

Deniable Internet Key Exchange*

Andrew C. Yao[†] Frances F. Yao[‡] Yunlei Zhao[§] Bin Zhu[¶]

Abstract

In this work, we develop a family of protocols for deniable Internet Key-Exchange (IKE) with the following properties:

- Highly practical efficiency, and conceptual simplicity and clarity.
- Forward and concurrent (non-malleable) deniability against adversaries with arbitrary auxiliary inputs, and better privacy protection of players' roles.
- Provable security in the Canetti-Krawczyk post-specified-peer model, and maintenance of essential security properties not captured by the Canetti-Krawczyk security model.
- Compatibility with the widely deployed and standardized SIGMA (i.e., the basis of IKEv2) and (H)MQV protocols, when parties possess DL public-keys.

Our protocols could potentially serve, in part, as either the underlying basis or a useful alternative for the next generation of IKE (i.e., IKEv3) of IPsec (in particular, when deniability is desired). In view of the wide deployment and use of IKE and increasing awareness of privacy protection (especially for E-commerce over Internet), this work is naturally of practical interest.

1 Introduction

Key-exchange (KE) is a traditional area of cryptography. But, key-exchange is also a quite special area of cryptography, in view of its seemingly simple yet error prone nature. Specifically, most key-exchange protocols seem to be very simple and even intuitive, and thus appearing to be easily designable, but the literature has been witnessing that the design of correct and secure KE turns out to be extremely error prone and could be notoriously subtle and difficult (the literature is filled with protocols that have been found to contain certain security flaws).

The Internet Key-Exchange (IKE) protocols [16, 18], pioneered by Photuris [17], SKEME [20], Oakley [26] and SIGMA [21], are the *core* cryptographic protocols to ensure Internet security, which specifies key exchange mechanisms used to establish shared keys for use in the Internet Protocol Security (IPsec) standards [19]. The IPsec and IKE are intended to protect messages communicated *in the IP layer*, i.e., “layer 3” of ISO-OSI, that processes the transmission of messages *using the network addresses possibly without knowing end-user peers' identities*. IKE and IPsec can in turn be used to offer confidentiality, authentication and privacy for communication protocols in the higher layers of ISO-OSI (note that many communication protocols including many authentication protocols which are invoked by *end-users with explicit peer's identity information* work at “layer 7” of ISO-OSI, i.e., the application layer) [24].

The standard of IKE has gone through two generations (versions). The first generation IKEv1 [16] uses public-key encryption as the authentication mechanism (with SKEME serving as the basis in part),

*The work described in this paper was supported in part by the National Basic Research Program of China Grant (973) 2007CB807900, 2007CB807901, and by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project Number CityU 122105) and CityU Research Grant (9380039).

[†]Institute for Theoretical Computer Science, Tsinghua University, Beijing, China. andrewcyao@tsinghua.edu.cn

[‡]Department of Computer Science, City University of Hong Kong. csfyao@cityu.edu.hk

[§]Corresponding author. Software School, Fudan University, Shanghai 200433, China. ylzhao@fudan.edu.cn

[¶]Microsoft Research Asia, Beijing, China. binzhu@microsoft.com

and IKEv2 [18] uses public-key signature as the authentication mechanism (with SIGMA serving as the basis in part).

Several salient features of SIGMA helped make it adopted as the basis of IKEv2: (1) it works in the post-specified-peer ID model; (2) it provides identity concealment and plausible deniability; etc.

Desirability for IKE with post-specified-peer ID [8]. Note that IPsec and IKE are intended to protect communication protocols in the IP layer that processes the transmission of messages *using network addresses (rather than end-users' IDs)*. This brings forth an issue: the peer's identity may be not available at the onset of the IKE, and may be learned by the party only after the protocol run evolves (*even learning from the last message in a run of the protocol*). More specifically, in IKE a party may be activated to exchange a key with an "address" of a peer without knowing the identity of the peer. This is quite common in practical applications, in particular for IKE. For example, a key-exchange session may take place with any one of a set of servers sitting behind a (url/ip) address specified in the session activation. Furthermore, even for protocols running in the higher layers (e.g, the application layer), a party may respond to a key-exchange request coming from a peer who, for the sake of preserving its privacy, is not willing to reveal its identity over the network and, sometimes, not even to the responder before the latter has authenticated itself (e.g., a roaming mobile user connecting from a temporary address, or a smart-card that authenticates the legitimacy of the card-reader before disclosing its own identity). Such a general setting for IKE is referred to as "post-specified peer" model [8].

Desirability for deniable IKE [12, 24]. Deniable communication is an essential privacy property, and has always been a central concern in personal and business communications, with off-the-record communication serving as an essential social and political tool. Given that many of these interactions now happen over digital media (email, instant messaging, web transactions, virtual private networks), it is of critical importance to provide these communications with "off-the-record" or deniability capability: the sender of a message should be able to deny, e.g., in court, that he or she has sent that message. In general, we can just run a deniable authentication protocol for each message to be sent when deniability of the message is desired. However, the beauty of using deniable key-exchange is that if the key-exchange protocol is deniable, then all transactions using the session-key produced by the key-exchange protocol will be automatically deniable. In addition, offering deniability at the IP layer is desirable because it enables various privacy services to be offered at the higher layers with uncompromised quality. Note that a privacy problem at the IP layer can cause irreparable privacy damage at the application layer. For example, an identity connected to an IP address, if not deniable, certainly nullifies an anonymous property offered by a fancy cryptographic protocol running at the application level. A *variant* of the SIGMA protocol of [21] (i.e., the underlying basis of public-key IKEv2) has been proved to satisfy some plausible partial deniability assuming key-aware derivation (typically, with random oracle assumption of hash functions) [12].

On the desirability of forward and concurrent deniability. In the traditional definition of deniable authentication, sender Alice wants, in protecting her privacy, to prevent (possibly malicious) receiver Bob from proving to a third party that he has received a message from Alice by requiring that the protocol is actually a (computational) zero-knowledge (ZK) protocol in the sense that Bob's view can be simulated. Furthermore, we hope the ZK property holds against malicious Bob interacting with Alice *concurrently* in many sessions, i.e., the concurrent ZK (CZK) property [14].

However, as clarified in [11], there are scenarios in which deniability is actually a concern to the receiver's privacy. Consider the following example presented in [11]. Alice and Bob are involved in some shady transaction, like drug-dealing or money laundering. Alice wants to make sure that her communications to Bob cannot be later linked to her, so she uses deniable authentication. Bob thinks that such communications are indeed deniable, stores all the messages in his hard disk. Later the operation is busted by the police and Alice and Bob end up in jail, and Bob's computer is seized. Alice is offered a sweet deal in exchange for her cooperation in linking Bob to the crime (Bob is claiming that the messages in his hard disk are not coming from Alice, that he never talked to her, actually does not even know her, all messages are just simulations!). Alice produces some piece of secret information (her secret-key for example) that indeed shows that the transcripts in Bob's hard disk are actually authentic and not simulated. Bob ends up in jail, cursing himself for dropping out of crypto class in graduate

school.

The above example shows that deniability is not just a concern for the sender, but also for the receiver. What we would like to happen is that if the sender acts honestly during the protocols, she should not be able at a later stage to claim the messages are authentic. This property is called *forward deniability*, as it has some affinity to the notion of forward secrecy. In particular, Di Raimondo and Gennaro have shown that computational ZK does not guarantee forward deniability [11], but statistical/perfect (concurrent) ZK guarantees forward (concurrent) deniability.

Another typical example where forward deniability is important is for electronic elections [11]. There, while it is important that both parties, the voting authority and the voter, authenticate each other (for the authority to know that the voter has the right to vote, for the voter to know that his or her vote will be counted), it is also mandatory to prevent either party from walking away with a non-repudiable proof of what the actual vote was (the message being authenticated). This application, in particular, shows the importance of forward deniability: if the voter (sender) is authenticating his or her vote to the authority, not only the latter should not be able to prove to a third party how the voter voted, but even more importantly, the voter should not be able to do so either at a later stage, to prevent coercion and vote-selling.

Some critiques and concerns on IPsec and IKE. One prominent criticism of IPsec and IKE, in the literature, is about their intensive system complexity and lack of clarity [24]. In particular, they contain too many options and too much flexibility. The high system complexity relates to an extreme obscurity in the system specification. This obscurity may confuse expert reviewers and blind them from seeing potential security weakness, or may mislead implementors and cause them to code flawed implementations.

Another main concern about IKE is its ability of protecting players' privacy over Internet (besides its ability of ensuring security). For Internet key exchange, both *security* and *privacy* are desired, which would also have been the major criterion underlying the evolution of IKE standards.

1.1 Our contributions

In this work, we present a family of protocols of deniable IKE (in accordance with the main model and the aggressive model of IKE in IPsec) with the following properties:

- Highly practical efficiency, and conceptual simplicity and clarity.
- Forward and concurrent (non-malleable) deniability against adversaries with arbitrary auxiliary inputs, and better privacy protection of players' roles.
- Provable security in the Canetti-Krawczyk post-specified-peer model [8], and maintenance of essential security properties not captured by the Canetti-Krawczyk security model.
- Compatibility with the widely deployed and standardized SIGMA [21] (i.e., the basis of IKEv2) and (H)MQV [23, 22] protocols, when parties possess DL public-keys.

We would like to remark that, though our deniable IKE seems to be simple and even intuitive, its design was thoroughly guided by the profound theory of concurrent non-malleability [13], universal composability [6], and the theory of secure key-exchange in the post-specified-peer model [?, 18].

The provable security of our protocols relies on the random oracle (RO) assumption and a (very reasonable and essentially weaker) variant of knowledge-of-exponent assumption (KEA), called “fresh-challenge KEA with arbitrary auxiliary input”. We note here that the RO and/or KEA assumptions are common in the stringent situation of deniable key-exchange and authentication [4, 3, 11, 22, 10, 12, 28], and actually most known KEA-based deniable authentication and key-exchange protocols are not secure against adversaries with arbitrary auxiliary input. More details can be found in Section 3.

Of independent interest of this work are the clarifications and revision of KEA assumption with arbitrary auxiliary input, and clarifications and observations on the *independent* Stinson-Wu key-exchange (SW-KE) protocol of [28].

1.2 Justification and comparison with SIGMA

In comparison with SIGMA protocol (i.e., the basis of IKEv2), our deniable IKE has the following advantages:

- Our deniable IKE is of very high conceptual simplicity and clarity (which can be seen from the following discussions and the clarifications presented in Section 3).
- Our deniable IKE provides full concurrent and forward (non-malleable) deniability, while SIGMA does not (due to the underlying signatures used). Besides full deniability, our deniable IKE also provides seemingly better privacy protection of players' roles than SIGMA. Detailed clarifications can be found in Appendix A.
- Our deniable IKE is essentially more practical than SIGMA in practice.

Specifically, in our deniable IKE, the player does not use signatures to show *proof-of-possession* of its secret-key; rather, it shows the stronger *proof-of-knowledge* of its secret-key via an approach such that the proof-procedure, and the verification-procedure respectively, only needs *one* modular exponentiation respectively (we do not know, to our knowledge, so practical secure signature scheme even in the RO model).

- Our deniable IKE is essentially (maybe controversially) more secure than SIGMA in practice.
 - **Security beyond the Canetti-Krawczyk model.** SIGMA is proved secure in accordance with the secure KE definition (SK-security) in Canetti-Krawczyk post-specified-peer model. According to the SK-security, a party that completes a session has the following guarantees: (1) if the peer to the session is uncorrupted then the session-key is unknown to anyone except this peer; (2) if the peer completes a matching session then the two parties have the same shared key. In comparison, the security of our deniable IKE goes one better with the following provable enhanced guarantee of (2): if a (possibly malicious) player completes a matching session, then the two parties, not only, have the same shared key, *but also and more importantly, the (possibly malicious) player does “know” the corresponding DH key exponent (i.e., the shared session-key) and the secret-key corresponding to its alleged public-key.* This provides extra much more robust and stronger security guarantee, which helps us provably rule out most essential security threats not captured by the Canetti-Krawczyk model (detailed clarifications are presented in Section 3). SIGMA lacks formal proofs for some essential security properties not captured by Canetti-Krawczyk model, at least under standard hardness assumptions.
 - **Essential security in practice.** With our deniable IKE solution, player's proof-of-knowledge of its secret-key and the DH-exponent is authentically bounded with session-specific information, which is masked by hash function H and then re-masked by MAC (with the output of H as the key of MAC). Player's proof-of-knowledge can only be verified by its peer, *but not seen directly by the peer or any other party*; furthermore, *correctly verifying player's proof-of-knowledge implies that the peer learns nothing from the proof-of-knowledge and can do no damage to the player's security. That is, player's proof-of-knowledge is of both (double-masked) hiding and authentic (non-malleable) binding.* Under CDH assumption and the underlying key-derivation function KDF is secure, adversarial efforts against our deniable IKE are essentially confined to attacks on the underlying MAC. When player's ID is put into the input of hash function (and/or some public session-specific information is put into the input of MAC), the adversary still gets no essential advantages even if the underlying MAC is broken (more details can be found in Section 3). This essentially leaves the adversary only the play to brute-force retrieve the proof-of-knowledge double-masked by hash function and MAC. Above discussed properties do not hold for SIGMA. In particular, in SIGMA, player's signature is seen and publicly-verifiable by anyone, and the security critically relies on the

security of both signatures and MAC. In this way, in practice, any essential way to violate the security of our deniable IKE can also be used to violate the security of SIGMA (unless one can effectively break the KEA assumption).

The relative disadvantages:

- The provable security of SIGMA in Canetti-Krawczyk post-specified-peer model is based on the security of the underlying cryptographic components: signatures, MAC and pseudorandom function (PRF), which can be based on the standard assumptions [8]. Note that the provable security of our protocol is based on KEA assumption in the random oracle model.

But, one should realize that the theoretical provable security of SIGMA in Canetti-Krawczyk model has some limitations in practice.

- As discussed above, there are a list of *essential* security properties of key-exchange that are not captured by Canetti-Krawczyk model [22, 25], which still need to be analyzed case by case. SIGMA lacks formal proofs for some of those un-captured essential security properties, *at least under standard assumptions*.
- Though SIGMA has theoretical provable security in Canetti-Krawczyk model under standard assumptions, when SIGMA are used in practice with practical signatures and MAC and PRFs, in most cases, the security still relies on the random oracle assumption and/or some non-standard assumption. *We remark that RO and/or KEA assumptions are common in the stringent situation of deniable key-exchange and authentication [4, 3, 11, 22, 10, 12, 28].* In particular, partial deniability for a variant of SIGMA, with 0 and 1 removed from inputs of signatures, is proved, assuming key-aware derivation (typically, with the RO assumption of hash functions and/or some other types of non-standard assumptions).
- The theoretical provable security in Canetti-Krawczyk model is for the basic version of SIGMA. Though it was discussed on the extension to some variants of SIGMA [8], the security analysis might not be extended to all variants of SIGMA proposed and used in practice. We note that the number of proposed variants of SIGMA is much more than that of our deniable IKE. We essentially propose and recommend only *one* variant, i.e., putting player’s ID also into the input of hash function and putting some public session-specific information also into the input of MAC. This proposed and *actually recommended* variant is still provably secure under KEA and RO assumptions, but provides seemingly much stronger security guarantee in practice (in particular, an adversary seem to get no extra advantages even if it can break the security of MAC). More details can be found in Section 3.
- Our deniable IKE requires players possessing DL public-keys. On the one hand, this limits the generality of public-keys used by players, in view that SIGMA can use any secure signature scheme; On the other hand, our deniable IKE is well compatible with the setting of widely deployed and standardized (H)MQV (which does require players possessing DL public-keys), and also, practical secure signature schemes with DL public-keys (in the RO model) are common in practice.

2 Preliminaries

Important note on presentation. *Although our deniable IKE protocols and their security properties can be rigorously formalized and analyzed (that can be found in the full version), we have chosen to present this work here in a somewhat informal and intuitive way to highlight basic ideas of our approach.*

Secure key-exchange in the post-specified-peer model [8]. Let $\langle A, B \rangle$ be a key-exchange protocol. Each activation of the protocol at a party results in a local procedure, called a *session*. Sessions can run concurrently and each session is identified by the party’s ID and a (unique) session identifier, e.g., the KE session (A, sid) run by A with session identifier sid . The output of the activation of the KE

session (A, sid) is $((A, sid, B), sk)$, where (A, sid, B) is the public output, and sk is the secret session-key (sessions can be aborted without producing the secret session-key sk , and in this case a special symbol is output). A session maintains a local state, e.g., the exponent x of a Diffie-Hellman (DH) component g^x , that is erased when the session completes (i.e., when the party gets the session-key). Each party may have additional state, such as a long-term signature key, which is accessed by different sessions and which is not part of any particular session state. If the session (A, sid) is a completed session with output $((A, sid, B), sk)$, then the session (B, sid) run at the side of B is called the *matching session* of (A, sid) if either (B, sid) is not completed or (B, sid) is finished with (B, sid, A) as its public output.

The attacker \mathcal{A} is a probabilistic polynomial-time (PPT) algorithm that controls the *unauthenticated* communication channels and can schedule session activations and message transmissions at will and sees the outputs of sessions except for the values of session keys. In addition, the attacker can have access to secret information via *session exposure* attacks of three types: session-state reveal, session-key queries, and party corruption. Session-state reveal is directed at an incomplete session and its result is that the attacker learns the session state, e.g., x of g^x , for that particular session. A session-key query can be performed against a completed session and the result is that the attacker learns the corresponding session-key. Finally, party corruption means that the attacker learns *all* information in the memory of that corrupted party; in addition, from the moment a party is corrupted all its actions are totally controlled by the attacker. Sessions can expire, and from the time a session has expired the attacker is not allowed to perform a session-key query or a state-reveal attack against the expired session, but is still allowed to corrupt the party that holds the session. (Protocols, ensuring that expired sessions are protected even in case of party corruption, are said to enjoy *perfect forward secrecy*.)

A KE protocol is secure, called *satisfying SK-security*, if the following two conditions hold: (1) If two uncorrupted parties complete matching sessions in a run of the key-exchange protocol, under the attacker \mathcal{A} , the session key output in the matching sessions is the same; (2) For any session, the attacker cannot distinguish session-key output of the session from an independent random value, provided the session and its matching sessions are *not exposed*.

The SK-security guarantees some essential security properties of key-exchange protocols, in the following sense: a party that completes a session has the guarantee that: (1) if the peer to the session is uncorrupted then the session-key is unknown to anyone except this peer; (2) if the peer completes a matching session then the two parties have the same shared key. These properties are guaranteed even if the attacker has corrupted other parties in the system and/or the attacker has learned the secret information from other sessions. Furthermore, as shown in [7, 8], the SK-security ensures secure communication channels via the shared session-keys.

Forward and concurrent deniable key-exchange. A key-exchange protocol $\langle A, B \rangle$ is called *concurrently forward deniable* for A , if for any PPT malicious \mathcal{A} running *concurrently* with instances of A in any polynomially many sessions (in any name in each session), there exists a PPT simulator S that, on public system parameters and public-keys, outputs a simulated transcript that is identical to (or statistically indistinguishable from) the real view of \mathcal{A} in its real concurrent executions (i.e., concurrent perfect/statistical zero-knowledge). *We stress that the real view of \mathcal{A} consists of all secret session-key outputs of all sessions, which means that the deniability of key-exchange implies deniability of subsequent communications via the shared session-keys.* Also, the public-keys used by \mathcal{A} in various names/IDs are not guaranteed to be well-generated (note that the names/identities claimed by \mathcal{A} in different sessions are not necessarily the same), though the association of player's ID and its public-key is certified by trusted CA. That is, \mathcal{A} may not know corresponding secret-keys of the various public-keys used by it.

The key-exchange protocol $\langle A, B \rangle$ is called *forward and concurrently deniable*, if the concurrent forward deniability property holds for both A and B *simultaneously*. That is, for any PPT man-in-the-middle (MIM) \mathcal{A} concurrently interacting with both instances of A and instances of B , its view can be perfectly simulated. Note that, for simplicity of presentation, w.l.o.g., we assume that the MIM adversary \mathcal{A} interacts concurrently with instances of only two parties with names/IDs: A and B . The formulation can be trivially extended to the general case of many parties.

Forward and concurrent deniable authentication. A forward and concurrent deniable authentication protocol $\langle P, V \rangle$ is an *interactive* protocol between a prover P with public-key PK_P and a verifier V (V

is not necessarily of public-key), satisfying the following: (1) It is concurrent perfect zero-knowledge for P (against possibly malicious V^*); (2) It is unforgeable against concurrent adaptive chosen-message attack. That is, a PPT adversary \mathcal{A} , who can run concurrently with polynomially many instances of P on inputs it chooses adaptively, still cannot authenticate to an honest V a new message that P did not ever authenticate in the concurrent executions. In a way, forward and concurrent deniable authentication could be viewed as an *interactive* public-key signature system with the additional property of *concurrent perfect zero-knowledge*.

3 Deniable Internet Key-Exchange: the Main Model

In this section, we present the deniable IKE protocol of the main model, revisit and clarify the subtleties of KEA assumptions with concurrent adversaries of arbitrary auxiliary input, clarify the essence of ideas and analyze security properties.

3.1 The protocol construction

The deniable IKE protocol of the main model is depicted in Figure 1. In the protocol specification, A and B stand for player IDs, x and y are randomly and independently chosen from Z_q . For simplicity of presentation, we do not specify the verification procedure details at each step of the protocol, which can be understood in a natural way. In particular, the non-one and order q of g^x , g^y , g^a and g^b will be checked (actually, the non-one and order q of players' public-keys can be checked during key registration).

<p>Public system parameters: (p, q, g, H), where p and q are large prime numbers such that $p = 2q + 1$, g is an element of Z_p^* of order q, and H is a hash function.</p> <p>Player keys: (g^a, a) is the public-key and secret-key of A, and (g^b, b) is the key pair of B, where a and b are randomly chosen from Z_q. The association of player's ID and its public-key is certificated by $Cert$ issued by a trusted CA. We stress that for a malicious player, its public-key can be any element of Z_p^* of order q. That is, <i>no proof-of-knowledge (of secret-keys) is required for key-registration and certification</i>. (Rather, each party proves its knowledge of the secret-key to its peer during the run of the protocol, in a non-malleable zero-knowledge (NMZK) authentic binding way.)</p>	
<p>Actual implementation:</p> <p>$A \xrightarrow{\quad sid, X = g^x \quad} B$</p> <p>$A \xleftarrow{\quad sid, B, Y = g^y, NMZK_{(B,y)} = MAC_{H(sid, g^y, g^x, g^{xy})}(B) \quad} B$</p> <p>$A \xrightarrow{\quad sid, (A, Cert_A), NMZK_{(a,x)} = MAC_{H(sid, g^x, g^y, Y^a, Y^x)}(A) \quad} B$</p> <p>$A \xleftarrow{\quad sid, Cert_B, NMZK_{(b,y)} = MAC_{H(sid, g^y, g^x, X^b, X^y)}(B) \quad} B$</p>	<p>Natural language interpretation:</p> <p>$A \xrightarrow{\quad \text{Hi, I wanna have a session with } g^x \quad} B$</p> <p>$A \xleftarrow{\quad \text{Okey, } g^y \text{ "is sent by } B \text{ " who knows } y \quad} B$</p> <p>$A \xrightarrow{\quad \text{Well, I TRULY know } (a, x) \quad} B$</p> <p>$A \xleftarrow{\quad \text{I TRULY know } (b, y) \text{ as well} \quad} B$</p>
<p>The session-key output is computed by $KDF(H(g^{xy}), \cdot)$, where KDF is a key-derivation function, e.g., $PRF_{H(g^{xy})}(0)$. In practice, MAC and PRF will be implemented with HMAC of [1].</p>	

Figure 1: Deniable IKE (the main model)

3.2 KEA assumption, revisited

The Knowledge-of-Exponent (KEA) assumption. The security analysis of the deniable IKE (depicted in Figure 1) relies on the KEA assumption in the random oracle model. The KEA assumption roughly says that: given a challenge $C = g^c$ for randomly chosen $c \in Z_q$, if a *uniform* PPT algorithm \mathcal{A} can output *valid* (X, Y) of order q each such that $Y = X^c$, then \mathcal{A} must *know* $x = \log_g^X$ and x can be efficiently extracted from the *internal states* of \mathcal{A} (i.e., *its codes, random coins and the transcript up to the point outputting (X, Y)*). Note that the KEA assumption is a *non-black-box* assumption by nature [3]. By “uniform” we mean that \mathcal{A} takes no auxiliary input. The KEA assumption was introduced in [9], and has been used in many subsequent works (e.g., [4, 3, 11, 22, 10, 12], etc). In particular, the KEA assumption, as well as the random oracle assumption, plays a critical role for provable knowledge-awareness (PA) of public-key encryption systems, e.g., Cramer-Shoup and OAEP [4, 22, 5], and for provable deniability of authentication and key-exchange [11, 22, 12].

Subtleties of KEA assumption for protocols running concurrently in the public-key model. As the KEA assumption is not a standard hardness assumption (though it seems to be very reasonable), we should take much care of the employment of the KEA assumption in practice, especially for IKE. When using KEA in the public-key model for provable security, e.g., for provable plain-text awareness and deniability [4, 11, 22, 10, 12], the challenge C is actually the player’s *public-key*. It would be fine w.r.t. a uniform adversary \mathcal{A} without auxiliary input (even with some well-predefined auxiliary input as is done in [22, 10, 12]). But, when \mathcal{A} takes an *arbitrary* auxiliary input and runs protocols *concurrently* in the *public-key* model (which is just the case of IKE), the KEA assumption can be totally meaningless. Specifically, in such cases, as the challenge C is player’s public-key, a valid (X, Y) could be just got from adversary’s arbitrary auxiliary input (modeling information collected/eavesdropped over the network from executions of other *different* protocols *prior to any interaction of the protocol at hand*) or from up-to-now transcript of concurrent executions of the protocol at hand. In other words, there could be practical ways to break the KEA assumption if the challenge C is public-key and \mathcal{A} is allowed with arbitrary auxiliary input and/or concurrent executions in the public-key model, which is however very natural and practical for cryptographic schemes when they are run concurrently in unpredicted adversarial settings like the Internet (in particular this is the case for IKE).

We note that the security proofs of PA properties in [4, 10] are critically based on the assumption that adversaries have no arbitrary (rather than well pre-specified) auxiliary input. In particular, the plaintext awareness properties of the public-key encryption schemes considered in [4, 10] will not hold w.r.t adversaries with arbitrary auxiliary input. Furthermore, it implies that the deniability of the PKE-based message-authenticator of [2] and the SKEME key-exchange protocol of [20], as proven in [12], will also not hold w.r.t adversaries with arbitrary auxiliary input when the underlying public-key encryption scheme is implemented by the Cramer-Shoup encryption scheme. (Note that the Cramer-Shoup encryption scheme is currently the only known PA2-secure public-key encryption scheme based on KEA but without random oracles [10].) This is the case even for the variants of the PKE-based message-authenticator and SKEME key-exchange protocols considered in [12], where fresh nonce is sent on the top to frustrate cooperating judge attack. We note that security proof (for the principal security of HMQV against revealing ephemeral exponent) in [22], *that uses both the KEA assumption and the random oracle assumption*, will also not work when adversaries have arbitrary auxiliary inputs.

Fresh challenge KEA assumption with arbitrary auxiliary input. According to the above clarifications, when using the KEA assumption for protocols running in the public-key model, we allow the algorithm \mathcal{A} to take arbitrary auxiliary input, but require that the challenge C is fresh in the sense that each C is answered by \mathcal{A} at most once. This implies that C cannot be player’s public-key. Such a revised KEA assumption is called *fresh challenge KEA assumption with arbitrary auxiliary input*, which is necessary and reasonable for protocols running concurrently in the public-key model.

Note that, as clarified, the arbitrary auxiliary input captures any information collected over the network from executions of arbitrary different protocols prior to any interaction of the protocol at hand. When using the KEA assumption for provable security of protocols running concurrently in the public-key model, the protocol should be designed that the transcript of the concurrent execution of

the protocol at hand should also not help \mathcal{A} break the KEA assumption, which is left in the security analysis of the protocol at hand and should be the responsibility and care of the protocol designer.

3.3 Discussion and analysis: security features and properties

Some notable features of our deniable IKE protocol. We first briefly point out some notable features of our deniable IKE protocol.

- Peer’s ID and public-key concealment: messages from one party do not include its peers’s ID and public-key. This perfectly matches the requirement of working in the post-specified-peer model and the principle of privacy-preserving.
- Concurrent non-malleable (and maybe universal composable) zero-knowledge (CNMZK) authentic binding: party’s ID (e.g., A), its DH-component (e.g., g^x), its proofs of knowledge of its secret-key (i.e., a or b) and the discrete logarithm (i.e., x or y) of its DH key component (i.e., X or Y), and the session-specific information (e.g., (sid, g^x, g^y)) are authentically bounded in a CNMZK way, through $NMZK_{(B,y)}$, $NMZK_{(a,x)}$ and $NMZK_{(b,y)}$. $NMZK_{(B,y)}$ denotes that it is a non-malleable zero-knowledge proof that the second round message is bounded to responder B who knows y . $NMZK_{(a,x)}$ (resp., $NMZK_{(b,y)}$) denotes that it is a NMZK proof of knowledge of (a, x) of initiator A (resp., (b, y) of responder B .)

This essentially implies that messages of one session cannot be malleated into other successful sessions, or messages of one successful session cannot be malleated from other sessions, other than just simply copying which however is not deemed to be a harmful activity. Here, we briefly point out that $NMZK_{(B,y)}$ plays multiple key roles with our deniable IKE: providing full deniability for initiator, preventing cutting-last-message attack, and preventing reflection attacks, etc. Detailed clarifications can be found in the following analysis.

- All authentic messages from B , i.e., $NMZK_{(B,y)}$ and $NMZK_{(b,y)}$, can be computed *merely* from the discrete logarithm, i.e., x , of it’s peer’s DH key component X (without knowing $(a, b, y)!$), and A does have proved the knowledge of x via $NMZK_{(a,x)}$; The authentic message $NMZK_{(a,x)}$ can be computed *merely* from the discrete logarithm, i.e., y , of it’s peer’s DH key component Y (without knowing $(a, b, x)!$), , and B does have proved the knowledge of y via $NMZK_{(B,y)}$ and $NMZK_{(b,y)}$. This is typically fair for the privacy concerns for both A and B .
- Better privacy protection of players’ roles: As clarified in Appendix A, in comparison with SIGMA, our deniable IKE provides seemingly better privacy protection of players’ roles. Specifically, one cannot distinguish player’s role only from $NMZK_{(a,x)}$ or $NMZK_{(b,y)}$; the whole transcript may reveal the roles of initiator and responder, but the whole transcript can be totally denied.

Security in the post-specified-peer model. We show that, under fresh-challenge KEA assumption with arbitrary auxiliary input and the random oracle assumption of H and computational Diffie-Hellman (CDH) assumption, the protocol depicted in Figure 1 satisfies SK-security in the post-specified-peer model of [8], which implies secure communication channels via the shared session-keys and perfect forward security. The security analysis follows, in a very essential way, the security analysis of the variant of SIGMA (combining MAC and signatures) presented in [8], which is quite complicated and lengthy and is deferred to the full version. Below, we provide some high-level hints of the proof. Our IKE protocol actually could correspond to a variant of SIGMA in which the MAC is combined into the signature, e.g., combining $SIG(sid, g^x, g^y)$ and $MAC(sid, A)$ into $SIG(MAC(sid, g^x, g^y), A)$. This variant of SIGMA is still secure in the post-specified-peer model [8]. In view of our protocol, we note that the first three rounds of our IKE protocol actually amount to a deniable authentication (i.e., interactive deniable public-key signature) with A as the signer, and the last three rounds amount to a deniable authentication with B as the signer (more details are deferred to the full version).

Forward and concurrent (non-malleable) deniability. We show that, under the fresh challenge KEA assumption with arbitrary auxiliary input and the RO assumption of H and the CDH assumption,

the protocol depicted in Figure 1 is forward and concurrently (non-malleable) deniable. For any PPT MIM-adversary \mathcal{A} concurrently interacting with instances of A and B with possible different IDs/names and public-keys in different sessions, we briefly hint at the security proof.

From the viewpoint of party A , for the *valid* second-round message $(sid, M, g^y, NMZK_{(M,y)} = MAC_{H(sid, Y=g^y, X=g^x, g^{xy})}(M))$ of the (A, sid) -session, received from \mathcal{A} in the name of M of public-key g^m w.r.t a fresh challenge g^x generated and sent by the simulator (in the name of A) at the first-round, if g^y was not generated by \mathcal{A} and $M \neq B$ and \mathcal{A} did not query the random oracle with (sid, g^y, g^x, g^{xy}) , there are essentially two possible ways for \mathcal{A} to compute out such message:

1. $NMZK_{(M,y)}$ is malleated by \mathcal{A} from some second message sent by honest B (who generated g^y and sent $NMZK_{(B,y)}$), by breaking the security of the underlying MAC. Note that the session-ID sid is unique to (A, sid) and its matching session (B, sid) , and the order of (g^y, g^x) or (g^x, g^y) indicates the role of player (i.e., whether responder B or initiator A).
2. Compute g^{xy} from g^x and g^y , and guess successfully the answer of the random oracle of H on (sid, g^y, g^x, g^{xy}) .

By the random-oracle of H and the security of the underlying MAC and the CDH assumption, \mathcal{A} must query the random oracle with (sid, g^y, g^x, g^{xy}) , which just contains a valid answer (g^y, g^{xy}) w.r.t. the freshly generated challenge of g^x . By the KEA assumption, the simulator (who programs the random oracle) can extract y , and then finish the (A, sid) -session perfectly. Note that in the third-round message, the values of (Y^a, Y^x) in $NMZK_{(a,x)}$ can be computed merely from the extracted value y by the simulator without knowing (a, x) .

Now, from the viewpoint of B , in the session (B, sid) , the simulator works as follows: after receiving the first-round message g^x from \mathcal{A} , it perfectly computes the second-round message in the name of B by generating g^y . Note that the second-round message, in particular $NMZK_{B,y}$ perfectly hides g^{xy} due to the random-oracle of H . For the third-round message sent by \mathcal{A} in the name of M with public-key g^m , similar to above arguments, \mathcal{A} must have queries the random oracle with $(sid, g^x, g^y, Y^m, Y^x)$. Then, by the KEA assumption w.r.t. fresh challenge $Y = g^y$, the value x can be extracted. Note that the fourth-message can then be perfectly simulated from the extracted value x by the simulator without knowing (b, y) .

On the dual roles of $(sid, X, Y)/(sid, Y, X)$. We remark that including $(sid, X = g^x, Y = g^y)$ or (sid, Y, X) into the input of H plays dual roles: On the one hand, the uniqueness of sid provides freshness to each session in order to prevent *replaying* attacks; the order of (g^x, g^y) or (g^y, g^x) indicates the role of player (i.e., whether the initiator A or the responder B), which is necessary for resistance against *reflection* attack. On the other hand, the inclusion of (sid, g^x, g^y) or (sid, g^y, g^x) plays a critical role in the proof of concurrent deniability (i.e., concurrent zero-knowledge CZK) as well as in the proof of security in Canetti-Krawczyk model. With the CZK proof as an illustration, as mentioned the KEA assumption is a *non-black-box* assumption, and the CZK simulator needs to extract x (resp., y), directly in a straight-line way, from the internal state (the codes, random coins, and partial transcript up to now) of \mathcal{A} just at the point where it queries the random oracle H with valid (Y, Y^x) (resp., (X, X^y)) w.r.t. a freshly-generated g^x (resp., g^y). But, if the input to H does not include (sid, g^x, g^y) or (sid, g^y, g^x) , the CZK simulator cannot realize whether random-oracle queries made by \mathcal{A} include valid Y^x (resp., X^y) to a fresh challenge g^x (resp., g^y). *Important note: \mathcal{A} may query the oracle with g^{xy} after receiving g^x but much earlier before sending the second-round message of that session that contains $Y = g^y$, and so the CZK-simulator cannot check the validness of g^{xy} before receiving Y .* This means that the CZK simulator has to rewind \mathcal{A} whenever it needs to extract y after receiving the second-round message in question (including Y) from \mathcal{A} (as the value y may not be extracted from a later state of \mathcal{A}). But, such rewindings cause exponential blow-up in the simulation time of the CZK-simulator [14]. Furthermore, even if y could still be extracted from a later state of \mathcal{A} (assuming \mathcal{A} has faithfully been keeping all its local states), the CZK-simulator has to be involved with highly inefficient computation/exponentiations in order to check whether \mathcal{A} ever queried with Y^x in the history of transcript, which causes highly inefficiency of the CZK simulation. In a way, (sid, X, Y) or (sid, Y, X) serves as a pointer to help the

CZK-simulator realize *immediately* whenever the random-oracle queries made by \mathcal{A} includes valid Y^x or X^y .

According to above discussions and clarifications, we have the following theorem:

Theorem 3.1 (informal) *Under fresh-challenge KEA assumption with arbitrary auxiliary input and the random oracle assumption of H and CDH assumption, the IKE protocol depicted in Figure 1 satisfies SK-security in the post-specified-peer model, and is forward and concurrently deniable.*

But, as noted in [22, 25], there are some important security properties of key-exchange that are not captured by the security model of [2, 8]. We analyze those properties case by case below.

Resistance against cutting-last-message attack. We remark that the cutting-last-message attack, suffered by IKEv2 and SIGMA, does not work on our IKE protocol (the cutting-last-message attack on IKEv2 and SIGMA is briefly recalled in Appendix A). Specifically, after receiving the second-round message $(B, Y = g^y, NMZK_{(B,y)})$ from B , the MIM-adversary \mathcal{A} cannot compute and send to A the message of $(M, Y = g^y, NMZK_{(M,y)})$ in the name of $M \neq B$, due to the security of MAC and random oracle assumption of H and CDH assumption. In other words, if A receives a valid second-round message $(M, Y' = g^{y'}, NMZK_{(M,y')})$ from \mathcal{A} , it implies $Y' \neq Y$ and Y' is generated by \mathcal{A} who knows y' . In this case, the third-round message from A will be w.r.t. Y' generated by \mathcal{A} (not w.r.t. Y generated by B), and cannot be forwarded by \mathcal{A} to B in order to fool B to believe it will share a session-key with A , as $Y \neq Y'$.

A NOTE ON INCLUDING B 'S PLAIN ID IN THE SECOND-ROUND. We remark that including B 's ID in the second-round message is critical to rule out the cutting-last-message attack. *Note that B does not identify it in the second-round message and does not reveal its certificate.* In a way, including B 's plain ID in the second-round could be viewed as a slight compromise between security and privacy. In case one insists perfect identify protections for responder, B 's ID can be removed from the second-round message and Theorem 3.1 still holds, but the cutting-last-message will come.

Resistance against key-compromise impersonation (KCI) attack. Informally, security against KCI attack says that [22]: even if the attacker knows the private-key of a party, e.g., the private-key a of A , still cannot compromise the session-key established between A and honest B (i.e., distinguishing it from independent random value), as long as the attacker is not actively controlling or observing the secret choices, i.e., x in this case, made by A for that session. The resistance against KCI attacks of our IKE protocol can be easily checked, from the CDH assumption and the pseudorandomness of the key-derivation function $KDF(H(g^{xy}), \cdot)$.

Resistance against reflection attack. In some scenarios, a party may want to establish a secure channel with itself (i.e., $A = B$ in this case). For example, a mobile user that communicates to its desktop computer, while both the mobile device and the desktop have the same identity in the form of the same signature certificate $Cert$. In a *reflection* attack, an attacker simply copies (authentic) messages from A and sends them back to A as the messages coming from the other copy of A . Such reflection attack fails with our IKE protocol, just by the inclusion of (sid, g^x, g^y) and (sid, g^y, g^x) that indicates the role of initiator and responder even for copies of the same party.

A NOTE ON THE ROLE OF $NMZK_{(B,y)}$ IN PREVENTING REFLECTION ATTACKS. Under the security of building components of our deniable IKE, the only way for an adversary to amount a successful reflection attack is to manage to set $Y = X$, i.e., $(X, Y) = (Y, X)$ in this case. Note that this is an effective reflection attack against the variant of SIGMA (depicted in Appendix A) *when "0" and "1" are removed from SIG and MAC*. But, this play is prevented in our deniable IKE, by briefly noting that the adversary cannot provide proof-of-knowledge of $y = x$ (i.e., $NMZK_{(B,y)}$) in second-round message. This is another key role played by $NMZK_{(B,y)}$. If one interests in a simplified version of our deniable IKE without $NMZK_{(B,y)}$, then "0" and "1" should be put into H to explicitly indicate the roles of initiator and responder OR change the order (X^b, X^y) in $NMZK_{(b,y)}$ into (X^y, X^b) (i.e., change $NMZK_{(b,y)}$ into $NMZK_{(y,b)}$).

Robust resistance against unknown key share (UKS) attack (even with secret-key compromise). A successful UKS attack is one in which two parties compute the same session key but have

different views of who the peer to the exchange was (in other words, the attacker successfully managed to make two different sessions, among different pairs of players, to output the same session-key). The resistance against UKS attack of our IKE protocol is from its NMZK authentic binding nature: player's ID (A or B), its DH key component (g^x or g^y), the NMZK proofs of knowledge of its secret-key (i.e., a or b) and the discrete logarithm (i.e., x or y) of its DH key component (i.e., X or Y), and the specific session information ((sid, g^x, g^y) or (sid, g^y, g^x)), in our IKE protocol, are bounded *in a non-malleable authentic way* through $NMZK_{(B,y)}$ and $NMZK_{(a,x)}$ and $NMZK_{(b,y)}$, in the sense that they adhere to the specific session and cannot be malleated into/from other sessions. Note that, a successful UKS attack implies that such non-malleable authentic binding is broken, which in turn implies that at least one of the KEA assumption, the security of MAC, RO assumption of H and the CDH assumption, will be broken.

We remark that the resistance against UKS attack of our IKE protocol holds even for adversaries who managed to register a public-key that is identical to that of some honest player but with different name, e.g., an adversary can register (M, g^a) to the CA for certification where g^a is actually public-key of A . *This holds even if player's secret-key (e.g., a of A) is compromised (known) by the adversary, which provides much more robust and stronger resistance against UKS attacks.*

Resistance against leakage of ephemeral DH exponents. That is, the leakage of ephemeral DH exponent, x or y , should not affect the security of sessions where this exponent is not used. In particular, the leakage of ephemeral DH exponents, x or y , of one session should not help the attacker to impersonate A or B in other sessions. This property is from the observation: the NMZK authentic binding of one session, between two honest parties, holds even if the attacker learns some ephemeral DH exponents from other sessions. In particular, note that our IKE implies deniable authentication (i.e., deniable interactive public-key signatures) both for A and for B , and thus an attacker cannot impersonate A or B without knowing the corresponding long-term secret-key (unless the CDH assumption or the security of the underlying MAC is broken).

Protocol variants. One recommendable protocol variant is to put one's own ID (the input of MAC) additionally into the input of H . This variant is still provably secure in the post-specified-peer model, but seemingly provides extra stronger security guarantee in practice. Specifically, even the security of MAC is broken, an adversary is still infeasible to output valid NMZK proofs. For example, given $NMZK_{(B,y)} = MAC_{H(sid, B, g^y, g^x, g^{xy})}(B)$ (in this case), a valid $NMZK_{(M,y)}$ output by an attacker with $M \neq B$ will have to be $MAC_{H(sid, M, g^y, g^x, g^{xy})}(M)$. In other words, with this variant, previous MAC values sent by B with MAC-key $H(sid, B, g^y, g^x, g^{xy})$ provides no extra advantage to an attack in order to forge a valid MAC value with different name $M \neq B$, as in this case the attack has to use MAC-key $H(sid, M, g^y, g^x, g^{xy})$ that is however independent of $H(sid, B, g^y, g^x, g^{xy})$. A further simplification of the above variant is to remove MAC, and only use hash function H to make authentic binding. In practice, we are against such over-simplified variant, as the security will critically rely on the *collision-resistance* and *one-way hiding* of hash function in practice (which we do not want that). In addition, some public session-specific information, e.g., (sid, X, Y) and (sid, Y, X) could also be put into the input of MAC. When our IKE protocol works in the pre-specified-peer setting and/or players do not much care about privacy preserving, peer's ID can also be put into the input of H (and/or the input of MAC) and/or "0" and "1" can be put into H for explicitly indicating initiator and responder.

4 Deniable Internet Key-Exchange: the Aggressive Model

We now present the 3-round variant of our deniable IKE protocol in accordance with the aggressive model of IKE, which is depicted in Figure 2 (page 13). For presentation simplicity, the public system parameters, player's keys and the key-derivation function KDF are omitted here, which remain the same as in the deniable IKE of the main model (depicted in Figure 1).

Most security properties of the deniable IKE of the plain model are essentially inherited by the above deniable IKE of the aggressive model, except for the full deniability for player B . Specifically, player B

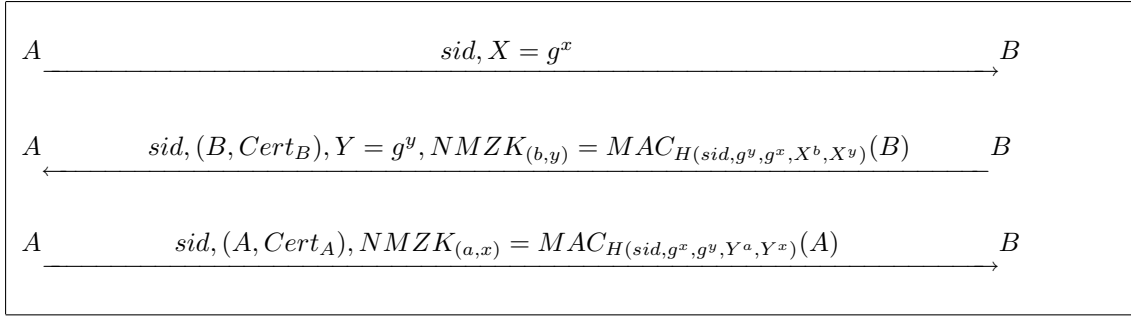


Figure 2: Deniable IKE (the aggressive model)

has only *completed-session deniability*, in the sense that if A^* completes the session then B 's deniability will be guaranteed. But if the (possibly malicious) A^* just aborts the session after receiving second-round message, then the deniability for B cannot be proved. Note that player A still has full deniability. We remark that completed-session deniability for responder is still very useful and reasonable. Consider the scenario that A is a client and B is a (bank or shop) server: in this case it is the client A who cares more about its privacy and full deniability does guarantee for it, while the server cares less about the deniability and a completed-session deniability is still deemed to be good enough for it.

5 Some Observations on the Independent SW-KE Protocol

The Stinson-Wu key-exchange (SW-KE) protocol of [28] is recalled in the left column of Figure 3. (In the description of [28], B is the initiator and A is the responder, but it does not pose any essential matter.) The public system parameters and player's keys (g^a for A and g^b for B) are identical with our deniable IKE protocol. We stress that our deniable IKE and authentication protocols are developed *fully independently* of [27, 28].

On key registration. As in our work, [28] also explicitly does *not* require player to prove knowledge of its secret-key, which is the common case with PKI in practice. Note that, as proof-of-knowledge complicates the already-complicated PKI, in most PKI standards, proof-of-knowledge is only optional rather than mandatory or standard. This, in particular, is of this case when we consider identification schemes within PKI. We remark that when we talk about identifications within PKI, it typically assumes no proof-of-knowledge/proof-of-possession during key-registration. Otherwise, in certain settings the simplest way to identify oneself might be just to show one's certificate (and we might not need to devise and run complicated identification protocols in this case in certain settings). Note that the goal of (mutual) identifications is very clearly and repeatedly claimed and stressed in [27, 28] (already started in Title, Abstract and Introduction). The formalizations of key-registration in [27, 28] clearly do not require proof-of-knowledge/proof-of-possession, and the security analysis implies this assumption.

SW-KE does not work in the post-specified-peer model. Note that the computation of the first-round message from A needs B 's public-key g^b , and the computation of the second-round message from B needs both A 's public-key g^a and B 's public-key g^b . We remark that mixing *both* parties' IDs in message computation violates the privacy-preserving principle of IKEv2, i.e., message from one party should not involve its peer's ID.

On limited deniability and privacy of SW-KE. B has no deniability over Internet (or against adversary with auxiliary input) *even for completed sessions*. Specifically, (X, g^{bx}) and/or $(X, H(g^{bx}))$ could just be collected over network (e.g., from the executions of other different protocols¹), or given by a judge, or eavesdrop from earlier executions of the SW-KE protocol possibly between B and other parties (rather than A). In this case, B cannot deny the second-round message *even for completed sessions*. Similarly, A also has no deniability against adversary with auxiliary input. Specifically, consider (Y, g^{ay}) is collected from network or given by a judge, then A cannot deny the sending of the

¹For example, an trusted party/organization identified itself using the protocol of [27].

third-round message $H(g^{ya})$.

With respect to the fairness concern of peer's privacy, note that for our deniable IKE, the authentic information $NMZK_{(a,x)}$ from A can be computed merely from y (without x, a, b), and all the authentic information from B , $NMZK_{(B,y)}$ and $NMZK_{(b,y)}$, can be merely computed from x (without y, a, b). In other words, messages from one party cannot be linked to its ID and its DH-component as well as its peer's ID. This is thus fair for the concern of privacy for both parties. The SW-KE protocol is clearly not of this case.

On provable concurrent security. We do not make sure the provable concurrent security of SW-KE, in view of the subtleties of KEA assumptions with public-keys in the concurrent sessions. As clarified in Section 3 for the dual role of (sid, X, Y) , for proving concurrent zero-knowledge (CZK), the simulator needs to extract x from the adversary impersonating A or extract y from the adversary impersonating B , by employing the KEA assumption. As KEA assumption is a non-black-box assumption, we want to extract, in straight-line way, the exponent from the local state of the adversary just at the point that it provides valid KEA answer (e.g., provides (X, g^{bx}) w.r.t. g^b). When SW-KE protocols are executed concurrently, note that the RO queries g^{bx} and g^{ay} w.r.t. players public-keys can be made at any arbitrary stages. Furthermore, the simulator cannot realize whether a valid KEA answer w.r.t some player's public-key has been provided by the adversary, as the adversary may make RO queries much earlier *before it reveals X or Y to the simulator* (and also there are numerous player's public-keys and concurrent sessions). This implies that, for successful and efficient CZK simulation, when the simulator receives some X or Y it may need to rewind the adversary to much earlier state to extract the corresponding exponents (otherwise, we may assume the adversary faithfully keeps all its past local states). But, rewinding causes exponential blow-up in simulation time. Such subtleties could also be applicable to the analysis of other concurrent security properties of SW-KE, as well as the protocol of [27]. The inclusion of $(sid, X, Y)/(sid, Y, X)$ can solve such subtleties. We remark that, at least, such subtleties were not mentioned or realized in [27, 28].

Dangerous UKS attack. Note that, as in the setup for our deniable IKE, the SW-KE protocol explicitly does *not* require proof-of-possession during parties' key-registrations with the CA (which is also the common case with public-key infrastructure in practice). Now, consider that an attacker \mathcal{A} registers to CA the public-key g^a (i.e., the public-key of the honest A) in the name of M . Then, the message flows of the UKS attack carried by \mathcal{A} is depicted in Figure 3.

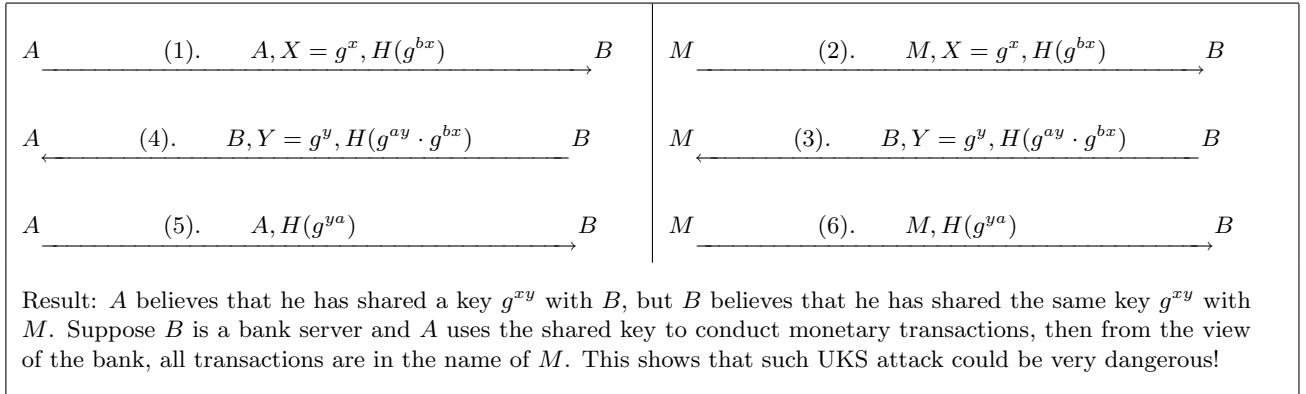


Figure 3: The UKS attack against SW-KE

A NOTE ON THE UKS ATTACK: We remark that even requiring proof-of-knowledge (of secret-key) during key-registration still does not fully prevent the above dangerous UKS attack. Specifically, such UKS attack can always be performed whenever the secret-key of A is compromised (known) by the adversary. This is in contrast to our deniable IKE, which resists UKS attacks even with secret-key compromise (as clarified in Section 3).

Effective reflection attacks. There are effective reflection attacks.

CONCURRENT MAN-IN-THE-MIDDLE ATTACK.

Note that, from an execution of SW-KE between A and itself, an adversary can collect $(X = g^x, H(g^{ax}))$ from the first-round message, and $(Y = g^y, H(g^{ya}))$ from the second and the third round messages. Denote the collected message by $(Z = g^z, H(g^{az}))$. Then, the adversary \mathcal{A} concurrently runs SW-KE in two sessions with A in the same name of A .

- \mathcal{A} initiates the first session with the honest A by playing the role of responder. After receiving the first-round message, $(X = g^x, H(g^{ax}))$, \mathcal{A} suspends the first session.
- \mathcal{A} initiates the second session with A by playing the role of initiator, and just sends to A the *collected* messages $(Z = g^z, H(g^{az}))$ as the first-round message. After receiving the second-round message from A , denoted by $(Y = g^y, h_2^2 = H(g^{ay} \cdot g^{az}) = H(g^{a(y+z)}))$, the adversary suspends the second session.
- \mathcal{A} continues the first session, and computes the second-round message, denoted (Y^1, h_2^1) , as follows: $Y^1 = Z \cdot Y \cdot X^{-1} = g^z \cdot g^y \cdot g^{-x} = g^{y+z} \cdot g^{-x}$ (where Z, Y are the messages of the second session, and X is in the first-round message of the first session), and $h_2^1 = h_2^2$. Note that, from the view of A , (Y^1, h_2^1) is a valid second-round message of the first session, as $H((Y^1)^a \cdot (g^a)^x) = H(g^{a(y+z)}) = h_2^2$. After receiving the third-round message from A , \mathcal{A} stops.

Consequences: the adversary successfully finishes the first session with A in the same name of A , BUT without knowing either the DH-exponent or the secret-key of A . Note that the transcript of the first session is significantly *different* than that of the second session. Also note that even proof-of-possession during key-registration does not prevent such reflection attack.

DIRECT REFLECTION ATTACKS. \mathcal{A} runs SW-KE with A in the same name of A , by playing the role of responder. After receiving the first-round message from A , $(X = g^x, H(g^{ax}))$, \mathcal{A} sends back $(Y = X^{-1}, h_2 = H(1))$ as the second-round message. Note that the second-round message is valid from the view of A , but \mathcal{A} knows neither A 's private key nor the DH-exponent.

It appears that explicitly checking $h_2 \neq H(1)$ can prevent above simple reflection attack. But, it does not. For example, \mathcal{A} forms the second-round message as follows: it selects $z \in \mathbb{Z}_q$, computes $(Y = X^{-1} \cdot g^z, h_2 = H(g^{az}))$. Note also that proof-of-possession during key-registration does not prevent such reflection attack.

Explicitly checking non-one of Y for identification of B . It is claimed in [28] that SW-KE is mutual identification protocol. Consider the following simple impersonating attack, in which an adversary \mathcal{A} impersonates B . After receiving the first-round message $(X, h_1 = H(g^{bx}))$ from A , \mathcal{A} sends back $(Y = 1, h_2 = H((g^a)^0 \cdot g^{xb}) = H(g^{bx}) = h_1)$. This is a valid second-round message without checking non-one of Y . This shows that checking non-one of Y is critical for identifying B as well as for key establishment. It seems that this attack was indeed not realized by Stinson and Wu [29], which is partially due to the separation of proof of DH-exponent and proof of secret-key (which seems to be not a very sound design principle). In SW-KE, players only abort when $X^q \neq 1$ or $Y^q \neq 1$. We briefly note that such attack does not work on our deniable IKE, *even without checking non-one of Y* .

On key confirmation. It is claimed that SW-KE is of authentic key confirmation (AKC) (in particular, a player, completing a session, must know the shared key). Suppose player A (possibly controlled by the adversary) collects some $(Z = g^z, H(g^{bz}))$ over Internet or eavesdrops from executions of SW-KW between B and another party other than A . Then, A sends $(Z, H(g^{bz}))$ as the first-round message, and finishes this session (by using its secret-key in the third-round). Note that A (controlled by adversary) does not know the shared secret-key.

A general weakness of SW-KE. Note that our deniable IKE guarantees that: if a player successfully completes a session with a (*possibly malicious*) peer, then the *possibly malicious* peer must know its DH exponent (i.e., knowing the shared session-key) and the secret-key corresponding to its alleged public-key. SW-KE loses this important security property in general. Consider the following attack: the adversary \mathcal{A} registers to CA with (M, g^{2a}) . That is, \mathcal{A} squares player's A 's public-key g^a , and

registers the malleated public-key g^{2a} with the ID M . Then, \mathcal{A} runs the SW-KE protocol with player A , in the name of M and public-key g^{2a} . After receiving $X = g^x$ from A in the first round, \mathcal{A} computes the second-round message as follows: it sets $Y = g^{-2x}$, and sends $H(1)$ to A . Note that, from the view of A , the second-round message from \mathcal{A} is valid, as $g^{-2x \cdot a} \cdot g^{2ax} = 1$. *BUT, \mathcal{A} knows neither the discrete logarithm of g^{2a} (i.e., the secret-key of its alleged public-key) nor the discrete logarithm (i.e., the DH-exponent $-2x$) of its DH-component Y* , which means SW-KE loses the authenticated key-confirmation property claimed in [28]. In general, for some polynomial-time computable functions $f(\cdot)$ and $g(\cdot)$, \mathcal{A} could potentially malleate its alleged public-key as $f(g^a)$ and its DH-component $Y = g(X)$, such that it can successfully provide second-round message without knowing the discrete logarithm of either $f(g^a)$ or $g(X)$. Such general attack works, even if player's ID and session-specific public information are put into the input of H .

A note: We would like to stress that the attacks and weakness identified in this section against SW-KE are mostly stemmed from the independent design experience of our deniable IKE by trying a rigorous provable security approach (actually, we still have not carefully investigated all (proof) details in [28]).

Some fixing suggestions of SW-KE. The above attacks and clarifications demonstrate the subtleties of obtaining correct and secure deniable key-exchange (we stress that our deniable IKE was developed independently of [27, 28]). One fixing suggestion is to combine $H(g^{bx})$ into the third-round as $H(g^{bx}, g^{ay})$, and change $H(g^{ay} \cdot g^{xb})$ of the second-round into $H(g^{ay}, g^{xb})$. In any case, we should put session-specific information into the input H and explicitly check non-one of Y . Specifically, for A to put (sid, A, B, g^x, g^y) and for B to put (sid, B, A, g^y, g^x) . This results in a more efficient and more secure KE protocol. If one insists using SW-KE, then proof-of-possession of secret-key should be critically required during key registration, and non-one of Y and X the value of $H(1)$ should also be explicitly checked. More details are given in the full version.

References

- [1] M. Bellare, R. Canetti and H. Krawczyk. Keying Hash Functions for Message Authentication. In *N. Koblitz (Ed.): Advances in Cryptology-Proceedings of CRYPTO 1996, LNCS 1109*, Springer-Verlag, 1996.
- [2] M. Bellare, R. Canetti and H. Krawczyk. A Modular Approach to the Design and Analysis of Authentication and Key Exchange Protocols. In *ACM Symposium on Theory of Computing*, pages 419-428, 1998.
- [3] M. Bellare and A. Palacio. The Knowledge-of-Exponent Assumptions and 3-Round Zero-Knowledge Protocols. In *M. Franklin (Ed.): Advances in Cryptology-Proceedings of CRYPTO 2004, LNCS 3152*, pages 273-289, Springer-Verlag, 2004.
- [4] M. Bellare and A. Palacio. Towards Plaintext-Aware Public-Key Encryption without Random Oracles. In *P. J. Lee (Ed.): Advances in Cryptology-Proceedings of Asiacrypt 2004, LNCS 3329*, pages 48-62, Springer-Verlag, 2004.
- [5] M. Bellare and P. Rogaway. Optimal Asymmetric Encryption. In *Advances in Cryptology-Proceedings of EUROCRYPT 1994, LNCS 950*, Springer-Verlag, 1994.
- [6] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *IEEE Symposium on Foundations of Computer Science*, pages 136-145, 2001.
- [7] R. Canetti and H. Krawczyk. Analysis of Key-Exchange Protocols and Their Use for Building Secure Channels. In *Advances in Cryptology-Proceedings of EUROCRYPT 2001, LNCS 2045*, Springer-Verlag, 2001. Available also from Cryptology ePrint Archive, Report No. 2001/040.
- [8] R. Canetti and H. Krawczyk. Security Analysis of IKE's Signature-Based Key-Exchange Protocol. In *M. Yung (Ed.): Advances in Cryptology-Proceedings of CRYPTO 2002, LNCS 2442*, pages 143-161, Springer-Verlag, 2002.
- [9] I. Damgård. Towards Practical Public-Key Systems Secure Against Chosen Ciphertext Attacks. In *J. Feigenbaum (Ed.): Advances in Cryptology-Proceedings of CRYPTO 1991, LNCS 576*, pages 445-456. Springer-Verlag, 1989.

- [10] A. Dent. Cramer-Shoup Encryption Scheme is Plaintext Aware in the Standard Model. In *Advances in Cryptology-Proceedings of EUROCRYPT 2006, LNCS 4004*, pages 289-307. Springer-Verlag, 2006.
- [11] M. Di Raimondo and R. Gennaro. New Approaches for Deniable Authentication. IN proc. of 12nd ACM Conference on Computer and Communications Security (ACM CCS'05), ACM Press, pages 112-121, 2005.
- [12] M. Di Raimondo, R. Gennaro and H. Krawczyk. Deniable Authentication and Key Exchange. To appear in ACM CCS'06. Full version appears in Cryptology ePrint Archive Report No. 2006/280.
- [13] D. Dolev, C. Dwork and M. Naor. Non-Malleable Cryptography. *SIAM Journal on Computing*, 30(2): 391-437, 2000. Preliminary version in *ACM Symposium on Theory of Computing*, pages 542-552, 1991.
- [14] C. Dwork, M. Naor and A. Sahai. Concurrent Zero-Knowledge. In *ACM Symposium on Theory of Computing*, pages 409-418, 1998.
- [15] S. Hada and T. Tanaka. On the Existence of 3-Round Zero-Knowledge Protocols. In *H. Krawczyk (Ed.): Advances in Cryptology-Proceedings of CRYPTO 1998, LNCS 1462*, pages 408-423, Springer-Verlag, 1998.
- [16] *D. Harkins and D. Carreal (Ed.): The Internet Key-Exchange (IKE)*, RFC 2409, Nov., 1998.
- [17] P. Karn and W.A. Simpson. The Photuris Session Key Management Protocol, draft-ietf-ipsec-photuris-03.txt, Sept., 1995.
- [18] C. Kaufman. Internet Key Exchange (IKEv2) Protocol. The Internet Engineering Task Force: INTERNET-DRAFT, October 2002.
- [19] S. Kent and R. Atkinson. Security Architecture for the Internet Protocol. Request for Comments 2401, 1998.
- [20] H. Krawczyk. SKEME: a Versatile Secure Key Exchange Mechanism for Internet In *proc. of 1996 IEEE Symposium on Network and Distributed System Security (SNDSS'96)*, pages 114-127.
- [21] H. Krawczyk. SIGMA: the "SIGn-and-MAC" Approach to Authenticated Diffie-Hellman and Its Use in the IKE-Protocols. Invited Talk at *D. Boneh (Ed.): Advances in Cryptology-Proceedings of CRYPTO 2003, LNCS 2729*, pages 400-425, Springer-Verlag, 2003.
- [22] H. Krawczyk. HMQV: A High-Performance Secure Diffie-Hellman Protocol. In *V. Shoup (Ed.): Advances in Cryptology-Proceedings of CRYPTO 2005, LNCS 3621*, pages 546-566. Springer-Verlag, 2005.
- [23] L. Law, A. Menezes, M. Qu, J. Solinas and S. Vanstone. An Efficient Protocol for Authenticated Key Agreement. *Designs, Codes and Cryptography*, 28: 119-134, 2003.
- [24] W. Mao. Modern Cryptography: Theory and Practice. Prentice Hall PTR, 2004.
- [25] A. Menezes. Another Look at HMQV. Cryptology ePrint Archive, Report No. 2005/205.
- [26] H. Orman. The OAKLEY Key Determination Protocol. Request for Comments 2412, Nov., 1998.
- [27] D. R. Stinson and J. Wu. An Efficient and Secure Two-Flow Zero-Knowledge Identification Protocol. Cryptology ePrint Archive, Report No. 2006/337, 16 Oct 2006.
- [28] D. R. Stinson and J. Wu. A Zero-Knowledge Identification and Key Agreement Protocol. Cryptology ePrint Archive, Report No. 2007/116, 30 Mar 2007.
- [29] Y. Zhao. Personal communications (in order to provide devoted clarification services to Stinson and Wu on Section 5). June 2007.
- [30] J. Zhou. Further Analysis of the Internet Key Exchange Protocol. *Computer Communications*, 23(2000): 1606-1612.

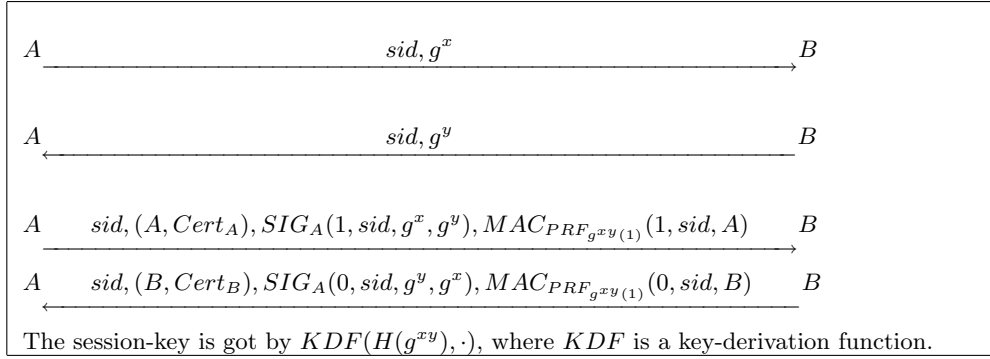


Figure 4: The SIGMA protocol (the main model)

A The SIGMA Protocol

SIGMA key-exchange protocol. We briefly recall the SIGMA protocol of [21] (the underlying basis of IKEv2), which is depicted in Figure 4 (page 18).

The “cutting-last-message” attack. The “cutting-last-message” attack works as follows [24]: A man-in-the-middle \mathcal{A} interacts with A in the name of M , while concurrently interacting with B in the name of A . M just relays messages between A and B , but aborts after receiving the last message from B . Such a simple attack results in authentication failure as follow: B is perfectly fooled to believe it has shared a session key with A , while A only thinks it ever took part in an aborted session with M . The above attack can be simply prevented by adding additional fifth-round of “acknowledgement” from A to B (as suggested in [8]), but increasing the system complexity.

On the role of “0” and “1” in SIGMA. We remark that including “0” and “1” in MAC and/or SIG in SIGMA, which explicitly indicate the roles of initiator and responder, plays a critical role in preventing reflection attacks. Specifically, for the variant of SIGMA when “0” and “1” are removed from MAC and SIG , there exists effective reflection attack: the adversary sets $Y = X$ and rebounds the third-round message back to A in the fourth-round in the same name of A . But, including “0” and “1” in MAC and/or SIG may (slightly) violate privacy protection of players’ roles. We note that our deniable IKE provides seemingly better privacy protection of players’ roles. Specifically, one cannot distinguish player’s role only from $NMZK_{(a,x)}$ or $NMZK_{(b,y)}$; the whole transcript may reveal the roles of initiator and responder, but the whole transcript can be totally deniable (which does not hold for SIGMA).

B Implication and Applications

Deniable message-authenticator.

THE BASIC VERSION. Our deniable IKE implies deniable message-authenticator: to authenticate a message m , A sends $\{sid, A, X = g^x, m\}$ to B in the first-round; in the second-round, B returns back $\{sid, B, Y = g^y, MAC_{H(sid, B, A, Y, X, X^y)}(m)\}$ (or $\{sid, B, Y = g^y, MAC_{H(sid, B, A, Y, X, X^y, g^{ay})}(m)\}$); A sends back $\{sid, A, MAC_{H(sid, A, B, X, Y, Y^a, Y^x)}(m)\}$ in the third-round.

SIMPLIFIED VERSION. In the first round, B sends $\{sid, B, Y = g^y, MAC_{H(sid, B, A, Y, g^a, g^{ay})}(B)\}$; and A responds back $\{sid, m, A, MAC_{H(sid, A, B, g^a, Y, Y^a)}(m)\}$.

NOTES: The variants discussed in Section 3 can also be applied to above deniable message authenticators. All above deniable message authenticator can also be slightly modified into identification protocols. For checking order q , though such check is mandatory for player’s public-key by CA during key-registration, in some settings and applications such checks on X and/or Y could be removed for more efficient implementations (though we do not recommend this).

On-line key-registration with proof-of-knowledge of secret-key. Let A be Certificate Authority with signature public-key PK_A and $X = g^x$, and signature secret-key SK_A , B wants to register

public-key g^b . In the first-round, A selects a random string r_A and sends $\{sid, A, r_A\}$ to B ; in the second-round, B returns back $\{sid, B, g^b, MAC_{H(sid, r_A, B, A, g^b, X, X^b)}(B)\}$; A sends back the certificate $\{sid, CERT_B = SIG_{SK_A}(r_A, B, g^b)\}$ in the third-round. (For simplicity, sid can be removed, and the random string r_A can play the role of sid .)

Actually, all above deniable message authenticator can be slightly modified into on-line key-registration with proof-of-knowledge. The basic idea is to replace Y by B 'S public-key g^b , and replace the last-round message from A by A 's signature on (sid, B, g^b) .

Formal formalization and analysis of this section are deferred to the full version.