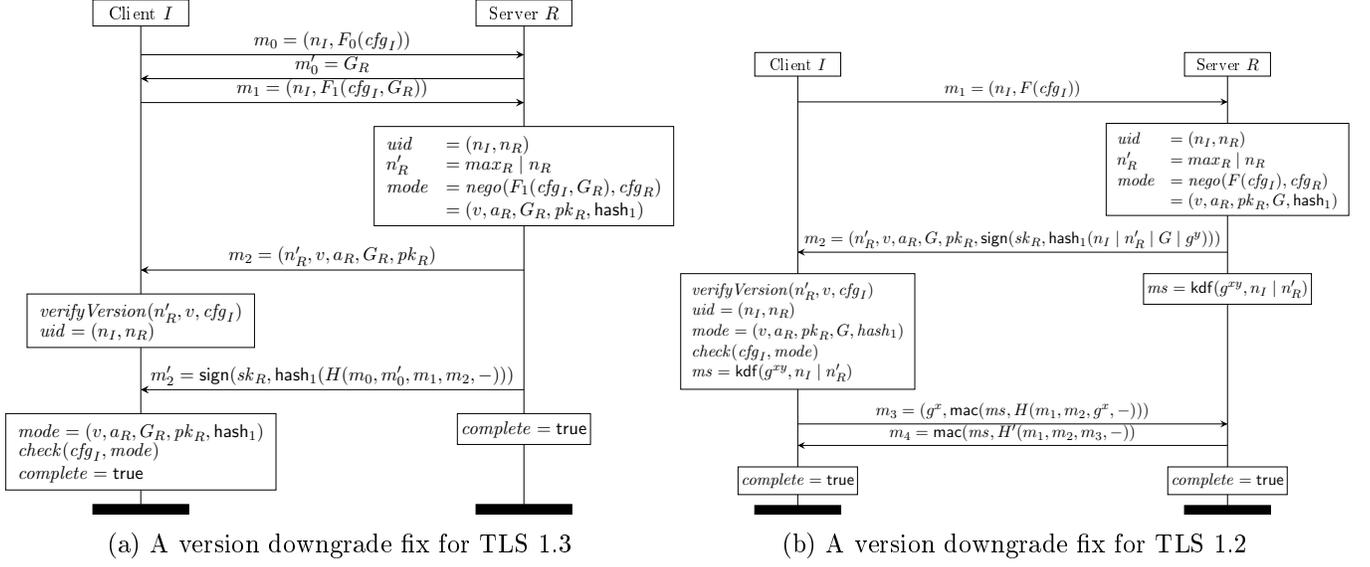
We can prevent all of these attacks by two countermeasures, both of which have been incorporated into TLS 1.3 Draft 11. See Fig. 15a. First, we continue the handshake hashes over retries. Second, TLS 1.3 servers always include their highest supported version number in the server nonce, even when they choose a lower version such as TLS 1.0.

Including the maximum version number into the server nonce of all versions yields *version* downgrade protection for clients. It is a simple patch (For the server, it amounts to changing how nonces are generated. The client needs to implement an equality check.) that can be incorporated into TLS versions without making them incompatible with TLS versions that do not implement the patch. If a server and a client both implement the patch, the client gets version downgrade protection.

We proceed in three steps: We show that when hashes continue over **Retry**, clients that interact with servers that just support TLS 1.3 achieve downgrade security. We then show that embedding the version number into the server’s nonce yields *version* downgrade protection from the client’s perspective. We then put the two results together and show that the composition of TLS 1.2 and TLS 1.3 with these countermeasures provides the same client-side downgrade protection as when servers just support TLS 1.3.

The downgrade protection sub-protocol uses the same session variables as for TLS 1.2, but defines *Nego* using the function *nego* from Fig. 15a. Let \mathcal{M} be the set of modes supported by TLS and $\mathcal{M}^* = \{Nego_{cfg.role}(cfg, cfg') | PS(cfg)\}$ be the modes negotiated between any pair of configurations for which the first guarantees partnering security. Let $\mathcal{P}_s = \{p \mid s, p = mode.\overline{sig} \wedge mode \in \mathcal{M}\}$ be the signature agility parameters for peer signature scheme s , \mathcal{H} be the set of all hash algorithms supported by TLS, and $\mathcal{H}^* = \{mode.hash \mid mode \in \mathcal{M}^*\}$ be the hash algorithms used by partnering secure modes. We now prove partnering security for TLS 1.2 and 1.3, and downgrade security for clients speaking to servers that implement the fix described in Fig. 15a. We then define version downgrade security and show that the fixes in Fig. 15a (TLS 1.3) and Fig. 15b (TLS 1.2) prevent version downgrade.

Theorem 8 (Partnering security of *TLS* for clients). *We consider a universe of configurations where*



RSA keys are used for signing or encryption, but not for both. Let PS be such that $PS(cfg)$ implies that $cfg.role = I$, that all public keys in the range of $cfg.PKs_R$ are honest and that cfg does not support RSA key transport. Given an adversary \mathcal{A} against the partnering security of TLS, we construct adversaries $\mathcal{B}_{s,p,i}$ and \mathcal{B}_h running in about the same time as \mathcal{A} such that $\text{Adv}_{TLS, PS}^{\text{partnering}}(\mathcal{A})$ is at most

$$\sum_{h \in \mathcal{H}^*} \text{Adv}_{h, \mathcal{H}}^{CR}(\mathcal{B}_h) + \sum_{(s,p) \in \text{sig}(\mathcal{M}^*)} \sum_{i=1}^{n_s} \text{Adv}_{s,p, \mathcal{P}_s}^{EUF-CMA}(\mathcal{B}_{s,p,i}),$$

where n_s keys are generated for signing scheme s .

The proof is the same as for Theorem 4 for SSH except that we only need to prove the property for clients and rely on the nonces being hashed and signed as part of messages m_0 and m_1 .

For downgrade security, we define $Nego$, \mathcal{M} , \mathcal{P}_s , and \mathcal{H} as before. However, we redefine \mathcal{M}^* , \mathcal{H}^* to use DP instead of PS , i.e., $\mathcal{M}^* = \{Nego_{cfg.role}(cfg, cfg') \mid DS(cfg, cfg')\}$ and $\mathcal{H}^* = \{mode.hash \mid mode \in \mathcal{M}^*\}$.

Theorem 9 (Downgrade security of *TLS1.3-sub*). *We consider a universe of configurations where RSA keys are used for signing or encryption, but not for both.*

*Let DP be such that $DP(cfg, cfg')$ implies that $cfg.role = I$, that all public keys in the range of $cfg.PKs_R$ are honest, that cfg and cfg' implement the countermeasure, such that cfg' only supports TLS 1.3 and cfg only supports TLS 1.3 or TLS 1.3 & 1.2 and does not support RSA key transport. Given an adversary \mathcal{A} against the downgrade security of *TLS1.3-sub*, we construct adversaries $\mathcal{B}_{s,p,i}$ and \mathcal{B}_h running in about the same time as \mathcal{A} such that $\text{Adv}_{TLS1.3-sub, DP}^{\text{downgrade}}(\mathcal{A})$ is at most*

$$\frac{n^2}{2^{|n_R|+1}} + \sum_{h \in \mathcal{H}^*} \text{Adv}_{h, \mathcal{H}}^{CR}(\mathcal{B}_h) + \sum_{(s,p) \in \text{sig}(\mathcal{M}^*)} \sum_{i=1}^{n_s} \text{Adv}_{s,p, \mathcal{P}_s}^{EUF-CMA}(\mathcal{B}_{s,p,i}),$$

where n is the number of sessions, n_s is the number of keys generated for signing scheme s , and $|n_R|$ is the size of the servers contribution to the unique identifiers. The current proposal is 24 bytes.

The proof again follows the same structure as for Theorem 5 for SSH except that we only need to prove the property for clients.

Proof sketch. Game G_0 is the original downgrade experiment. In game G_1 we guarantee that there is at most one partnered server session. By the collision probability of server nonces we have that $|\Pr_{G_0}[S] - \Pr_{G_1}[S]| \leq n^2(2^{-|n_R|-1})$. In games G_2 and G_3 we abort on agile hash function collisions and signature forgeries for modes in \mathcal{M}^* .

The final reasoning step follows a different pattern than SSH. As the server signs all information that is taken into account by *nego* together with the corresponding nonces, the signature over the protocol messages must have been generated by the server session with the unique identifier corresponding to these nonces. Moreover, if there had been any modification to the messages of the server and the client that influence *nego*, then the client would have aborted in Game G_3 . \square

We define version downgrade security similarly to downgrade security via a function $Version_r$ that maps two opposite-role configurations (which include the version numbers) to the version number negotiated (if any) in the absence of active adversaries. Formally, if a session π talking to a session π' completes, it must be the case that $\pi.v = Version_r(cfg_r, cfg_{\bar{r}})$. Akin to downgrade security, our definition of version downgrade security is parameterized by a version downgrade protection predicate VDP on pairs of configurations. When $VDP(cfg_r, cfg_{\bar{r}})$ holds, we expect that the local instance r is protected. For TLS, we will only consider version downgrade protection from the client's perspective.

Definition 11 (Version downgrade security). *The advantage $\mathbf{Adv}_{\Pi, VDP}^{version}(\mathcal{A})$ of \mathcal{A} against the version downgrade security of Π is the probability that, when \mathcal{A} terminates after interacting with protocol Π through its oracles, there exists a session π such that $\pi.complete = \mathbf{true}$ and there is a partnered session π' such that $VDP(\pi.cfg, \pi'.cfg)$ but $\pi.v \neq Version_{\pi.role}(\pi.cfg, \pi'.cfg)$.*

In the variant of Version Downgrade Security that includes version agreement, the adversary also wins if there exists a session π such that $\pi.complete = \mathbf{true}$ and there is a partnered session π' such that $VDP(\pi.cfg, \pi'.cfg)$ but π does not agree with π' on v .

Theorem 10 (Version downgrade security of $TLS1.3-TLS1.2-sub$). *We consider a universe of configurations where RSA keys are used for signing or encryption, but not for both. Let VDP be such that $VDP(cfg, cfg')$ implies that $cfg.role = I$ and that all public keys in the range of $cfg.PKs_R$ are honest and such that: (a) cfg activates the countermeasure for all its versions and does not support RSA key transport. (b) cfg' activates the countermeasure for all its versions.*

Given an adversary \mathcal{A} against the version downgrade security of $TLS1.3-TLS1.2-sub$, we construct adversaries $\mathcal{B}_{s,p,i}$ and \mathcal{B}_h running in about the same time as \mathcal{A} such that $\mathbf{Adv}_{TLS1.3-TLS1.2-sub, VDP}^{version}(\mathcal{A})$ is at most

$$\frac{n^2}{2^{|n_R|+1}} + \sum_{h \in \mathcal{H}^*} \mathbf{Adv}_{h, \mathcal{H}}^{CR}(\mathcal{B}_h) + \sum_{(s,p) \in \mathit{sig}(\mathcal{M}^*)} \sum_{i=1}^{n_s} \mathbf{Adv}_{s,p, \mathcal{P}_s}^{EUF-CMA}(\mathcal{B}_{s,p,i}),$$

where n is the number of sessions, n_s is the number of keys generated for signing scheme s , and $|n_R|$ is the size of the servers contribution to the unique identifiers. The current proposal is 24 bytes.

If the server configuration supports at most TLS 1.3 and TLS 1.2, the above bound also holds for the variant of version downgrade security that assures version agreement.

Proof sketch. Let S be the event that there exists a session π that completed and is version downgraded, i.e., there exists a partner session π' with the same *uid* and opposite role and such that $\pi.v \neq \text{Version}_{\pi,\text{role}}(\pi.\text{cfg}, \pi'.\text{cfg})$

Let \mathbf{G}_0 be the original version downgrade security experiment. Game \mathbf{G}_1 aborts whenever two server sessions assign the same n_R . Let n be the total number of sessions. The probability that n such random values give rise to a collision is upper bounded by $n^2 2^{-|n_R|-1}$. As \mathbf{G}_0 and \mathbf{G}_1 are identical up to collisions, we have $|\Pr_{\mathbf{G}_0}[S] - \Pr_{\mathbf{G}_1}[S]| \leq n^2 2^{-|n_R|/2-1}$. Game \mathbf{G}_2 behaves as \mathbf{G}_1 , except that it aborts whenever two sessions hash different transcripts to the same *log*. Using \mathcal{A} , we can construct for each $h \in \mathcal{H}^*$ an adversary \mathcal{B}_h such that $|\Pr_{\mathbf{G}_0}[S] - \Pr_{\mathbf{G}_1}[S]| \leq \sum_{h \in \mathcal{H}^*} \mathbf{Adv}_{h, \mathcal{H}}^{\text{CR}}(\mathcal{B}_h)$. Observe that in \mathbf{G}_2 , every n_R is signed at most once by the responder.

Game \mathbf{G}_3 is the same as Game \mathbf{G}_2 , except that sessions with honest peer keys and $(s, p) \in \overline{\text{sig}}(\mathcal{M})$ abort without completing whenever they verify a signature that was not produced by another session.

We again bound the difference in the success probabilities using the lemma of Bhargavan et al. [14, Lemma 4] and have

$$|\Pr_{\mathbf{G}_2}[S] - \Pr_{\mathbf{G}_3}[S]| \leq \sum_{(s,p) \in \overline{\text{sig}}(\mathcal{M}^*)} \sum_{i=1}^{n_s} \mathbf{Adv}_{s,p,\mathcal{P}_s}^{\text{EUF-CMA}}(\mathcal{B}_{s,p,i}).$$

Observe that, since the signatures computed by initiators and responders cannot be confounded, in Game \mathbf{G}_3 we never verify a message that was not signed by its unique partner.

For the initiator we have that message *log* was signed by its partner and contains the responder's nonce. As the *VDP* holds, the responder's nonce contains the maximal supported version number and the initiator verifies that this matches the one that the initiator either proposed or received, the initiator yields the preferred protocol version.

Generally, we do not get agreement on v , because the responder does not sign the initiator's maximal version number and hence, even if the initiator completes, he has no guarantees that the server actually received the highest protocol version. However, if the server only supports TLS 1.3 and TLS 1.2, then the domain separation property of signatures guarantees that initiator and responder have the same protocol version. \square

Note that we would generally obtain agreement on the version number if the client embeds his highest supported version number into his nonce.

For a downgrade predicate DP and a version predicate VDP such that $DP \subseteq VDP$, let DP^{+VDP} be the predicate that holds for pair of configurations in DP , with server configurations extended to also support configurations of lower version protocols that by VDP should never be negotiated. Putting Theorem 10 and Theorem 9 together, we get that when both client and server implement the countermeasures, then supporting TLS 1.3 & 1.2 is as good as supporting only TLS 1.3.

Corollary 1 (Downgrade security of *TLS1.3-TLS1.2-sub*). *Let $DP \subseteq VDP$ and let VDP be such that $VDP(\text{cfg}, \cdot)$ implies that $\text{cfg.role} = I$, that all public keys in the range of cfg.PKs_R are honest and that*

1. *cfg only supports TLS 1.3 or TLS 1.3 & 1.2, activates the countermeasure for all its versions, and does not support RSA key transport,*
2. *cfg' only supports TLS 1.3 or TLS 1.3 & 1.2 and activates the countermeasure for all its versions.*

Given an adversary \mathcal{A} against the downgrade security of TLS1.3-TLS1.2-sub, we construct adversaries \mathcal{B} and \mathcal{C} running in about the same time as \mathcal{A} such that

$$\mathbf{Adv}_{\text{TLS1.3-1.2-sub}, DP'}^{\text{downgrade}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{TLS1.3-1.2-sub}, VDP}^{\text{version}}(\mathcal{B}) + \mathbf{Adv}_{\text{TLS1.3-sub}, DP}^{\text{downgrade}}(\mathcal{C})$$

where $DP' = DP \cup DP^{+VDP}$.

Proof sketch. Theorem 10 ensures that if the client and server configurations both support TLS 1.3 and TLS 1.2, then they get TLS 1.3 and Theorem 9 ensures that a client session that supports TLS 1.3 and TLS 1.2 gets the preferred mode when interacting with a server that only supports TLS 1.3.

The probability that the server partner session of a completing client that supports TLS 1.3 or TLS 1.3 and TLS 1.2 assigns a different version variable than 1.3 is upper bounded by $\mathbf{Adv}_{\text{TLS1.3-1.2-sub}, VDP}^{\text{version}}(\mathcal{B})$. Note that crucially we rely on agreement on the version variable.

The main proof idea is that servers that support TLS 1.3 and 1.2 and server that support only TLS 1.3 behave the same way *after* receiving 1.3 as the client’s highest version numbers. Hence, in a first game-hop, we wait until the first Send query to a server session, and if this first query contains version number 1.3, we initialize the server with a configuration that only supports 1.3 (and besides has the configuration that the adversary provided). Else, we initialize the server with the configuration that the adversary provided. Now, all servers that run TLS 1.3 in a particular session only support TLS 1.3.

The probability that a client that supports either TLS 1.3 or TLS 1.3 and 1.2 when interacting with such a server is downgraded is upper bounded by $\mathbf{Adv}_{\text{TLS1.3-sub}, DP}^{\text{downgrade}}(\mathcal{C})$. \square

8 Related Work

Downgrade as an attack vector The importance of downgrades when building practical exploits against key exchange protocols has been widely recognized [40, 10, 44, 2].

The lessons learned from this are less clear. There is a disconnect between the IETF, implementers, penetration testers, and protocol analysts. Browser developers are frequently criticized for prioritized interoperability over security and standard compliance.

RFC 7507 [39] proposes the Signaling Cipher Suite Value (SCSV) extension for TLS to prevent version downgrade attacks when the key-exchange of all versions provides transcript authentication. SSL2.0 and SSL3.0 are being deprecated, partly to prevent version downgrade attacks as these versions do not support said extensions [42, 5] and SSL2.0 in any case does not provide reliable transcript authentication. Similarly, ciphersuite hygiene is frequently discussed in standard documents [35, 34]

Previous downgrade security theorems about TLS 1.2 Dowling and Stebila [23] model ciphersuite and version negotiation for the TLS protocol up to version 1.2 in the multi-ciphersuite setting introduced by Bergsma et al. [9] (discussed below). In our model, their result corresponds to a proof of downgrade security for a $DP(\pi.cfg, \pi'.cfg)$ predicate that guarantees that all negotiable ciphersuites and versions are strong enough to provide ACCE security and that all public keys are honest and used at most by one negotiable ciphersuite. Their optimality function ω is a more limited variant of our *Nego* function and does not include entity identifiers. Their main theorem states that under such strong conditions multi-mode authentication implies downgrade security.

This is a rather weak form of downgrade security, but as shown by our attack, TLS 1.2 does not provide much stronger protection at least for clients. Servers that authenticate clients can however receive stronger guarantees.

Related work for SSH Bergsma et al. [9] previously analyzed SSH in a multi-ciphersuite setting. They split the protocol into a negotiation phase \mathbf{NP} and key-exchange phase \mathbf{SP} , one for each value of $\pi.mode$. They show that if each combination $\mathbf{NP}\|\mathbf{SP}$ is ACCE secure, then $\mathbf{NP}\|\vec{\mathbf{SP}}$ is multi-ciphersuite ACCE secure. While they do not prove downgrade security per se, the result of [23] adapted to SSH corresponds to a proof of downgrade security for a $DP(\pi.cfg, \pi'.cfg)$ predicate that guarantees that all negotiable ciphersuites and versions are strong enough to provide ACCE security.

The sharing of the public key is admissible under the condition that each sub-protocol provides sufficient oracle access to the long-term key functionality, e.g., signing, to simulate all other sub-protocols. In our terminology, the protocols $\mathbf{NP}\|\mathbf{SP}$ of [9] are single mode restrictions of $\mathbf{NP}\|\vec{\mathbf{SP}}$. After their extensions with oracles providing sufficient access to long-term key functionalities, they are also sub-protocols in our sense.

We prove downgrade protection for a predicate DP that includes a much larger set of configurations. Combined with the result of Bergsma et al., our result allows to prove multi-ciphersuite ACCE security when not all sub-protocols in $\vec{\mathbf{SP}}$ are ACCE secure, as long as we restrict the protocol to configurations in DP that do not negotiate them (cf. Theorem 1).

9 Conclusion and Future Work

In this paper we put forward a methodology to analyze the downgrade security of real-world key exchange protocols. Our approach breaks down the complexity of analyzing a full protocol by considering only a core sub-protocol that abstracts away details that are irrelevant for negotiation of a protocol mode. We showed that proving a simulatability property for a sub-protocol is sufficient for ensuring the *soundness* of our methodology: proving the absence of downgrade attacks on the sub-protocol is enough to guarantee the downgrade security of the full protocol. In contrast, our methodology does not provide *completeness*: it may very well be the case that a particular choice of sub-protocol abstracts too much and ends up allowing attacks that are impossible to turn into attacks on the full protocol. Indeed, sometimes sieving through false positives helped us to refine our choice of sub-protocols.

Our analysis of exemplary protocols shows that many designs fail to appropriately address downgrade security. We advocate incorporating downgrade security as an integral part of security models of key exchange protocols.

We believe that analyzing the downgrade security of typical sub-protocols is within reach of automated tools. Symbolic analysis tools like ProVerif [17], Scyther [19] and Tamarin [38] seem particularly well suited to detect attacks on sub-protocols, helping analysts to find attacks against the full protocol or converge toward a sub-protocol that rules out false positives. Computationally sound tools like CryptoVerif [16], on the other hand, provide a means to prove the downgrade security of sub-protocols and, provided the sub-protocol is a sound abstraction of the full protocol, conclude that the corresponding full protocol also enjoys downgrade security. Finding a simulator that witnesses the correctness of a sub-protocol appears to be a more difficult task that may require ingenuity. While this may be out of reach of fully automated tools, interactive proofs can be constructed and machine-checked with tools like e.g. EasyCrypt [6]. Exploring the use of automated tools could increase the confidence in our proofs of downgrade security, and perhaps find other simpler or more practical attacks on protocols for which we only showed negative results.

References

- [1] B. Aboba, L. Blunk, J. Vollbrecht, J. Carlson, and H. Levkowetz. Extensible Authentication Protocol (EAP). IETF RFC 3748, 2004.
- [2] D. Adrian, K. Bhargavan, Z. Durumeric, P. Gaudry, M. Green, J. A. Halderman, N. Heninger, D. Springall, E. Thomé, L. Valenta, B. VanderSloot, E. Wustrow, S. Zanella-Béguelin, and P. Zimmermann. Imperfect forward secrecy: How Diffie-Hellman fails in practice. In *22nd ACM Conference on Computer and Communications Security*, pages 5–17, 2015.
- [3] N. J. AlFardan and K. G. Paterson. Lucky thirteen: Breaking the TLS and DTLS record protocols. In *2013 IEEE Symposium on Security and Privacy*, pages 526–540, 2013.
- [4] N. J. AlFardan, D. J. Bernstein, K. G. Paterson, B. Poettering, and J. C. N. Schuldt. On the security of RC4 in TLS. In *22th USENIX Security Symposium*, pages 305–320, 2013. ISBN 978-1-931971-03-4.
- [5] R. Barnes, M. Thomson, A. Pironti, and A. Langley. Deprecating Secure Sockets Layer Version 3.0. IETF RFC 7568, 2015.
- [6] G. Barthe, B. Grégoire, S. Heraud, and S. Zanella-Béguelin. Computer-aided security proofs for the working cryptographer. In *Advances in Cryptology, CRYPTO 2011*, pages 71–90, 2011.
- [7] G. Barthe, J. M. Crespo, Y. Lakhnech, and B. Schmidt. Mind the gap: Modular machine-checked proofs of one-round key exchange protocols. In *Advances in Cryptology – EUROCRYPT 2015*, pages 689–718, 2015.
- [8] M. Bellare and P. Rogaway. Entity authentication and key distribution. In *CRYPTO 1993*, pages 232–249, 1993.
- [9] F. Bergsma, B. Dowling, F. Kohlar, J. Schwenk, and D. Stebila. Multi-ciphersuite security of the secure shell (SSH) protocol. In *21st ACM Conference on Computer and Communications Security*, pages 369–381, 2014.
- [10] B. Beurdouche, K. Bhargavan, A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, A. Pironti, P.-Y. Strub, and J. Zinzindohoue. A Messy State of the Union: Taming the Composite State Machines of TLS. In *2015 IEEE Symposium on Security and Privacy*, pages 535–552, 2015.
- [11] K. Bhargavan and G. Leurent. Transcript Collision Attacks: Breaking Authentication in TLS, IKE, and SSH. Unpublished manuscript, 2015.
- [12] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P. Strub. Implementing TLS with verified cryptographic security. In *2013 IEEE Symposium on Security and Privacy*, pages 445–459, 2013.
- [13] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, P. Strub, and S. Zanella-Béguelin. Proving the TLS handshake secure (as it is). In *CRYPTO 2014*, pages 235–255, 2014.
- [14] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, P.-Y. Strub, and S. Zanella-Béguelin. Proving the tls handshake secure (as it is). Cryptology ePrint Archive, Report 2014/182, 2014. <http://eprint.iacr.org/2014/182/>.

- [15] S. Blake-Wilson, M. Nystrom, D. Hopwood, J. Mikkelsen, and T. Wright. Transport Layer Security (TLS) Extensions. IETF RFC 3546, 2003.
- [16] Blanchet. A computationally sound mechanized prover for security protocols. *Dependable and Secure Computing, IEEE Transactions on*, 5(4):193–207, 2008.
- [17] B. Blanchet. An efficient cryptographic protocol verifier based on prolog rules. In *14th IEEE Computer Security Foundations Workshop, CSFW 2014*, pages 82–96, 2001.
- [18] R. Canetti and H. Krawczyk. Security analysis of IKE’s signature-based key-exchange protocol. In *CRYPTO 2002*, pages 143–161, 2002.
- [19] C. Cremers. The Scyther tool: Verification, falsification, and analysis of security protocols. In *20th International Conference on Computer Aided Verification, CAV 2008*, pages 414–418, 2008.
- [20] C. J. F. Cremers. Key exchange in ipsec revisited: Formal analysis of ikev1 and ikev2. In *Computer Security - ESORICS 2011 - 16th European Symposium on Research in Computer Security, Leuven, Belgium, September 12-14, 2011. Proceedings*, pages 315–334, 2011.
- [21] A. Delignat-Lavaud and K. Bhargavan. Network-based origin confusion attacks against HTTPS virtual hosting. In *24th International Conference on World Wide Web, WWW 2015*, pages 227–237, 2015.
- [22] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. Internet Draft, 2015.
- [23] B. Dowling and D. Stebila. Modelling ciphersuite and version negotiation in the TLS protocol. In *20th Australasian Conference on Information Security and Privacy, ACISP 2015*, pages 270–288, 2015. doi: 10.1007/978-3-319-19962-7_16.
- [24] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In *22nd ACM Conference on Computer and Communications Security*, pages 1197–1210, 2015.
- [25] T. Duong and J. Rizzo. Here come the \oplus ninjas. Available at http://nerdoholic.org/uploads/dergln/beast_part2/ssl_jun21.pdf, May 2011.
- [26] C. Garman, K. G. Paterson, and T. V. der Merwe. Attacks only get better: Password recovery attacks against RC4 in TLS. In *24th USENIX Security Symposium*, pages 113–128, 2015. ISBN 978-1-931971-232.
- [27] D. Harkins and D. Carrel. The internet key exchange (IKE). RFC 2409, 1998. URL <http://www.ietf.org/rfc/rfc2409.txt>.
- [28] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DHE in the standard model. In *CRYPTO 2012*, pages 273–293, 2012.
- [29] M. Just and S. Vaudenay. Authenticated multi-party key agreement. In *ASIACRYPT 1996*, pages 36–49, 1996.
- [30] C. Kaufman, P. Hoffman, Y. Nir, and P. Eronen. Internet Key Exchange Protocol Version 2 (IKEv2). RFC 5996, 2010. URL <http://www.ietf.org/rfc/rfc5996.txt>.

- [31] K. Kobara, S. Shin, and M. Streffer. Partnership in key exchange protocols. In *2009 ACM Symposium on Information, Computer and Communications Security, ASIACCS 2009*, pages 161–170, 2009.
- [32] H. Krawczyk. SIGMA: the 'SIGn-and-MAc' approach to authenticated Diffie-Hellman and its use in the IKE protocols. In *CRYPTO 2003*, pages 400–425, 2003.
- [33] H. Krawczyk, K. G. Paterson, and H. Wee. On the security of the TLS protocol: A systematic analysis. In *CRYPTO 2013*, pages 429–448, 2013.
- [34] N. M. Langley, A. and B. Moeller. Transport Layer Security (TLS) False Start. Internet Draft, 2010.
- [35] R. H. M. Salter, E. Rescorla. Suite B Profile for Transport Layer Security (TLS). IETF RFC 5430, 2009.
- [36] N. Mavrogiannopoulos, F. Vercauteren, V. Velichkov, and B. Preneel. A cross-protocol attack on the TLS protocol. In *19th ACM Conference on Computer and Communications Security*, pages 62–72, 2012. ISBN 978-1-4503-1651-4. doi: 10.1145/2382196.2382206.
- [37] C. Meadows. Analysis of the internet key exchange protocol using the NRL protocol analyzer. In *1999 IEEE Symposium on Security and Privacy, Oakland, California, USA, May 9-12, 1999*, pages 216–231, 1999.
- [38] S. Meier, B. Schmidt, C. Cremers, and D. Basin. The tamarin prover for the symbolic analysis of security protocols. In *25th International Conference on Computer Aided Verification, CAV 2013*, pages 696–701, 2013.
- [39] B. Moeller and A. Langley. TLS Fallback Signaling Cipher Suite Value (SCSV) for Preventing Protocol Downgrade Attacks. IETF RFC 7507, 2015.
- [40] B. Möller, T. Duong, and K. Kotowicz. This POODLE Bites: Exploiting The SSL 3.0 Fallback. Available at <https://www.openssl.org/~bodo/ssl-poodle.pdf>, 2014.
- [41] J. Rizzo and T. Duong. The CRIME Attack, September 2012.
- [42] S. Turner and T. Polk. Prohibiting Secure Sockets Layer (SSL version 2.0). IETF RFC 6176, 2011.
- [43] University of Michigan. Tracking the FREAK Attack. Available at <https://freakattack.com/>, November 2015.
- [44] D. Wagner and B. Schneier. Analysis of the SSL 3.0 protocol. In *2nd USENIX Workshop on Electronic Commerce, WOEC 1996*, pages 29–40, 1996.
- [45] T. Ylonen and C. Lonvick. The secure shell (SSH) authentication protocol. RFC 4252, 2006.
- [46] T. Ylonen and C. Lonvick. The secure shell (SSH) transport layer protocol. RFC 4253, 2006.
- [47] P. Zimmermann. RFC 6189bis: ZRTP: Media Path Key Agreement for Unicast Secure RTP. Available at <http://zfone.com/docs/ietf/rfc6189bis.html>, November 2012.