

Multi-Ciphersuite Security of the Secure Shell (SSH) Protocol

Florian Bergsma
Horst Görtz Institute
Ruhr-Universität Bochum
Bochum, Germany
florian.bergsma@rub.de

Benjamin Dowling
Electrical Engineering and
Computer Science
Queensland University of
Technology
Brisbane, Australia
b1.dowling@qut.edu.au

Florian Kohlar
Horst Görtz Institute
Ruhr-Universität Bochum
Bochum, Germany
florian.kohlar@rub.de

Jörg Schwenk
Horst Görtz Institute
Ruhr-Universität Bochum
Bochum, Germany
joerg.schwenk@rub.de

Douglas Stebila
Electrical Engineering and
Computer Science
Mathematical Sciences
Queensland University of
Technology
Brisbane, Australia
stebila@qut.edu.au

ABSTRACT

The Secure Shell (SSH) protocol is widely used to provide secure remote access to servers, making it among the most important security protocols on the Internet. We show that the signed-Diffie–Hellman SSH ciphersuites of the SSH protocol are secure: each is a secure authenticated and confidential channel establishment (ACCE) protocol, the same security definition now used to describe the security of Transport Layer Security (TLS) ciphersuites.

While the ACCE definition suffices to describe the security of individual ciphersuites, it does not cover the case where parties use the same long-term key with many different ciphersuites: it is common in practice for the server to use the same signing key with both finite field and elliptic curve Diffie–Hellman, for example. While TLS is vulnerable to attack in this case, we show that SSH is secure even when the same signing key is used across multiple ciphersuites. We introduce a new generic multi-ciphersuite composition framework to achieve this result in a black-box way.

Categories and Subject Descriptors

C.2.0 [Computer–Communication Networks]: General
—security and protection

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS'14, November 3–7, 2014, Scottsdale, Arizona, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2957-6/14/11 ...\$15.00.

<http://dx.doi.org/10.1145/2660267.2660286>.

Keywords

Secure Shell (SSH); key agility; cross-protocol security; multi-ciphersuite; authenticated and confidential channel establishment

1. INTRODUCTION

Communication on the Internet is protected by a variety of cryptographic protocols: while the Transport Layer Security (TLS) protocol (also known as the Secure Sockets Layer (SSL) protocol) secures web communication, as well as e-mail transfer and many other network protocols, the Secure Shell (SSH) protocol¹ provides secure remote login and rudimentary virtual private network (VPN) access. It is of paramount importance to have strong cryptographic assurances of these protocols.

These and other real-world protocols tend to be far more complex than protocols typically studied in the academic literature. These protocols include both key exchange and secure channel communication, support negotiation of many combinations of cryptographic algorithms and a variety of authentication modes, and have additional functionality such as renegotiation and error reporting. All of these can affect the practical and theoretical security of the protocol.

At a high level, the parties run a cryptographic protocol to establish a secure channel, then communicate arbitrary application data over that channel. More precisely, execution begins with a channel establishment phase, in which parties negotiate which set of cryptographic parameters they intend to use, establish a shared session key, use long-term keys for entity authentication (either server-only or mutual), and send key confirmation messages. This is followed by the communication of application data over a secure channel which provides confidentiality and integrity using the session key from the channel establishment phase. The secure channel is called the *binary packet protocol* in SSH. A complicating factor for SSH (as well as TLS) is that some portions

¹In this paper, we refer exclusively to SSHv2 [39, 37, 40].

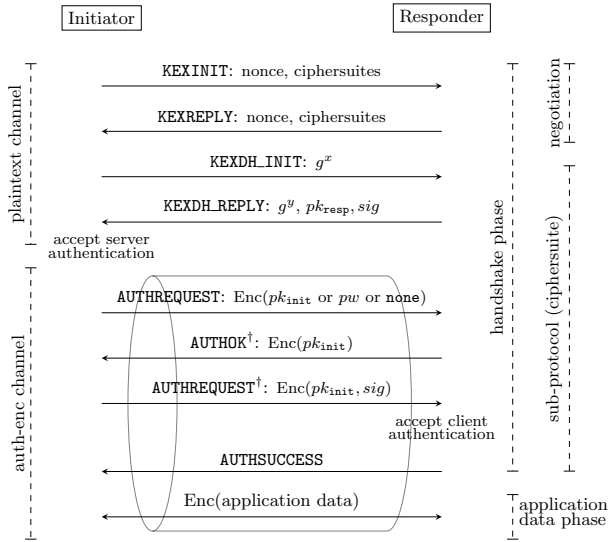


Figure 1: Overview of SSH protocol flow.

† denotes messages omitted for server-only/password auth.

of the channel establishment phase take place in plaintext, and other portions are sent over the secure channel. The overlap between the channel establishment phase and the secure channel can cause complications in the analysis of these protocols.

For precision, we will use the following terminology:

- *plaintext channel*: communication that is not sent via authenticated encryption using the session key;
- *auth-enc channel*: communication that is sent via authenticated encryption using the session key;
- *handshake phase*: communication of protocol messages to perform entity authentication and establish a secure channel, consisting of a negotiation phase and a sub-protocol (or *ciphersuite*) phase;²
- *application data phase*: communication of application data using the auth-enc channel.

Figure 1 shows a simplified version of the SSH protocol with mutual authentication; details appear in Section 4.

Provable security of real-world protocols. Standard *authenticated key exchange (AKE)* models [7, 13, 27] are not appropriate for modelling protocols such as SSH and TLS for several reasons. First, the auth-enc channel for secure application data communication is quite important but is not included in AKE definitions. Moreover, even the handshake phase cannot be analyzed as an AKE protocol: AKE security requires indistinguishability of session keys, but in both SSH and TLS, in the handshake phase, a key confirmation message is sent over the auth-enc channel which allows an attacker to distinguish a random session key from the real one. Some work has shown that truncated forms of the SSH [36] and TLS [22, 29] handshakes are secure AKE protocols, but this does not necessarily imply security of the entire protocol.

²We note that *ciphersuite* happens to be a TLS-centric term. SSH does not define a single ciphersuite, instead separately negotiating key exchange, encryption, and MAC algorithms. For consistency, in the case of SSH we refer to a single combination of these algorithms as a ciphersuite.

It has also been observed that standard notions of *authenticated encryption* are not quite appropriate for the auth-enc channels in SSH or TLS either. The security property that the auth-enc channel in SSH aims to meet is *buffered stateful authenticated encryption* [6, 1, 30], which includes confidentiality and integrity of ciphertexts and protection against reordering, along with details associated with byte-wise processing of received data.

Analysis of TLS proceeded in a similarly separate manner, until, in 2012, the first security proof of a full, unmodified TLS ciphersuite appeared. Jager et al. [20] showed that (mutually authenticated) signed-Diffie–Hellman TLS ciphersuites were secure *authenticated and confidential channel establishment (ACCE)* protocols under reasonable assumptions on the cryptographic building blocks. ACCE essentially combines AKE and authenticated encryption notions to obtain a single notion in which parties establish a channel that provides confidentiality and integrity of ciphertexts. Subsequently, ciphersuites based on RSA key transport and static Diffie–Hellman, with mutual and server-only authentication, have been shown ACCE secure by both Kohlar et al. [24] and Krawczyk et al. [26]. The ACCE notion was extended by Giesen et al. [17] to cover renegotiation, in which parties can establish a new ciphersuite or change authentication credentials in an existing connection. Alternative approaches for proving the full security of TLS include a composability approach [11] and formal verification of security properties of an implementation [9], but ACCE seems the dominant approach at present, and thus our choice for analyzing SSH.

Multi-protocol security. As noted above, both SSH and TLS support the negotiation of different combinations of cryptographic algorithms—ciphersuites—for both the handshake phase and the auth-enc channel. SSH’s possible negotiated algorithms are noted in Section 4, and TLS supports more than 300 different combinations of algorithms. A note on terminology: we will talk about SSH or TLS as a single “protocol” consisting of different “ciphersuites”; hence we are interested in “multi-ciphersuite” security.

The previous works on ACCE security of TLS all focus on ciphersuites running in isolation: in a cryptographic sense, each ciphersuite is a different “protocol”. Most ciphersuites of TLS have been proven secure, but only in a world where they have no interaction with other ciphersuites. In practice, servers and clients often share a single long-term key across multiple ciphersuites: For example, in SSH, the server may have a single 2048-bit RSA signing key that it uses with various key exchange and authenticated encryption mechanisms.

As first identified by Kelsey et al. [23], re-use or sharing of keys across multiple primitives or protocols can potentially be insecure; this is variously called a *chosen protocol attack*, *cross-protocol attack*, or *multi-protocol attack*. Very early work on SSL by Wagner and Schneier [35] identified a theoretical cross-ciphersuite attack on TLS: in ciphersuites with signed key exchange, the data structure that is signed (**ServerKeyExchange**) does not contain an identifier of its type, so it is theoretically possible that a data structure signed for one key exchange method could be interpreted as valid in another key exchange method. While Wagner and Schneier were not able to translate this into a concrete attack, Mavrogiannopoulos et al. [28] were able to make use of this observation to interpret a set of ECDH parameters as valid DH parameters. Cross-protocols attacks have been studied in a variety of contexts for protocols in the literature [2, 34,

[15] and in practice [21]; notably, Cremers [15] studied 30 AKE protocols from the literature and found cross-protocol attacks on 23 of them. In these lines of work, attacks arose from a common fundamental principle: messages signed or decrypted using long-term keys did not have sufficiently different structure to prevent misuse in other protocols.

There have been several works considering the joint security of protocols with shared or re-used keys, sometimes called *key agility*.³ In their original paper on chosen protocol attacks, Kelsey et al. [23] state five design principles that aim to render chosen protocol attacks impossible; Canetti et al. [14] similarly discuss requirements for security in multi-protocol environments. Thayer-Fabrega et al. [33] proposed the use of *strand spaces*, a type of formal logic for protocol execution, to identify under which conditions a protocol could be composed with other protocols (re-using the same long-term public key) without compromising security; enhancements to this approach have followed [18, 5]. Datta et al. [16] and Andova et al. [4] both give an alternative protocol composition logic. A common characteristic to these approaches is defining some form of *independence* of protocols, and then using a composition theorem where protocols that are secure in isolation and which are independent remain secure when used together, even with re-used long-term keys. Bhargavan et al. [10] analyze TLS in a multi-ciphersuite setting, constructing a generic protocol where some—but not all—algorithms can be combined while sharing long-term keys.

Contributions. Our main contribution is a provable security analysis of the SSH protocol. In particular, we show the various signed-Diffie–Hellman ciphersuites of SSH are ACCE secure in isolation, under reasonable assumptions on the underlying cryptographic primitive. We also show, using a newly created framework for analyzing the security of multi-ciphersuite protocols, that SSH is secure even when these ciphersuites share the same long-term key. Our multi-ciphersuite ACCE framework can be applied to analyze the security of other ACCE protocols.

1. *Provable security of signed-Diffie–Hellman SSH ciphersuites in isolation.* We provide the first proof that SSH is ACCE-secure. In particular, we show that the signed-Diffie–Hellman ciphersuites in SSH are ACCE-secure, under reasonable assumptions on the cryptographic primitives used. (Although RSA-key-transport-based ciphersuites have been standardized for SSH [19], OpenSSH, the most prominent implementation of SSH, does not support them as of this writing⁴, so we omit them.) We give results for both server-only and mutually authenticated variants.

For mutual authentication, we only provide a formal treatment of client authentication using public keys. While SSH does support client authentication using passwords [37, §8], this is a non-cryptographic form of password authentication: after establishing a server-to-client auth-enc channel, the client simply sends her username and password directly over the auth-enc channel. Thus, having analyzed the server-only variant, there is no value in further analyzing the case of password authentication. Note as well that SSH allows multiple connections to be multiplexed in a single encrypted

³Contrast this with the universal composability (UC) framework [12], where secure AKE protocols [13] can be composed with other protocols but long-term keys are in general not re-used across functionalities.

⁴<http://www.openssh.org/cgi-bin/cvsweb/src/usr.bin/ssh/kex.h?rev=1.64>

tunnel [38], but from a cryptographic perspective this is all just application data.

2. *Framework for analyzing multi-ciphersuite protocols.* We begin by adapting Jager et al.’s authenticated and confidential channel establishment (ACCE) definition [20]: we define a multi-ciphersuite ACCE protocol: a short negotiation phase is used to agree on one of several ciphersuites, which is then used in the subsequent handshake phase and auth-enc channel. We next define what it means for a multi-ciphersuite ACCE protocol to be secure: it should be hard to break authentication or channel security in any ciphersuite.

We then develop in Section 6 a generic approach for proving multi-ciphersuite security from single ciphersuite security. It will not be possible to prove in general that, if individual ciphersuites are ACCE-secure in isolation, then the collection is multi-ciphersuite-secure even when long-term keys are re-used across ciphersuites: the aforementioned attack by Mavrogianopoulos et al. [28] on the signed-DH and signed-ECDH ciphersuites in TLS serves as a counterexample to such a theorem, so we need some additional alteration to the standard ACCE definition.

Moreover, when long-term keys are shared, there are challenges in the standard simulation approach to proof. For example, consider the case of two different ciphersuites that use the same long-term keys for authentication. A standard simulation approach to proving multi-ciphersuite security would be to assume one ciphersuite is secure in isolation, then simulate the other ciphersuite. However, if long-term keys are shared between the two ciphersuites, then it is in general not possible to simulate the long-term private key operations in the second, simulated ciphersuite, because those keys are internal to the first ciphersuite.

These are the main problems our technical approach must solve. We achieve a composition theorem as follows:

1. Define a variant of ACCE in which the adversary has access to an *auxiliary oracle* that does operations using the long-term secret key, as long as queries to that oracle do not violate a certain condition.
2. Suppose for each ciphersuite SP_i there exists an auxiliary algorithm $Aux_i(sk, \cdot)$ and condition Φ_i such that:
 - (a) SP_i is secure even if an adversary makes queries to $Aux_i(sk, \cdot)$, provided the queries do not violate Φ_i (i.e., in the sense of item 1 above); and
 - (b) if SP_j shares long-term keys with SP_i , then SP_j can be simulated using Aux_i without violating Φ_i .
3. Then the collection of ciphersuites is secure, even when long-term keys are re-used across ciphersuites.

Item 1 can be viewed as “opening up” the ACCE definition a little bit, providing access to the secret key to do operations that “don’t affect security”. With carefully chosen auxiliary algorithms and conditions, items 2(a) and 2(b) work together to bypass the aforementioned challenge in proving a composition theorem using a simulation argument. Our approach seems to provide substantial compositional power without making proofs much harder in practice.

Our multi-ciphersuite ACCE approach contrasts with the key agility methodology of Bhargavan et al. [10] for analyzing TLS. As noted above, TLS is not multi-ciphersuite secure in general due to the cross-ciphersuite attack [28], so Bhargavan et al. develop a more “fine-grained” approach to key agility in TLS: they explicitly model TLS as a protocol with multiple signature, KEM, and PRF algorithms,

and then prove the joint security of key-agile TLS under reasonable assumptions on the individual building blocks. Our approach is more “coarse-grained”: we can compose several whole ACCE-secure ciphersuites in a nearly black-box manner, and the ciphersuites to be composed need not be as “cleanly” related to each other as in Bhargavan et al.. In fact, one could conceivably prove that key re-use in entirely unrelated protocols (e.g., the same signing key in SSH and (a revised form of) TLS) is secure using our framework.

3. *Multi-ciphersuite security of SSH.* Our composition framework can be readily applied to signed-Diffie–Hellman ciphersuites in SSH, yielding multi-ciphersuite security even when long-term signing keys are re-used across ciphersuites. To do so, we describe how to instantiate the auxiliary oracle Aux_i and predicate Φ_i in a way that maintains security in condition 2(a) above, yet still allows cross-protocol simulation as per condition 2(b) above. The composition theorem then immediately yields multi-ciphersuite security.

2. PRELIMINARIES

In this section, we define notation used in the paper and review the cryptographic assumptions used in the proofs.

Notation. Different typefaces are used to represent different types of objects: Algorithms (also \mathcal{A} and \mathcal{B}); **Queries**; **Protocols**; *variables*; security-notions; **constants**; vector notation \vec{x} is used for ordered lists. We use \emptyset to denote the empty string, and $[n] = [1, n] = \{1, \dots, n\} \subset \mathbb{N}$ for the set of integers between 1 and n . If A is a set, then $a \stackrel{\$}{\leftarrow} A$ denotes that a is drawn uniformly at random from A . If \mathcal{A} is a probabilistic algorithm, then $x \stackrel{\$}{\leftarrow} \mathcal{A}(y)$ denotes the output x of \mathcal{A} when run on input y and randomly chosen coins.

Standard security notions. We define in the standard way (omitted for space) the advantages of an algorithm \mathcal{A} in solving the decisional Diffie–Hellman (DDH) problem in a group of prime order q generated by g ($\text{Adv}_{g,q}^{\text{ddh}}(\mathcal{A})$), finding collisions in an unkeyed hash function H ($\text{Adv}_H^{\text{coll}}(\mathcal{A})$) [31], breaking existential unforgeability of a signature scheme SIG under chosen message attack ($\text{Adv}_{\text{SIG}}^{\text{euf-cma}}(\mathcal{A})$), and distinguishing a pseudorandom function F from random ($\text{Adv}_F^{\text{prf}}(\mathcal{A})$) [26, full version, p. 43–45].

Buffered stateful authenticated encryption. Paterson et al. [30] introduced *buffered stateful authenticated encryption (BSAE)* for appropriately modeling the security of the SSH auth-enc channel. The main difference of BSAE to previous definitions for authenticated encryption schemes is that the *decryption* oracle buffers *partial* ciphertexts until a complete ciphertext block is received, before answering a decryption query. A full definition of BSAE and the security game can be found in the full version [8].

3. MULTI-CIPHERSUITE ACCE

In the original ACCE formulation, an ACCE protocol is defined implicitly by however the experiment responds to the **Send** queries. In the multi-ciphersuite setting, there are many different ciphersuite algorithms to consider, so we begin by more formally defining a multi-ciphersuite protocol in several portions. There will be a *negotiation protocol*, which is common to all ciphersuites, and which is typically used to negotiate which ciphersuite is used. Each party then proceeds with the negotiated one of several *sub-protocols*, each of which represents a different ciphersuite. Each execution of

the protocol is called a *session* and will maintain and update a collection of *per-session variables*.

DEFINITION 1 (PER-SESSION VARIABLES). Let π denote the following collection of per-session variables:

- $\rho \in \{\text{init}, \text{resp}\}$: The party’s role in this session.
- $c \in \{1, \dots, n_{\text{SP}}, \perp\}$: The identifier of the sub-protocol chosen for this session, or \perp .
- $\text{pid} \in \{1, \dots, n_P, \perp\}$: The identifier of the alleged peer of this session, or \perp for an unauthenticated peer.
- $\alpha \in \{\text{in-progress}, \text{reject}, \text{accept}\}$: The status.
- k : A session key, or \perp . Note that k consists of two sub-keys: bi-directional authenticated encryption keys k_e and k_d (which themselves may consist of encryption and MAC sub-keys).
- sid : A session identifier defined by the protocol.
- st_e, st_d : State for the stateful authenticated encryption and decryption algorithms.
- Any additional state specific to the protocol.
- Any additional state specific to the security experiment.

We can now define an ACCE protocol. It will be convenient to explicitly name the different algorithms that are executed at different times in the protocol.

DEFINITION 2 (ACCE PROTOCOL). An ACCE protocol is a tuple of algorithms. The key generation algorithm $\text{KeyGen}() \stackrel{\$}{\rightarrow} (sk, pk)$ outputs a long-term secret key / public key pair. The handshake algorithms AlgI_ℓ and AlgR_ℓ , $\ell = 1, \dots$, take as input (sk, pk) and an incoming message m , update per-session variables π , and output an outgoing message m' . The handshake algorithms eventually set the variables for the peer identifier $\pi.\text{pid}$, the session status $\pi.\alpha$, the session key $\pi.k$, and the session identifier $\pi.\text{sid}$. There are also stateful authenticated encryption and decryption algorithms $\text{Enc}(\pi.k_e, m, \pi.\text{st}_e) \stackrel{\$}{\rightarrow} (C, \pi.\text{st}_e)$ and $\text{Dec}(\pi.k_d, C, \pi.\text{st}_d) \rightarrow (m', \pi.\text{st}_d)$. All algorithms are assumed to take as implicit input any global protocol parameters, including the list of all trusted peer public keys.

Having defined a single ACCE protocol, we now turn to the multi-ciphersuite setting.

DEFINITION 3 (MULTI-CIPHERSUITE PROTOCOL). A multi-ciphersuite ACCE protocol $\text{NP} \parallel \vec{\text{SP}}$ is the protocol obtained by first running a negotiation protocol NP , which outputs per-session variables π and a ciphersuite choice c , then running subprotocol $\text{SP}_c \in \vec{\text{SP}}$. A negotiation protocol NP is a tuple of algorithms, denoted either $\text{NP}.\text{AlgI}_\ell$ or $\text{NP}.\text{AlgR}_\ell$ for initiator or responder algorithms, respectively, for $\ell = 1, \dots$. All algorithms take as input an incoming message m , update per-session variables π , and output an outgoing message m' . The first algorithms for both the initiator and responder also take as input a vector \vec{s}_P of ciphersuite preferences that the party should use in this session. The final negotiation algorithm for both parties sets the ciphersuite choice variable $\pi.c$. Each sub-protocol SP_c is a tuple of algorithms corresponding to an ACCE protocol as in Definition 2, namely $\text{SP}_c.\text{KeyGen}, \text{SP}_c.\text{AlgI}_\ell, \text{SP}_c.\text{AlgR}_\ell, \text{SP}_c.\text{Enc}, \text{SP}_c.\text{Dec}$. Note that the execution of the negotiation protocol and the chosen subprotocol may be slightly interleaved, in that the responder may send the last negotiation message and the first sub-protocol message together.

3.1 Execution environment

The security experiment for a multi-ciphersuite ACCE protocol is similar to that of individual ACCE protocols [20], except that parties initially establish multiple long-term keys, the adversary can activate parties with an ordered list of sub-protocols, and the encryption/decryption is buffered stateful authenticated encryption, rather than a stateful length-hiding authenticated encryption. Let $\text{NP}\|\overline{\text{SP}}$ be a multi-ciphersuite ACCE protocol, with $|\overline{\text{SP}}| = n_{\text{SP}}$.

Parties and long-term key generation. The execution environment consists of n_P parties, P_1, \dots, P_{n_P} , each of whom is a potential protocol participant. At the beginning of the experiment, the variable $\delta_{i,\{c,d\}}$ is set to 1 or 0 and represents whether party P_i re-uses the same long-term key for SP_c and SP_d ; note that $\delta_{i,\{c,d\}}$ must be 0 if $\text{SP}_c.\text{KeyGen} \neq \text{SP}_d.\text{KeyGen}$, namely if there exists at least one input on which the two algorithms differ (for the same randomness). Observe that $\delta_{i,\{c,d\}}$ is symmetric in c and d . Each party P_i generates long-term private key / public key pairs $(sk_{i,c}, pk_{i,c})$ for each sub-protocol SP_c using $\text{SP}_c.\text{KeyGen}()$, but, for all $d > c$ such that $\delta_{i,\{c,d\}} = 1$, sets $(sk_{i,d}, pk_{i,d}) = (sk_{i,c}, pk_{i,c})$. We say that there is *no key re-use* if all $\delta_{i,\{c,d\}} = 0$.

Sessions. Each party can execute multiple sessions of the protocol, either concurrently or subsequently. We will denote the s th session of a protocol at party P_i by π_i^s , where $s \in \{1, \dots, n_S\}$. We overload the notation so that π_i^s also denotes the per-session variables π for this session. Each session within a party has read access to the party's long-term keys. The per-session variables $\pi_i^s.(c, pid, \alpha, k, sid)$ are initialized to $(\perp, \perp, \text{in-progress}, \perp, \perp)$. For the purposes of defining ciphertext indistinguishability and integrity, each session upon initialization chooses a uniform random bit $\pi_i^s.b \xleftarrow{\$} \{0, 1\}$. Each session also maintains additional variables for stateful encryption/decryption as required in Figure 2.

Adversary interaction. The adversary controls all communications between parties: it directs parties to initiate sessions, delivers messages to parties, and can reorder, alter, delete, and create messages. The adversary can also compromise certain long-term and per-session values of parties. The adversary interacts with parties using the following queries.

The first query models normal, unencrypted communication of parties during session establishment.

- **Send** $(i, s, m) \xrightarrow{\$} m'$: The adversary sends message m to session π_i^s . Party P_i processes message m according to the protocol specification and its per-session state π_i^s , updates its per-session state, and optionally outputs an outgoing message m' .

There is a distinguished initialization message which allows the adversary to activate the session with certain information. In particular, the initialization message consists of: the role ρ the party is meant to play in this session; the ordered list \vec{sp} of sub-protocols the party should use in this session; and optionally the identity pid of the intended partner of this session.

This query may return error symbol \perp if the session has entered state $\alpha = \text{accept}$ and no more protocol messages are transmitted over the unencrypted channel.

The next two queries model adversarial compromise of long-term and per-session secrets.

- **Reveal** $(i, s) \rightarrow k$: Returns session key $\pi_i^s.k$.
- **Corrupt** $(i, c) \rightarrow sk$: Returns party P_i 's long-term secret key $sk_{i,c}$ for sub-protocol c . Note the adversary does not

Encrypt (i, s, m_0, m_1) :

1. $u \leftarrow u + 1$
2. $(C^{(0)}, st_e^0) \xleftarrow{\$} \text{SP}_c.\text{Enc}(k_e, m_0, st_e^0)$
3. $(C^{(1)}, st_e^1) \xleftarrow{\$} \text{SP}_c.\text{Enc}(k_e, m_1, st_e^1)$
4. **if** $C^{(0)} = \perp$ **or** $C^{(1)} = \perp$ **then**
5. **return** \perp
6. $C[u] \leftarrow C^{(b)}$
7. **return** $C^{(b)}$

Decrypt (i, s, C) :

1. $(j, t) \leftarrow \pi_i^s.pid, v \leftarrow v + 1$
2. $(m, st_d) \leftarrow \text{SP}_c.\text{Dec}(k_d, C, st_d)$
3. **if** $m = \perp_p$ **then return** \perp
4. **if** $b = 0$ **then return** \perp
5. **if** $v > \pi_j^t.u$ **or** $C \neq \pi_j^t.C[v]$ **then**
6. **phase** $\leftarrow 1$
7. **if** **phase** = 1 **then return** m
8. **return** \perp

Figure 2: Encrypt and Decrypt queries in the multi-ciphersuite ACCE security experiment.

Note that $b, c, C[], k_d, k_e, st_d, st_e, u, v$ denote the values stored in the per-session variables π_i^s . Although $\pi_i^s.pid$ only contains the party identifier j , once π_i^s has accepted every session π_j^t has a unique matching session π_j^t known to the challenger. The Decrypt query accounts for buffering in the third line; this is the difference from ACCE's original stateful length-hiding definition [20, 26].

take control of the corrupted party, but can impersonate P_i in later sessions of sub-protocol c .

The final two queries model communication over the encrypted channel. The adversary can cause plaintexts to be encrypted as outgoing ciphertexts, and can cause ciphertexts to be delivered and decrypted as incoming plaintexts.

- **Encrypt** $(i, s, m_0, m_1) \xrightarrow{\$} C$: This query takes as input two messages m_0 and m_1 . If $\pi_i^s.k = \perp$, the query returns \perp . Otherwise, it proceeds as in Figure 2, depending on the random bit $\pi_i^s.b$ sampled by π_i^s at the beginning of the game and the state variables of π_i^s .
- **Decrypt** $(i, s, C) \rightarrow m$ **or** \perp : This query takes as input a ciphertext C . If $\pi_i^s.k = \perp$, the query returns \perp . Otherwise, it proceeds as in Figure 2. Note in particular that decryption can be buffered, meaning a decryption state may be maintained containing unprocessed bytes of a partial ciphertext.

Together, these two oracles model the BSAE notion, which simultaneously captures (i) indistinguishability under chosen ciphertext attack, (ii) integrity of ciphertexts, and (iii) buffered in-order delivery of ciphertexts. The hidden bit $\pi_i^s.b$ is leaked to the adversary if any of these goals is violated.

3.2 Security definitions

Security of ACCE protocols is defined by requiring that (i) the protocol is a secure authentication protocol, and (ii) the encrypted channel provides authenticated and confidential communication in the sense of buffered stateful authenticated encryption. In the multi-ciphersuite setting, security is further augmented by requiring that the parties agree on the sub-protocol used.

Multi-ciphersuite ACCE security experiment. The security experiment is played between an adversary \mathcal{A} and a challenger who implements all parties according to the multi-ciphersuite ACCE execution environment. The adversary sets the values of the long-term key re-use variables $\delta_{i,\{c,d\}}$. After the challenger initializes long-term keys based on $\delta_{i,\{c,d\}}$, the adversary receives the long-term public keys of all parties, then interacts with the challenger using **Send**, **Reveal**, **Corrupt**, **Encrypt**, and **Decrypt** queries. Finally, the adversary outputs a triple (i, s, b') and terminates. We begin by defining when sessions match.

DEFINITION 4 (MATCHING SESSIONS). *We say that session π_j^t matches π_i^s if*

- $\pi_i^s.\rho \neq \pi_j^t.\rho$;

- $\pi_i^s.c = \pi_j^t.c$; and
- $\pi_i^s.sid$ prefix-matches $\pi_j^t.sid$, meaning that (i) if π_i^s sent the last message in $\pi_i^s.sid$, then $\pi_j^t.sid$ is a prefix of $\pi_i^s.sid$, or (ii) if π_j^t sent the last message in $\pi_i^s.sid$, then $\pi_i^s.sid = \pi_j^t.sid$.

Note that for SSH, session IDs consist of a single value and thus not only prefix-match, but must be identical: $\pi_i^s.sid = \pi_j^t.sid$. Thus the “matching” relation is symmetric and thus easier to handle.

Next we give mutual and server-only authentication definitions, based on the existence of matching sessions. For server-only authentication, we are only concerned about clients accepting without a matching server session.

DEFINITION 5 (AUTHENTICATION). *Let π_i^s be a session. We say that π_i^s accepts maliciously for sub-protocol c^* if*

- $\pi_i^s.\alpha = \text{accept}$;
- $\pi_i^s.c = c^*$; and
- $\pi_i^s.pid = j \neq \perp$, where no $\text{Corrupt}(j, c^*)$ query was issued before π_i^s accepted, nor $\text{Corrupt}(j, d)$ for any d such that $\delta_{j, \{c^*, d\}} = 1$,

but there is no unique session π_j^t which matches π_i^s .

Define $\text{Adv}_{\text{NP} \parallel \overline{\text{SP}}, c^*}^{\text{mcs-acce-auth}}(\mathcal{A})$ as the probability that, when \mathcal{A} terminates in the multi-ciphersuite ACCE experiment for $\text{NP} \parallel \overline{\text{SP}}$, there exists a session that has accepted maliciously for sub-protocol c^* .

Define $\text{Adv}_{\text{NP} \parallel \overline{\text{SP}}, c^*}^{\text{mcs-acce-so-auth}}(\mathcal{A})$ as the probability that, when \mathcal{A} terminates in the multi-ciphersuite ACCE experiment for $\text{NP} \parallel \overline{\text{SP}}$, there exists an initiator session (i.e., with $\pi_i^s.\rho = \text{init}$) that has accepted maliciously for sub-protocol c^* .

Channel security is defined by the ability to break confidentiality or integrity of the channel. Formally, this is defined as the ability of the adversary to guess the bit b used in the **Encrypt** and **Decrypt** queries of an uncompromised session. “Uncompromised” means that the adversary did not reveal the session key at either the session or any matching session, and that that adversary did not corrupt the long-term keys of either party in the session. We give variants for mutually and server-only authenticated channels.

DEFINITION 6 (CHANNEL SECURITY).

Suppose \mathcal{A} outputs (i, s, b') in the multi-ciphersuite ACCE experiment. We say that \mathcal{A} answers the encryption challenge correctly for subprotocol c^ if*

- $\pi_i^s.\alpha = \text{accept}$;
- $\pi_i^s.c = c^*$;
- no $\text{Corrupt}(i, c^*)$ query was ever issued, nor $\text{Corrupt}(i, d)$ for any d such that $\delta_{i, \{c^*, d\}} = 1$;
- no $\text{Corrupt}(j, c^*)$ query was ever issued for any j such that π_j^t matches π_i^s , nor $\text{Corrupt}(j, d)$ for any d such that $\delta_{j, \{c^*, d\}} = 1$;
- no **Reveal** (i, s) query was issued;
- no **Reveal** (j, t) query was issued for any π_j^t that matches π_i^s ; and
- $\pi_i^s.b = b'$.

Define $\text{Adv}_{\text{NP} \parallel \overline{\text{SP}}, c^*}^{\text{mcs-acce-so-aenc}}(\mathcal{A})$ as $|p - 1/2|$, where p is the probability that \mathcal{A} answers the encryption challenge correctly for subprotocol c^* .

Define $\text{Adv}_{\text{NP} \parallel \overline{\text{SP}}, c^*}^{\text{mcs-acce-aenc}}(\mathcal{A})$ as $|p - 1/2|$, where p is the probability that \mathcal{A} answers the encryption challenge correctly

for subprotocol c^* and either $\pi_i^s.\rho = \text{init}$ or both $\pi_i^s.\rho = \text{resp}$ and there exists a session that matches π_i^s .

DEFINITION 7 (MULTI-CIPHERSUITE-ACCE-SECURE).

A multi-ciphersuite protocol $\text{NP} \parallel \overline{\text{SP}}$ is $\bar{\epsilon}$ -multi-ciphersuite-ACCE-secure against an adversary \mathcal{A} if, for all c^ , we have that $\text{Adv}_{\text{NP} \parallel \overline{\text{SP}}, c^*}^{\text{mcs-acce-auth}}(\mathcal{A}) \leq \epsilon_{c^*}$ and $\text{Adv}_{\text{NP} \parallel \overline{\text{SP}}, c^*}^{\text{mcs-acce-aenc}}(\mathcal{A}) \leq \epsilon_{c^*}$. We define an analogous notion for server-only authentication.*

When $n_{\text{SP}} = 1$, the multi-ciphersuite ACCE protocol and security definitions are equivalent to the original ACCE definitions (albeit with slightly different notation), except for the change to buffered stateful authenticated encryption.

4. THE SSH PROTOCOL

In this section, we describe the SSH protocol using signed Diffie–Hellman.

There are several cryptographic components that may be negotiated in SSH, and the collective choice of these components constitutes a ciphersuite. A party’s preferences are represented as a vector \overline{sp} , and the initiator and responder preferences $\overline{sp}_C, \overline{sp}_S$ are inputs to the negotiation function $\text{neg}(\overline{sp}_C, \overline{sp}_S) \rightarrow c$ specified by the standard [40, §7.1] which selects the first element in \overline{sp}_C that is also in \overline{sp}_S .

Each ciphersuite SSH_c can use different cryptographic components. The signature scheme SIG_c for server and client authentication may be either RSA, DSA, ECDSA [32], or Ed25519. The key exchange method is Diffie–Hellman over either a finite field or elliptic curve cyclic group G_c of prime order q_c generated by g_c . The hash function H_c can be either SHA-1 or SHA-256. The buffered stateful encryption scheme StE_c can be composed of a variety of encryption and MAC algorithms, including TripleDES in CBC mode or AES in CBC or CTR mode and HMAC with MD5, SHA-1, SHA-256, or SHA-512; or ChaCha20 with Poly1305.

During the negotiation phase, **KEXINIT** and **KEXREPLY** exchange nonces and negotiate the ciphersuite. During the key-exchange portion of the sub-protocol phase, **KEXDH_INIT** and **KEXDH_REPLY** exchange key-material, generate session keys and authenticate the responder to the initiator via the negotiated digital certificates and ciphersuites. During the authentication portion of the sub-protocol phase, the responder verifies if the chosen authentication mode is authorised for the given initiator, and authenticates the initiator via passwords, public-keys or no client authentication at all.

The basic outline of the SSH protocol is given in Figure 1 in the introduction; the detailed message flow and processing for the signed-Diffie–Hellman handshake phase with server-only or mutual public key authentication can be found in Figure 3. For details on the authenticated encryption we refer to the standard [40] and Albrecht et al. [1].

4.1 The SSH PRF

The PRF_c function described in Figure 4 is used in the SSH protocol to compute two values: H , which will be used as the session ID (this value is later signed in the **KEXDH_REPLY** and **AUTHREPLY** messages); and $k_1 \parallel k_2 \parallel k_3 \parallel k_4 \parallel k_5 \parallel k_6$ (which are later used as encryption keys, IVs, and authentication keys). PRF_c computes these values using the hash function H_c negotiated by the ciphersuite. While PRF_c is superficially similar to HMAC, it varies sufficiently that it merits independent analysis.

We cannot prove security for SSH from the assumption that H_c is a collision-resistant hash function: in SSH the

Negotiation

1. **init** \rightarrow **resp**: KEXINIT
 1. $r_C \xleftarrow{\$} \{0, 1\}^{\mu=128}$
 2. send KEXINIT $\leftarrow (r_C, \bar{s}p_C)$
 3. $\pi.\rho \leftarrow \text{init}$
 4. $\pi.\alpha \leftarrow \text{in-progress}$
2. **resp** \rightarrow **init**: KEXREPLY
 1. $r_S \xleftarrow{\$} \{0, 1\}^{\mu}$
 2. send KEXREPLY $\leftarrow (r_S, \bar{s}p_S)$
 3. $\pi.\rho \leftarrow \text{resp}$
 4. $\pi.\alpha \leftarrow \text{in-progress}$
 5. $\pi.c \leftarrow \text{neg}(\bar{s}p_C, \bar{s}p_S)$
3. **init**
 1. $\pi.c \leftarrow \text{neg}(\bar{s}p_C, \bar{s}p_S)$

Signed-Diffie–Hellman sub-protocol (common to all authentication modes)

4. **init** \rightarrow **resp**: KEXDH_INIT
 1. $x \xleftarrow{\$} \mathbb{Z}_{q_{\pi.c}}$
 2. $e \leftarrow g_{\pi.c}^x$
 3. send KEXDH_INIT $\leftarrow e$
5. **resp** \rightarrow **init**: KEXDH_REPLY and NEWKEYS
 1. $y \xleftarrow{\$} \mathbb{Z}_{q_{\pi.c}}$
 2. $f \leftarrow g_{\pi.c}^y$
 3. $K \leftarrow e^y$
 4. $(\pi.sid, \pi.k) \leftarrow \text{PRF}_{\pi.c}(K, V_C \| V_S \| \text{KEXINIT} \| \text{KEXREPLY} \| pk_{S, \pi.c} \| e \| f)$
 5. $\sigma_S \leftarrow \text{SIG}_{\pi.c}.\text{Sign}(sk_{S, \pi.c}, \pi.sid)$
 6. send KEXDH_REPLY $\leftarrow (f, pk_{S, \pi.c}, \sigma_S)$
 7. send NEWKEYS
6. **init** \rightarrow **resp**: NEWKEYS
 1. $K \leftarrow f^x$
 2. $(\pi.sid, \pi.k) \leftarrow \text{PRF}_{\pi.c}(K, V_C \| V_S \| \text{KEXINIT} \| \text{KEXREPLY} \| pk_{S, \pi.c} \| e \| f)$
 3. **if** $\text{SIG}_{\pi.c}.\text{Vfy}(pk_{S, \pi.c}, \sigma_S, \pi.sid) = 0$ **then**
 4. $\pi.\alpha \leftarrow \text{reject}$ and terminate
 5. $\pi.pid \leftarrow S$, where P_S has public key $pk_{S, \pi.c}$
 6. send NEWKEYS

Note V_C and V_S are client and server version strings.

Signed-Diffie–Hellman sub-protocol, server-only authentication mode

7. **init** \rightarrow **resp**: AUTHREQUEST
 1. send AUTHREQUEST $\leftarrow \text{username} \| \text{service} \| \text{service} \| \text{none}$
8. **resp** \rightarrow **init**: AUTHSUCCESS or AUTHFAILURE
 1. **if** **none** authentication is authorised for *username* for *service* **then**
 2. $\pi.\alpha \leftarrow \text{accept}$; send AUTHSUCCESS
 3. **else**
 4. $\pi.\alpha \leftarrow \text{reject}$; send AUTHFAILURE
11. **init**
 1. **if** AUTHFAILURE **then**
 2. $\pi.\alpha \leftarrow \text{reject}$ and terminate
 3. **else if** AUTHSUCCESS **then**
 4. $\pi.\alpha \leftarrow \text{accept}$

Signed-Diffie–Hellman sub-protocol, mutual authentication mode

7. **init** \rightarrow **resp**: AUTHREQUEST
 1. send AUTHREQUEST $\leftarrow \text{username} \| \text{service} \| \text{public-key} \| 0 \| \text{alg} \| pk_{C, \pi.c}$ (where *alg* is the name of the public key algorithm (RSA, DSA, ECDSA) and $pk_{C, \pi.c}$ is the client's public key for this ciphersuite)
8. **resp** \rightarrow **init**: AUTHOK or AUTHFAILURE
 1. **if** *username* is not allowed access to *service* by public-key authentication **then**
 2. $\pi.\alpha \leftarrow \text{reject}$ and terminate
 3. **if** $\pi.\alpha = \text{in-progress}$ **then**
 4. send AUTHOK $\leftarrow \text{alg} \| pk_{C, \pi.c}$
 5. **if** $\pi.\alpha = \text{reject}$ **then**
 6. send AUTHFAILURE and terminate
9. **init** \rightarrow **resp**: AUTHREQUEST
 1. $A \leftarrow \text{username} \| \text{service} \| \text{public-key} \| 1 \| \text{alg} \| pk_{C, \pi.c}$
 2. $\sigma_C \leftarrow \text{SIG}_{\pi.c}.\text{Sign}(sk_{C, \pi.c}, \pi.sid, A)$
 3. send AUTHREQUEST $\leftarrow A \| \sigma_C$
10. **resp** \rightarrow **init**: AUTHSUCCESS or AUTHFAILURE
 1. $A' \leftarrow \text{username} \| \text{service} \| \text{public-key} \| 1 \| \text{alg} \| pk_{C, \pi.c}$
 2. **if** $A' \neq A$ **then**
 3. $\pi.\alpha \leftarrow \text{reject}$
 4. **if** $\text{SIG}_{\pi.c}.\text{Vfy}(pk_{C, \pi.c}, \sigma_C, \pi.sid, A) = 0$ **then**
 5. $\pi.\alpha \leftarrow \text{reject}$
 6. **if** $\pi.\alpha = \text{in-progress}$ **then**
 7. $\pi.\alpha \leftarrow \text{accept}$
 8. **if** $\pi.\alpha = \text{accept}$ **then**
 9. send AUTHSUCCESS
 10. **else if** $\pi.\alpha = \text{reject}$ **then**
 11. send AUTHFAILURE and terminate
11. **init**
 1. **if** AUTHFAILURE **then**
 2. $\pi.\alpha \leftarrow \text{reject}$ and terminate
 3. **else if** AUTHSUCCESS **then**
 4. $\pi.\alpha \leftarrow \text{accept}$

Figure 3: SSH handshake phase protocol: negotiation protocol and signed-Diffie–Hellman sub-protocol

$\text{PRF}_c(K, x)$:

1. $H \leftarrow H_c(x \| K)$
2. $\text{label} \leftarrow [A, B, C, D, E, F]$
3. **for** $i \in \{1, \dots, 6\}$ **do**
4. $k_i \leftarrow H_c(K \| H \| \text{label}_i \| H)$
5. **return** $(H, k_1 \| k_2 \| k_3 \| k_4 \| k_5 \| k_6)$

Figure 4: Computation of PRF_c using H_c .

hash value H to be signed by both parties not only contains a transcript of the most important exchanged messages, but also the secret Diffie-Hellman key K computed by both parties. If H_c leaks information about K , the protocol cannot be proven secure. Therefore we need the assumption that PRF_c is a secure PRF, which is how our security proof of SSH proceeds in the rest of this section.

Under the assumption that H_c is a random function, it is straightforward to see that PRF_c is a secure PRF.

Analysis of PRF_c under weaker, standard-model assumptions on H_c is more challenging. One way of analyzing key derivation functions is Krawczyk's extract-then-expand paradigm [25]. In this paradigm, first a pseudorandom key $K \leftarrow H(SK_M)$ is *extracted* from the secret key material (such as the Diffie–Hellman shared secret) SK_M using a hash function H , then application keying material $KM \leftarrow F(K, \text{"1"} \| \text{info}) \| F(K, \text{"2"} \| \text{info}) \| \dots$ is *expanded* from the pseudorandom key K using a PRF F . Although PRF_c

does seemingly have an extract phase (line 1) and then an expand phase (line 4), the extract-then-expand paradigm does not directly apply because the pseudorandom key (H , in the case of PRF_c) is subsequently used in another area of the SSH protocol: H is signed by the signature scheme and the signature is transmitted over the channel. Thus H and the signature on H must not leak anything about the Diffie–Hellman shared secret.

It may be possible to adapt extract-then-expand to analyze the SSH PRF, but we leave that as future work. Our main security proof of SSH is entirely standard model, so any future work improving the analysis of PRF_c from random oracle model to standard model immediately yields a full standard model proof of SSH.

5. ACCE SECURITY OF SSH

In this section, we analyze the security of single signed-DH SSH ciphersuites, in isolation. We first note a few challenges we faced in the proofs, then show authentication and channel security in the server-only and mutual authentication modes.

5.1 Challenges with security proofs for SSH

ACCE. As noted in the Introduction, challenges are often encountered when trying to analyze real-world protocols. The first problem that arises when analyzing SSH is the fact that the messages needed for client authentication are

sent encrypted, allowing the adversary to trivially win key indistinguishability in a standard authenticated key exchange security experiment; this is resolved by switching to ACCE.

Session IDs vs. matching conversations. A secure authentication protocol can, loosely speaking, be defined as a protocol where the success probabilities of active and passive adversaries are equal, up to a negligible difference. There are two main possible formalizations of this concept: session IDs and matching conversations. We initially tried to base our proof that SSH is a secure authentication protocol on the classical notion of *matching conversations* (when the two parties have the same transcript of communication), in order to make our result comparable to previous work. However, SSH itself makes this impossible, because of a special option to negotiate keys more quickly: the SSH client may choose to start an abbreviated handshake, by guessing which cryptographic parameters the server would accept, sending messages KEXINIT and KEXDH_INIT simultaneously. The SSH server however may refuse this option, and in this case, the KEXDH_INIT message is discarded, and replaced by a new message KEXDH_INIT' to the server. In such a scenario, an adversary may simply change the original value of KEXDH_INIT arbitrarily, thus breaking the matching conversations condition, and nevertheless make both sessions accept. Instead, the SSH specification itself suggests the use of a protocol-specific session ID, a hash value H over the initial handshake messages. This hash value is then used to generate and verify the signatures both on client and server side.

PRF-ODH. Readers may wonder why we do not need the PRF-Oracle-Diffie–Hellman (PFR-ODH) assumption used in the analysis of signed-DH in TLS [20, 26]. In TLS, a hypothetical adversary who can solve the CDH problem can make a client oracle accept maliciously by intercepting all messages after `ServerKeyExchange`, and then faking a valid `ServerFinished` message. At the same time, this adversary can refuse to cooperate in breaking DDH by testing if the DDH challenge was embedded in the current session. All this may not happen in SSH: here, acceptance directly depends on signature verification.

5.2 Server-only-authentication mode

In this section we drop subscripts for ciphersuites: SSH denotes a single ciphersuite of the signed-Diffie–Hellman SSH protocol described in Section 4, with signature scheme SIG, Diffie–Hellman group of prime order q generated by g , and hash function H , and the BSAE scheme StE.

The following theorem shows that, if the hash function H is collision-resistant, the signature scheme SIG is euf-cma-secure, the DDH problem for (g, q) is hard, the PRF is a secure PRF, and the symmetric encryption is a secure BSAE scheme, then the (single ciphersuite) signed-Diffie–Hellman SSH protocol is a secure server-only ACCE protocol.

THEOREM 1 (SSH IS SERVER-ONLY-ACCE-SECURE).

Let μ be the length of the nonces in KEXINIT and KEXREPLY ($\mu = 128$), n_P the number of participating parties and n_S the maximum number of sessions per party. The algorithms $\mathcal{B}_1, \dots, \mathcal{B}_5$ given in the proof of the theorem are such that, for all algorithms \mathcal{A} , $\text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A}) \leq \frac{(n_P n_S)^2}{2^\mu} + \text{Adv}_{\text{H}}^{\text{cr}}(\mathcal{B}_1^{\mathcal{A}}) + n_P \text{Adv}_{\text{SIG}}^{\text{euf-cma}}(\mathcal{B}_2^{\mathcal{A}})$ and $\text{Adv}_{\text{SSH}}^{\text{acce-so-aenc}}(\mathcal{A}) \leq \text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A}) + n_P n_S (\text{Adv}_{g,q}^{\text{ddh}}(\mathcal{B}_3^{\mathcal{A}}) + \text{Adv}_{\text{PRF}}^{\text{prf}}(\mathcal{B}_4^{\mathcal{A}}) + \text{Adv}_{\text{StE}}^{\text{bsae}}(\mathcal{B}_5^{\mathcal{A}}))$ and $\mathcal{B}_1^{\mathcal{A}}, \dots, \mathcal{B}_5^{\mathcal{A}}$ have approximately the same running time as \mathcal{A} .

In order to prove the theorem, we first obtain a bound on the server-only authentication advantage, then on the channel security advantage.

PROOF OF $\text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A})$ BOUND. The essence of the proof is the observation that acceptance of a client session is the result of a successful signature verification. To be able to use this fact, we have to make sure that all session IDs are different (by aborting if a nonce is chosen twice or if a collision occurs in the hash computation of the session ID).

Let $\text{break}_\delta^{(0)}$ be the event that occurs when a client session accepts maliciously in Game δ in the sense of Definition 5.

Game 0. The game equals the ACCE security experiment described in Section 3.2. Thus, $\text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A}) = \Pr(\text{break}_0^{(0)})$.

Game 1. In this game we add an abort rule for non-unique nonces r_i . Specifically the challenger collects a list L of all cookies r_i sampled by the challenger during the simulation. If one cookie appears twice, we abort the simulation. Thus $\Pr(\text{break}_0^{(0)}) \leq \Pr(\text{break}_1^{(0)}) + \frac{(n_P n_S)^2}{2^\mu}$.

Game 2. In this game we exclude hash collisions. Note that in this game we can compute all session keys and session identifiers honestly, and we maintain a list $Coll$, where all the input/output pairs of all executions of the hash function H are recorded. We abort if at any time a pair $(in, H(in))$ is added to $Coll$ such that there already exists an entry $(in', H(in'))$ in $Coll$ with $H(in) = H(in')$ but $in \neq in'$. Now we construct $\mathcal{B}_1^{\mathcal{A}}$ as follows: \mathcal{B}_1 simulates the SSH protocol and interacts with \mathcal{A} . Whenever \mathcal{A} wins the acce-so-auth game, \mathcal{B}_1 inspects the recorded simulation to see if a hash collision occurred. If it did, \mathcal{B}_1 outputs this collision. Since \mathcal{B}_1 finds a collision, we have that $\Pr(\text{break}_1^{(0)}) \leq \Pr(\text{break}_2^{(0)}) + \text{Adv}_{\text{H}}^{\text{cr}}(\mathcal{B}_1^{\mathcal{A}})$.

Game 3. In this game we exclude signature forgeries. We abort the simulation if some session $\pi_{i^*}^s$ accepts after it receives a signature which was never output of a session with a matching session identifier. Note that we have excluded nonce and hash collisions, so from now on all values to be signed are different. Thus any abort event is related to a signature forgery.

Technically, we construct an algorithm $\mathcal{B}_2^{\mathcal{A}}$ which simulates the SSH protocol as in Game 1. \mathcal{B}_2 interacts with \mathcal{A} . \mathcal{B}_2 receives a public key pk from an euf-cma signature challenger for SIG, guesses which public key pk_{j^*} the session will use to verify the signature (which costs us a factor n_P in the reduction) and sets $pk_{j^*} = pk$. Since the signing key has to be uncorrupted it is no problem for the reduction that the secret signing key is unknown. If \mathcal{B}_2 needs to sign a message on behalf of party P_{j^*} , it makes a signing query to the euf-cma challenger. If the session $\pi_{i^*}^s$ maliciously accepts in the sense of definition 5 in Game 3, we know from the discussion above that the maliciously accepting session has verified a signature σ' over a session ID H where there is no session $\pi_{i^*}^s$ with the same session ID, thus this signature was not generated with a call to the signature challenger. Thus \mathcal{B}_2 has found (H, σ') as a signature forgery, so $\Pr(\text{break}_2^{(0)}) \leq \Pr(\text{break}_3^{(0)}) + n_P \text{Adv}_{\text{SIG}}^{\text{euf-cma}}(\mathcal{B}_2^{\mathcal{A}})$.

Final analysis. Now all signatures are computed by legitimate parties only, and are all computed for different session IDs. Thus there is no way for a session to accept maliciously, and we have $\Pr(\text{break}_3^{(0)}) = 0$. \square

PROOF OF $\text{Adv}_{\text{SSH}}^{\text{acce-so-aenc}}(\mathcal{A})$ BOUND. Let $\text{break}_\delta^{(1)}$ be the event that occurs when \mathcal{A} answers the encryption challenge correctly in Game δ in the sense of Definition 6.

Game 0. This game equals the ACCE security experiment described in Section 3.2.

Game 1. This game is identical to Game 3 of the previous proof and we abort if some session accepts maliciously. With the previous sequence of games we ensured unique nonces, excluded hash collisions and signature forgeries. Thus, in this game any session that accepts non-maliciously in the sense of Definition 5 has a unique uncorrupted partner session. From the previous proof, we have $\Pr(\text{break}_0^{(1)}) \leq \Pr(\text{break}_1^{(1)}) + \text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A})$. From now on, we always have a matching session for the session $\pi_{j^*}^{s^*}$ where the adversary tries to guess the random bit: for server sessions through Definition 5, and for client sessions through this game.

Game 2. In this game, we guess the session for which the adversary outputs the bit b' . We guess two indices $(i^*, s^*) \in [n_P] \times [n_S]$ and abort if the adversary outputs (i, s, b') with $(i^*, s^*) \neq (i, s)$. This happens with probability $\frac{1}{n_P n_S}$. We then exploit that no client session maliciously accepts due to Game 1, so we have that there exists a unique partner session $\pi_{j^*}^{s^*}$ which can be easily determined by the simulator. Thus we have: $\Pr(\text{break}_1^{(1)}) \leq n_P n_S \cdot \Pr(\text{break}_2^{(1)})$.

Game 3. In this game we replace the value $K = g^{x^y}$ computed by $\pi_{i^*}^{s^*}$ and $\pi_{j^*}^{s^*}$ with a random value K^* . Since we have excluded maliciously accepting sessions, and since $\pi_{j^*}^{s^*}$ fulfills all conditions from Definition 6, the adversary cannot influence these values. Any adversary \mathcal{A} that can distinguish this game from the previous game can directly be used to construct an adversary \mathcal{B}_3^A that can break the DDH assumption: let (g, g^u, g^v, g^w) be the DDH challenge. We set $g^x := g^u$ and $g^y := g^v$, and $K^* := g^w$. If $w = uv$, then we have $K^* = K$, and we are in Game 2, otherwise we are in Game 3. Thus $\Pr(\text{break}_2^{(1)}) \leq \Pr(\text{break}_3^{(1)}) + \text{Adv}_{g,q}^{\text{ddh}}(\mathcal{B}_3^A)$.

Game 4. In this game we replace the values H, k_1, \dots, k_6 computed by $\pi_{i^*}^{s^*}$ and $\pi_{j^*}^{s^*}$ as $\text{PRF}(K^*, \text{sid})$ with random values H^*, k_1^*, \dots, k_6^* . Any adversary \mathcal{A} that can distinguish this game from the previous game can directly be used to construct an adversary \mathcal{B}_4^A that can break the PRF assumption: let $S = H \| k_1 \| \dots \| k_6$ be the output of PRF, and let $S^* = H^* \| k_1^* \| \dots \| k_6^*$ be a random string of the same length. For S we are in Game 3, and for S^* in Game 4. Thus $\Pr(\text{break}_3^{(1)}) \leq \Pr(\text{break}_4^{(1)}) + \text{Adv}_{\text{PRF}}^{\text{prf}}(\mathcal{B}_4^A)$.

Final analysis. We now have that the keys k_1^*, \dots, k_6^* are information-theoretically independent from the key exchange messages. Thus any adversary \mathcal{A} that can guess (i^*, s^*, b') correctly can directly be used to construct an adversary \mathcal{B}_5^A that breaks the BSAE scheme. Technically we exploit the fact that all keys for the encryption scheme are independent from the handshake and embed a BSAE challenger. Now we simply have to forward \mathcal{A} 's output to the challenger and thus we have $\Pr(\text{break}_4^{(1)}) \leq \text{Adv}_{\text{StE}}^{\text{bsae}}(\mathcal{B}_5^A)$. \square

Remark 1. Forward secrecy. The ACCE definition of Jager et al. [20] can be extended to include forward secrecy, meaning that the adversary in the channel security definition is allowed to corrupt the long-term key of the owner of the target session or its peer after the target session has accepted. We have omitted forward secrecy from this paper for simplicity, but Definition 6 can be easily extended to cover the case of forward secrecy, and the proof of the channel security bound can be readily adapted using the techniques in [20].

Remark 2. Mutual authentication mode. In a similar manner, it can be shown that the (single ciphersuite) signed-Diffie–Hellman SSH protocol has secure mutual authentication when the client uses public key authentication if the building blocks of SSH are secure, and thus is a secure ACCE protocol with mutual authentication.

6. COMPOSITION THEOREM FOR MULTI-CIPHERSUITE SECURITY

As noted in the Introduction, if two ciphersuites with the same long-term key generation algorithm have been proven individually secure (i.e., if $\text{SP}_1.\text{KeyGen} = \text{SP}_2.\text{KeyGen}$, $\text{NP} \parallel \text{SP}_1$ is ACCE-secure, and $\text{NP} \parallel \text{SP}_2$ is ACCE-secure), it does not necessarily follow that they are collectively secure when parties use the same long-term secret key in both ciphersuites.

We still hope however to be able to prove some security properties of individual ciphersuites separately and then compose them together using some generic theorem, rather than having to directly prove security of the whole multi-ciphersuite combination all at once. Some intuition for our composition framework follows.

Suppose a user supports two ACCE-secure ciphersuites (the “apple” ciphersuite and the “orange” ciphersuite) with authentication in both cases provided by use of the same digital signature scheme, and that in each ciphersuite, the signed data clearly and unambiguously identifies the ciphersuite (for example, starting with the word “apple” or the word “orange”, respectively). As well, suppose that during authentication in each ciphersuite, the receiver verifies that the signed data is for the ciphersuite in question (it really does start with the word “apple” or the word “orange”, respectively).

Intuitively, then, obtaining signatures from one ciphersuite should not help in breaking the second ciphersuite, even if they are both signed using the same long-term keys.

We are now able to consider the security of the two ciphersuites together. Since “apple” signatures will not affect the security of the “orange” ciphersuite, and “orange” signatures will not affect the security of the “apple” ciphersuite, the two ciphersuites remain secure even if they share long-term keys. A theorem for the security of the two ciphersuites together should say: if both the “apple” and “orange” ciphersuites are being used and users are possibly sharing long-term keys between them, and the adversary breaks some session in the “apple” ciphersuite, then the “apple” ciphersuite was not secure even in isolation; and similarly for “orange”.

To prove security, our simulator will be given a challenger for the “apple” ciphersuite and must simulate the others. The simulation can simulate ciphersuites that use keys not shared with “apple” because it can choose those keys itself. Only ciphersuites that share keys with “apple” are tricky; in this case, the simulator asks the “apple” challenger to sign an “orange” message, which should not affect the security of the “apple” ciphersuite but allows the simulator to simulate the “orange” ciphersuite. We have to introduce a few small technical conditions to ensure that the simulation goes through, but this is the main idea.

6.1 Single ciphersuite security with auxiliary oracle

We begin by “opening up” the ACCE security definition a little bit, to consider security of a single ciphersuite in isolation, but with additional access to secret key operations.

As shown in Definition 8, we extend the ACCE security experiment to allow the adversary access to an *auxiliary oracle* that runs a specified private key operation $\text{Aux}(sk, \cdot)$ (in the case of signed-DH SSH, a signing oracle that signs arbitrary messages). If the adversary breaks the original ACCE security goals without asking a query x to Aux that violates the *constraint* or *predicate* Φ , then the adversary wins. For example, if we are studying the “orange” ciphersuite, then the predicate $\Phi(x)$ would test if x starts with the word “orange”. As long as the adversary’s signing queries did not start with the word “orange”, they should not help him win the security experiment.

DEFINITION 8 (ACCE-SECURE W/AUXILIARY ORACLE). Let P be an ACCE protocol. Let $\text{Aux} : (sk, x) \mapsto y$ be an algorithm. Augment the ACCE experiment giving the adversary access to an additional oracle $\text{Aux}(i, x)$ which outputs $\text{Aux}(sk_i, x)$. Let $\Phi(x)$ be a predicate on a value x .

Define $\text{Adv}_{\mathsf{P}, \text{Aux}, \Phi}^{\text{acce-auth-aux}}(\mathcal{A})$ as the probability that, when \mathcal{A} terminates in the above augmented ACCE experiment for P with auxiliary oracle, there exists a session that has accepted maliciously, with the additional constraint that, for all x such that \mathcal{A} queried $\text{Aux}(\pi_i^s, \text{pid}, x)$, $\Phi(x) = 0$.

Similarly, define $\text{Adv}_{\mathsf{P}, \text{Aux}, \Phi}^{\text{acce-aenc-aux}}(\mathcal{A})$ as $|p - 1/2|$, where p is the probability that \mathcal{A} answers the encryption challenge correctly in the above augmented ACCE experiment for P with auxiliary oracle, again with the additional constraint that, for all x such that \mathcal{A} queried $\text{Aux}(\pi_i^s, \text{pid}, x)$, $\Phi(x) = 0$.

We define analogous notions for server-only authentication.

6.2 Multi-ciphersuite composition

Once we have that each ciphersuite is individually secure, we want to use a composition theorem to show that their multi-ciphersuite combination is secure, even if long-term keys are shared across ciphersuites. For ciphersuites that do not re-use long-term keys, security of the combination is trivial. For ciphersuites that do re-use long-term keys, reducing the security of the combination to the security of the individual ciphersuites requires that we be able to *simulate* the other ciphersuites. We can do so using the above auxiliary signing oracle, as long as we do not violate the predicate. For example, we need to be able to simulate the “apple” ciphersuite using the “orange” signing oracle, without asking queries that start with the word “orange”. This simulatability condition is modelled in Definitions 9 and 10. Our composition theorem (Theorem 2) is then shown using such a simulation argument.

DEFINITION 9 (SIMULATABLE). We say a sub-protocol SP is simulatable using auxiliary algorithm Aux and helper algorithms $\{\text{HI}_\ell, \text{HR}_\ell\}$ if, for all ℓ , $\text{HI}_\ell^{\text{Aux}(sk, \cdot)}(pk, \pi, m) = \text{SP.AlgI}_\ell(sk, pk, \pi, m)$ and $\text{HR}_\ell^{\text{Aux}(sk, \cdot)}(pk, \pi, m) = \text{SP.AlgR}_\ell(sk, pk, \pi, m)$.

DEFINITION 10 (FRESHLY SIMULATABLE). We say that auxiliary algorithm Aux and helper algorithms $\{\text{HI}_\ell, \text{HR}_\ell\}$ provide a fresh simulation of SP under condition Φ if Definition 9 is satisfied and, for all $A \in \{\text{HI}_\ell, \text{HR}_\ell\}$, there exist no inputs to A that cause A to make a call $\text{Aux}(\cdot, x)$ such that $\Phi(x) = 1$.

THEOREM 2 (MULTI-CIPHERSUITE COMPOSITION).

Let $\text{NP} \parallel \vec{\text{SP}}$ be a multi-ciphersuite ACCE protocol. Let $\vec{\text{Aux}}$ be a vector of auxiliary algorithms and let $\vec{\Phi}$ be a vector of conditions. Suppose that:

1. for all $c, d \in [n_{\text{SP}}]$, $d \neq c$, there exist helper algorithms $\{\text{HI}_\ell^{d,c}, \text{HR}_\ell^{d,c}\}$ such that Aux_c and these helper algorithms provide a fresh simulation of SP_d under Φ_c ; and
2. after observing the messages output by the negotiation protocol, one can efficiently reconstruct the complete per-session variables updated by those algorithms.

Then the algorithm \mathcal{B} explicitly given in the proof of the theorem is such that, for all algorithms \mathcal{A} and for all c , $\text{Adv}_{\text{NP} \parallel \vec{\text{SP}}, c}^{\text{mcs-acce-auth}}(\mathcal{A}) \leq n_{\text{SP}} \text{Adv}_{\text{NP} \parallel \vec{\text{SP}}, c, \text{Aux}_c, \Phi_c}^{\text{acce-auth-aux}}(\mathcal{B}^{\mathcal{A}})$ even under key re-use across ciphersuites. Moreover, $\mathcal{B}^{\mathcal{A}}$ has at most approximately the same running time as \mathcal{A} .

Similarly, $\text{Adv}_{\text{NP} \parallel \vec{\text{SP}}, c}^{\text{mcs-acce-aenc}}(\mathcal{A}) \leq n_{\text{SP}} \text{Adv}_{\text{NP} \parallel \vec{\text{SP}}, c, \text{Aux}_c, \Phi_c}^{\text{acce-aenc-aux}}(\mathcal{B}^{\mathcal{A}})$ for all c , even under key re-use across ciphersuites.

Moreover, analogous versions of the theorem apply for server-only authentication.

PROOF. We will specify an algorithm \mathcal{B} that uses \mathcal{A} . Whenever \mathcal{A} breaks authentication or channel security for ciphersuite c^* in the multi-ciphersuite ACCE experiment for multi-ciphersuite ACCE protocol $\text{NP} \parallel \vec{\text{SP}}$, the algorithm $\mathcal{B}^{\mathcal{A}}$ will, with probability $1/|\vec{\text{SP}}|$, break authentication or channel security in the ACCE with auxiliary oracle experiment for the ACCE protocol $\text{NP} \parallel \text{SP}_{c^*}$.

Let \mathcal{A} be an adversary in the multi-ciphersuite ACCE experiment. Recall that \mathcal{A} starts the experiment by setting key re-use variables $\delta_{i, \{c, d\}}$, which is 1 if party P_i is to re-use long-term keys between SP_c and SP_d , namely if $\text{SP}_c.\text{KeyGen} = \text{SP}_d.\text{KeyGen}$ and party P_i is to set $sk_{i,c} = sk_{i,d}$; $\delta_{i, \{c, d\}} = 0$ otherwise.

Algorithm \mathcal{B} simulates an multi-ciphersuite ACCE experiment for $\text{NP} \parallel \vec{\text{SP}}$ as follows. First, \mathcal{B} chooses $\hat{c} \xleftarrow{\$} \{1, \dots, n_{\text{SP}}\}$. \mathcal{B} interacts with a challenger for the ACCE experiment for $\text{NP} \parallel \text{SP}_{\hat{c}}$ with auxiliary oracle $\text{Aux}_{\hat{c}}$.

\mathcal{B} obtains the parties’ public keys for sub-protocol c from the $\text{NP} \parallel \text{SP}_{\hat{c}}$ experiment. For each party P_i and each sub-protocol d , if $\delta_{i, \{c, d\}} = 1$ then \mathcal{B} sets party P_i ’s public key for sub-protocol d equal to its public key in sub-protocol c , otherwise it generates a fresh key pair using $\text{SP}_d.\text{KeyGen}$. \mathcal{B} gives all of these public keys to \mathcal{A} .

\mathcal{B} now runs \mathcal{A} . \mathcal{A} can make any **Send**, **Corrupt**, **Reveal**, **Encrypt**, or **Decrypt** queries specified in the multi-ciphersuite ACCE experiment. \mathcal{B} needs to answer all of these. The basic idea of \mathcal{B} ’s simulation is as follows.

\mathcal{B} will start off every session by relaying it down to the challenger for the ACCE-security of $\text{NP} \parallel \text{SP}_{\hat{c}}$ with auxiliary oracle. If a session ends up negotiating sub-protocol c , then \mathcal{B} continues relaying all queries for that session to the $\text{NP} \parallel \text{SP}_{\hat{c}}$ challenger.

If a session ends up negotiating a sub-protocol d other than \hat{c} , \mathcal{B} needs to simulate it. It can do so as follows. By pre-condition 2 of the theorem, it can reconstruct the per-session variables used by the negotiation protocol in the challenger, so it can construct its own per-session variables from the output of the negotiation protocol. If the query is directed towards a party P_i such that P_i is using the same key for sub-protocols \hat{c} and d (i.e., if $\delta_{i, \{c, d\}} = 1$), then \mathcal{B} simulates the session for party P_i using the helper algorithms $\{\text{HI}_\ell^{d, \hat{c}}, \text{HR}_\ell^{d, \hat{c}}\}$ for SP_d using the auxiliary oracle Aux of the challenger; by pre-condition 1 of the theorem, this provides a correct simulation of $\text{NP} \parallel \text{SP}_d$. If $\delta_{i, \{c, d\}} = 0$, then \mathcal{B} can simulate the session for party P_i itself since it generated P_i ’s secret key for this sub-protocol.

For parties and sessions where \mathcal{B} relayed the complete session to the challenger, \mathcal{B} also relays the **Corrupt**, **Reveal**, **Encrypt**, and **Decrypt** queries to the challenger; otherwise \mathcal{B} answers them itself.

\mathcal{B} 's simulation of the multi-ciphersuite ACCE experiment for $\text{NP}\|\overline{\text{SP}}$ to \mathcal{A} is perfect.

Suppose \mathcal{A} breaks authentication in $\text{NP}\|\overline{\text{SP}}$. In particular, there exists in the multi-ciphersuite ACCE experiment some $c^* \in [n_{\text{SP}}]$ and some session π_i^s that has accepted maliciously for sub-protocol c^* with peer identifier j , but there is no unique session π_j^t (in the multi-ciphersuite ACCE experiment) with which π_i^s has a matching session. With probability $1/n_{\text{SP}}$, $\hat{c} = c^*$. In this case, there correspondingly exists in the ACCE-aux challenger a session π_i^s that has accepted with peer identifier j but there is no unique session π_j^t (in the ACCE-aux challenger) with which π_i^s has had a matching session. Note in particular that \mathcal{B} has not violated the condition Φ_{c^*} for $\text{NP}\|\text{SP}_{c^*}$ because SP_d is freshly simulatable under Φ_{c^*} due to pre-condition 1 of the theorem. Thus \mathcal{B} has caused a session in the ACCE-aux challenger to accept maliciously, and thus has broken authentication in $\text{NP}\|\text{SP}_{c^*}$. Hence, $\text{Adv}_{\text{NP}\|\overline{\text{SP}}, c^*}^{\text{mcs-acce-auth}}(\mathcal{A}) \leq n_{\text{SP}} \text{Adv}_{\text{NP}\|\text{SP}_{c^*}, \text{Aux}_{c^*}, \Phi_{c^*}}^{\text{acce-auth-aux}}(\mathcal{B}^{\mathcal{A}})$.

Similarly, if \mathcal{A} breaks channel security of $\text{NP}\|\overline{\text{SP}}$ by answering the encryption challenge correctly, then with probability $1/n_{\text{SP}}$ \mathcal{B} can answer its encryption challenge correctly and break the channel security of $\text{NP}\|\text{SP}_{c^*}$. Note that \mathcal{B} has not made any prohibited queries in the channel security definition: **Reveal** queries that would have made the ACCE challenger unrefresh also would have made the multi-ciphersuite ACCE experiment unrefresh; and similarly to the authentication case above, \mathcal{B} has not violated the condition Φ_{c^*} . Hence, $\text{Adv}_{\text{NP}\|\overline{\text{SP}}, c^*}^{\text{mcs-acce-aenc}}(\mathcal{A}) \leq n_{\text{SP}} \text{Adv}_{\text{NP}\|\text{SP}_{c^*}, \text{Aux}_{c^*}, \Phi_{c^*}}^{\text{acce-aenc-aux}}(\mathcal{B}^{\mathcal{A}})$.

This yields the result. Note the same reasoning yields the results for server-only authentication. \square

7. SSH IS MULTI-CIPHERSUITE SECURE

In order to use the composition theorem to show that signed-Diffie–Hellman SSH ciphersuites are multi-ciphersuite secure, even with re-use of long-term keys across ciphersuites, we need to define the auxiliary algorithm Aux and the condition Φ , show that the preconditions of Theorem 2 are satisfied, and show that individual ciphersuites are ACCE-secure with Aux .

Let SSH_c denote a ciphersuite of SSH, using signature scheme SIG_c . Recall from Section 4 that both the initiator and responder use the long-term signing key as follows. First, they compute the session ID as a hash of a session identification string and the shared secret:

$$\pi.\text{sid} \leftarrow \text{H}_c(V_C \| V_S \| \text{KEXINIT} \| \text{KEXREPLY} \| pk_d \| e \| f \| K) . \quad (1)$$

Finally, they compute a signature $\sigma \leftarrow \text{SIG}_c.\text{Sign}(sk, \pi.\text{sid})$. (If Sign is a hash-then-sign scheme, this means that the session identification string is hashed twice.) Recall further that KEXINIT and KEXREPLY contain the initiator and responder's respective preference-ordered list of ciphersuites. (These are actually separate lists for key exchange, compression, signature, MAC, and symmetric encryption algorithms, but from these we can infer a ciphersuite.)

We define the auxiliary algorithm $\text{Aux}_c(sk, x)$ as computing $\text{SIG}_c.\text{Sign}(sk, \text{H}_c(x))$. For a ciphersuite c , we define $\Phi_c(x) = 1$ if, when x is parsed as in (1) and the ordered ciphersuite preferences $\vec{s}p_C$ and $\vec{s}p_S$ are parsed from KEXINIT

and KEXREPLY , $c = \text{neg}(\vec{s}p_C, \vec{s}p_S)$; in other words, if c is the ciphersuite that is mutually most preferred by the initiator and responder.

7.1 Preconditions of composition theorem

It is straightforward to show that signed-DH SSH ciphersuites satisfy precondition 1 (freshly simulatable under the condition) and precondition 2 (reconstruction of per-session variables after negotiation) of Theorem 2. We sketch an argument of this; the formal version is omitted due to space.

Consider two different signed-DH SSH ciphersuites SSH_c and SSH_d . Precondition 1 can be seen by constructing helper algorithms $\{\text{HL}_\ell, \text{HR}_\ell\}$ that, with access to the above auxiliary algorithm Aux_c for SSH_c defined above, perfectly simulate the signed-DH ciphersuite and showing that these helper algorithms provide a fresh simulation of SSH_d under the above condition Φ_c . These helper algorithms are exactly the algorithms described in Section 4, except that constructions of signatures by hashing the session string and then using the signature scheme's Sign_c algorithm are replaced with calls to the auxiliary oracle Aux which implements the above algorithm Aux_c . Moreover, these helper algorithms for SSH_d do not violate condition Φ_c for SSH_c because Φ_c for SSH_c only outputs 1 if the input $V_c \| V_s \| \text{KEXINIT} \| \text{KEXREPLY} \| pk_d \| e \| f \| K$ to Aux_c is such that $\text{neg}(\vec{s}p_C, \vec{s}p_S) = c$. However, the helper algorithms will only make queries where the first mutual ciphersuite is $d \neq c$. Thus, the helper algorithms for SSH_d provide a fresh simulation of SSH_d using auxiliary algorithm Aux_c under condition Φ_c .

Precondition 2 requires that there exists an efficient algorithm that, after observing the messages output by the negotiation protocol, can reconstruct the complete per-session variables at that point in time. Notice in the description of SSH in Section 4 that the negotiation stage does not construct or use any secret values—at the end of negotiation, all variables in π are based either on the **Send** initialization values from \mathcal{A} or the contents of the plaintext messages KEXINIT or KEXREPLY .

7.2 Security of SSH with auxiliary oracle

THEOREM 3 (SSH IS SECURE W/AUX. ORACLE).

Let SSH_c be a signed-DH SSH ciphersuite with signature scheme SIG_c , hash function H_c ; define Aux_c and Φ_c as above. Let μ be the length of the nonces in KEXINIT and KEXREPLY ($\mu = 128$), n_P the number of participating parties and n_S the maximum number of sessions per party. The algorithms $\mathcal{B}_1, \dots, \mathcal{B}_5$ given in the proof of the lemma, are such that, for all algorithms \mathcal{A} , $\text{Adv}_{\text{SSH}_c, \text{Aux}_c, \Phi_c}^{\text{acce-so-auth-aux}}(\mathcal{A}) \leq \frac{(n_P n_S)^2}{2^\mu} + \text{Adv}_{\text{H}_c}^{\text{cr}}(\mathcal{B}_1^{\mathcal{A}}) + n_P \text{Adv}_{\text{SIG}_c}^{\text{uf-cma}}(\mathcal{B}_2^{\mathcal{A}})$, and $\text{Adv}_{\text{SSH}_c, \text{Aux}_c, \Phi_c}^{\text{acce-so-aenc-aux}}(\mathcal{A}) \leq \text{Adv}_{\text{SSH}_c, \text{Aux}_c, \Phi_c}^{\text{acce-so-auth-aux}}(\mathcal{A}) + n_P n_S \left(\text{Adv}_{g_c, q_c}^{\text{ddh}}(\mathcal{B}_3^{\mathcal{A}}) + \text{Adv}_{\text{PRF}_c}^{\text{prf}}(\mathcal{B}_4^{\mathcal{A}}) + \text{Adv}_{\text{StE}_c}^{\text{bsae}}(\mathcal{B}_5^{\mathcal{A}}) \right)$, and $\mathcal{B}_1^{\mathcal{A}}, \dots, \mathcal{B}_5^{\mathcal{A}}$ have approximately the same running time as \mathcal{A} .

We obtain a bound on the server-only authentication advantage; the bound on channel security proof proceeds identically to the bound on $\text{Adv}_{\text{SSH}}^{\text{acce-so-aenc}}(\mathcal{A})$ in Section 5.2.

PROOF OF $\text{Adv}_{\text{SSH}_c, \text{Aux}_c, \Phi_c}^{\text{acce-so-auth-aux}}(\mathcal{A})$ BOUND. Games 0, 1, and 2 proceed exactly as in the bound on $\text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A})$.

Game 0. The game equals the multi-ciphersuite ACCE security experiment described in Section 3.2.

Game 1. In this game we proceed identically to Game 1 in the bound on $\text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A})$, adding an abort rule for non-unique nonces, and get the same result.

Game 2. In the next two games we will exclude adversarial modifications of all messages (KEXINIT to KEXDH_INIT) by using a successful adversary to either output a hash collision (in this game) or a signature forgery (next game). In this game we proceed exactly as Game 2 in the bound on $\text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A})$, adding an abort rule for hash collisions, and get the same result.

Game 3. In this game we ensure an adversary cannot use signature forgery to make some session accept maliciously. If the session $\pi_{i^*}^s$ maliciously accepts in the sense of Definition 5, we know from the discussion in the proof for the bound on $\text{Adv}_{\text{SSH}}^{\text{acce-so-auth}}(\mathcal{A})$ that \mathcal{A} has modified at least one of the key exchange messages and computed a valid signature σ' over the hash of the correspondingly modified session string. In order to do this, either \mathcal{A} has computed a valid signature itself, or \mathcal{A} has utilised the auxiliary signing algorithm (for the negotiated ciphersuite c) Aux_c to compute a hash and signature on the modified session string. In order for the ACCE-with-auxiliary-oracle experiment to remain fresh, for all x that \mathcal{A} queries to Aux_c , we must have that $\Phi_c(x) = 0$; in particular, when x is parsed as a session string as given in equation (1), the negotiated ciphersuite $\text{neg}(\vec{s}p_C, \vec{s}p_S) \neq c$. But all sessions that accept have negotiated ciphersuite equal to c , and thus no query to the auxiliary oracle helps make any session accept maliciously. We now embed a euf-cma signature challenger, receive a public key pk , guess the public-key pk_{j^*} the oracle will use for signature verification, (again costing our reduction by a factor of n_P) and replace pk with pk_{j^*} . We know any maliciously accepting oracle has verified a signature σ' over a session string where there exists no other oracle $\pi_{j^*}^{t^*}$ with the same session string. Thus σ' was generated by the adversary, and we can forward (sid', σ') as a signature forgery to the euf-cma signature challenger, and we get: $\Pr(\text{break}^{(0)}) \leq \Pr(\text{break}_3^{(0)}) + n_P \text{Adv}_{\text{SIG}_c}^{\text{euf-cma}}(\mathcal{B}_2^{\mathcal{A}})$.

Final analysis. After Game 3, all of the server's relevant key-exchange messages are authenticated via the signature σ_S , and since Game 3 aborts when a session accepts maliciously, consequently we have $\Pr(\text{break}_3^{(0)}) = 0$. \square

7.3 Final result: Multi-ciphersuite SSH

Combining Theorem 3 from the previous subsection with the composition theorem (Theorem 2) immediately yields that the SSH protocol is multi-ciphersuite secure, even with key re-use across ciphersuites.

COROLLARY 1 (SSH IS MULTI-CIPHERSUITE SECURE).
Let SSH be the multi-ciphersuite SSH protocol with each of the n_{SP} ciphersuites SSH_c being a signed-Diffie-Hellman ciphersuite as in Section 4. The algorithms $\mathcal{B}_1, \dots, \mathcal{B}_5$ inferred from the proof are such that, for all algorithms \mathcal{A} :

$$\text{Adv}_{\text{SSH},c}^{\text{mcs-acce-so-auth}}(\mathcal{A}) \leq n_{\text{SP}} \left(\frac{(n_P n_S)^2}{2^\mu} + \text{Adv}_{\text{H}_c}^{\text{cr}}(\mathcal{B}_1^{\mathcal{A}}) + n_P \text{Adv}_{\text{SIG}_c}^{\text{euf-cma}}(\mathcal{B}_2^{\mathcal{A}}) \right),$$

$$\text{Adv}_{\text{SSH},c}^{\text{mcs-acce-so-aenc}}(\mathcal{A}) \leq \text{Adv}_{\text{SSH},c}^{\text{mcs-acce-so-auth}}(\mathcal{A}) + n_{\text{SP}} n_P n_S \cdot \left(\text{Adv}_{g_c, q_q}^{\text{ddh}}(\mathcal{B}_3^{\mathcal{A}}) + \text{Adv}_{\text{PRF}_c}^{\text{prf}}(\mathcal{B}_4^{\mathcal{A}}) + \text{Adv}_{\text{StE}_c}^{\text{bsae}}(\mathcal{B}_5^{\mathcal{A}}) \right),$$

and $\mathcal{B}_1^{\mathcal{A}}, \dots, \mathcal{B}_5^{\mathcal{A}}$ have approximately the same running time as \mathcal{A} . Moreover, analogous versions of the theorem apply for mutual authentication.

8. DISCUSSION

Although we encountered some challenges in proving the ACCE-security of SSH, overall SSH seemed somewhat easier to prove secure compared with the proofs of signed-DH ciphersuites in TLS [20]. As in both cases, mutual authentication comes from digital signatures, but in SSH the object that is signed is the (hash of the) session identifier which is a significant portion of the transcript, whereas in TLS only the ephemeral keys are signed. This means the security guarantees from verifying the signature in SSH more readily lead to the proof of entity authentication for the whole session.

Another nice consequence of how SSH uses signatures is that it enabled us to readily prove multi-ciphersuite security. Even though long-term signing keys may be used by multiple ciphersuites, in every case the object that is signed uniquely identifies the ciphersuite it is intended to be used for. This reinforces the importance of the long-held cryptographic wisdom of ‘signing what you mean to sign’. Our multi-ciphersuite composition framework precisely captures Anderson and Needham’s [3] principle 3 on protocol design: “Be careful when signing or decrypting data that you never let yourself be used as an oracle by your opponent.”

It is instructive to examine the TLS cross-ciphersuite attack in the context of our multi-ciphersuite framework: the full version of this paper [8] describes the attack of Mavrogiannopoulos et al. [28] in more detail and the conditions of our composition theorem that TLS does not satisfy. Mavrogiannopoulos et al. suggested including the ciphersuite and handshake transcripts in what is signed in TLS as a countermeasure. If future versions of TLS do indeed do this, for example moving the server signature to just before the server’s `Finished` message and including the complete transcript, then it should be straightforward to adapt existing security analyses of TLS. Moreover, such an adaptation should easily have a proof of single-ciphersuite security even with an auxiliary signing oracle, at which point our composition theorem can readily be applied to yield multi-ciphersuite security. With discussions for a new version of TLS beginning on the IETF’s mailing list, we hope that TLS 1.3 will indeed incorporate this suggestion.

Acknowledgements

The authors gratefully acknowledge helpful discussions with Tibor Jager. The research leading to these results has received funding from the European Community (FP7/2007-2013) under grant agreement number ICT-2007-216646—European Network of Excellence in Cryptology II (ECRYPT II), the Australian Technology Network—German Academic Exchange Service (ATN-DAAD) Joint Research Co-operation Scheme, and the Australian Research Council (ARC) Discovery Project scheme under grant DP130104304.

9. REFERENCES

- [1] M. R. Albrecht, K. G. Paterson, and G. J. Watson. Plaintext recovery attacks against SSH. In *2009 IEEE Symposium on Security and Privacy*, pages 16–26. IEEE Computer Society Press, May 2009.
- [2] J. Alves-Foss. Multi-protocol attacks and the public key infrastructure. In *Proc. 21st National Information Systems Security Conference*, pages 566–576, October 1998.

- [3] R. J. Anderson and R. M. Needham. Robustness principles for public key protocols. In D. Coppersmith, editor, *CRYPTO'95*, volume 963 of *LNCS*, pages 236–247. Springer, Aug. 1995.
- [4] S. Andova, C. Cremers, K. Gjøsteen, S. Mauw, S. F. Mjølsnes, and S. Radomirović. A framework for compositional verification of security protocols. *Information and Computation*, 206:425–459, 2008.
- [5] G. Bela and I. Ignat. Verifying the independence of security protocols. In *Proc. 2007 IEEE International Conference on Intelligent Computer Communication and Processing*, pages 155–162. IEEE, 2007.
- [6] M. Bellare, T. Kohno, and C. Namprempre. Breaking and provably repairing the SSH authenticated encryption scheme: A case study of the encode-then-encrypt-and-MAC paradigm. *ACM Transactions on Information and System Security*, 7(2):206–241, May 2004. Extended abstract published in *ACM CCS 2002*.
- [7] M. Bellare and P. Rogaway. Entity authentication and key distribution. In D. R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 232–249. Springer, Aug. 1993.
- [8] F. Bergsma, B. Dowling, F. Kohlar, J. Schwenk, and D. Stebila. Multi-ciphersuite security of the Secure Shell (SSH) protocol (full version). Cryptology ePrint Archive, Report 2013/813, 2014. <http://eprint.iacr.org/>.
- [9] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P.-Y. Strub. Implementing TLS with verified cryptographic security. In *2013 IEEE Symposium on Security and Privacy*, pages 445–459. IEEE Computer Society Press, May 2013.
- [10] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, P.-Y. Strub, and S. Zanella-Béguelin. Proving the TLS handshake secure (as it is). In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014*, volume 8617 of *LNCS*, pages 235–255. Springer, 2014.
- [11] C. Brzuska, M. Fischlin, N. P. Smart, B. Warinschi, and S. C. Williams. Less is more: Relaxed yet composable security notions for key exchange. *International Journal of Information Security*, 12(4):267–297, August 2013.
- [12] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, Oct. 2001.
- [13] R. Canetti and H. Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In B. Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, May 2001.
- [14] R. Canetti, C. Meadows, and P. Syverson. Environmental requirements for authentication protocols. In M. Okada, B. C. Pierce, A. Scedrov, H. Tokuda, and A. Yonezawa, editors, *Proc. Meet-NSF-JSPS Internaional Symposium on Software Security (ISSS) – Theories and Systems, Part 9*, volume 2609 of *LNCS*, pages 339–355. Springer, 2002.
- [15] C. J. F. Cremers. Feasibility of multi-protocol attacks. In *Proc. 1st International Conference on Availability, Reliability, and Security (ARES) 2006*, pages 287–294. IEEE, 2006.
- [16] A. Datta, A. Derek, J. C. Mitchell, and D. Pavlovic. Secure protocol composition. *Electronic Notes in Theoretical Computer Science*, 83(15), 2004.
- [17] F. Giesen, F. Kohlar, and D. Stebila. On the security of TLS renegotiation. In A.-R. Sadeghi, V. D. Gligor, and M. Yung, editors, *ACM CCS 13*, pages 387–398. ACM Press, Nov. 2013.
- [18] J. D. Guttman and F. J. Thayer Fabrega. Protocol independence through disjoint encryption. In *Proceedings 13th IEEE Computer Security Foundations Workshop (CSFW-13)*, pages 24–34. IEEE, 2000.
- [19] B. Harris. RSA Key Exchange for the Secure Shell (SSH) Transport Layer Protocol. RFC 4432 (Proposed Standard), Mar. 2006.
- [20] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DHE in the standard model. In R. Safavi-Naini and R. Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 273–293. Springer, Aug. 2012.
- [21] T. Jager, K. G. Paterson, and J. Somorovsky. One bad apple: Backwards compatibility attacks on state-of-the-art cryptography. In *Proc. Internet Society Network and Distributed System Security Symposium (NDSS) 2013*, 2013.
- [22] J. Jonsson and B. S. Kaliski Jr. On the security of RSA encryption in TLS. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 127–142. Springer, Aug. 2002.
- [23] J. Kelsey, B. Schneier, and D. Wagner. Protocol interactions and the chosen protocol attack. In B. Christianson, B. Crispo, M. Lomas, and M. Roe, editors, *Proc. 5th International Workshop on Security Protocols*, volume 1361 of *LNCS*, pages 91–104. Springer, 1997.
- [24] F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DH and TLS-RSA in the standard model. Cryptology ePrint Archive, Report 2013/367, 2013. <http://eprint.iacr.org/2013/367>.
- [25] H. Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 631–648. Springer, Aug. 2010.
- [26] H. Krawczyk, K. G. Paterson, and H. Wee. On the security of the TLS protocol: A systematic analysis. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 429–448. Springer, Aug. 2013.
- [27] B. A. LaMacchia, K. Lauter, and A. Mityagin. Stronger security of authenticated key exchange. In W. Susilo, J. K. Liu, and Y. Mu, editors, *ProvSec 2007*, volume 4784 of *LNCS*, pages 1–16. Springer, Nov. 2007.
- [28] N. Mavrogiannopoulos, F. Vercauteren, V. Velichkov, and B. Preneel. A cross-protocol attack on the TLS protocol. In T. Yu, G. Danezis, and V. D. Gligor, editors, *ACM CCS 12*, pages 62–72. ACM Press, Oct. 2012.
- [29] P. Morrissey, N. P. Smart, and B. Warinschi. A modular security analysis of the TLS handshake protocol. In J. Pieprzyk, editor, *ASIACRYPT 2008*, volume 5350 of *LNCS*, pages 55–73. Springer, Dec. 2008.
- [30] K. G. Paterson and G. J. Watson. Plaintext-dependent decryption: A formal security treatment of SSH-CTR. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 345–361. Springer, May 2010.
- [31] P. Rogaway. Formalizing human ignorance. In P. Q. Nguyen, editor, *Progress in Cryptology - VIETCRYPT 06*, volume 4341 of *LNCS*, pages 211–228. Springer, Sept. 2006.
- [32] D. Stebila and J. Green. Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer. RFC 5656 (Proposed Standard), Dec. 2009.
- [33] F. J. Thayer Fabrega, J. Herzog, and J. D. Guttman. Mixed strand spaces. In *Proceedings 12th IEEE Computer Security Foundations Workshop (CSFW-12)*, pages 72–82, 1999.
- [34] W.-G. Tzeng and C.-M. Hu. Inter-protocol interleaving attacks on some authentication and key distribution protocols. *Information Processing Letters*, 69(6):297–302, March 1999.
- [35] D. Wagner and B. Schneier. Analysis of the SSL 3.0 protocol. In *Proc. 2nd USENIX Workshop on Electronic Commerce*, 1996.
- [36] S. C. Williams. Analysis of the SSH key exchange protocol. Cryptology ePrint Archive, Report 2011/276, 2011. <http://eprint.iacr.org/2011/276>.
- [37] T. Ylonen and C. Lonvick. The Secure Shell (SSH) Authentication Protocol. RFC 4252 (Proposed Standard), Jan. 2006.
- [38] T. Ylonen and C. Lonvick. The Secure Shell (SSH) Connection Protocol. RFC 4254 (Proposed Standard), Jan. 2006.
- [39] T. Ylonen and C. Lonvick. The Secure Shell (SSH) Protocol Architecture. RFC 4251 (Proposed Standard), Jan. 2006.
- [40] T. Ylonen and C. Lonvick. The Secure Shell (SSH) Transport Layer Protocol. RFC 4253 (Proposed Standard), Jan. 2006. Updated by RFC 6668.