

ANOTHER LOOK AT HMQV

ALFRED MENEZES

ABSTRACT. HMQV is a ‘hashed variant’ of the MQV key agreement protocol. It was recently introduced by Krawczyk, who claimed that HMQV has very significant advantages over MQV: (i) a security proof under reasonable assumptions in the (extended) Canetti-Krawczyk model for key exchange; and (ii) superior performance in some situations.

In this paper we demonstrate that HMQV is insecure by presenting realistic attacks in the Canetti-Krawczyk model that recover a victim’s static private key. We propose HMQV-1, a patched version of HMQV that resists our attacks (but does not have any performance advantages over MQV). We also identify the fallacies in the security proof for HMQV, critique the security model, and raise some questions about the assurances that proofs in this model can provide.

1. INTRODUCTION

The MQV protocols [26] are a family of efficient Diffie-Hellman authenticated protocols that have been widely standardized [2, 3, 17, 18, 32]. The two-pass protocol in this family (which is what the term “*the* MQV protocol” refers to) provides implicit key authentication. A three-pass variant adds key confirmation, thus providing explicit key authentication, while a one-pass variant is useful for applications such as email where only the sender may be online. The protocols were designed to meet a variety of security goals even in the presence of active attacks.

The security of the MQV protocols has been intensively studied since 1995. However, until recently there was no attempt to ‘prove’ that the protocols actually meet their security objectives. In a recent paper (see [24] and the expanded version [25]), Krawczyk undertook the ambitious task of trying to prove the security of the MQV protocols in the Canetti-Krawczyk model for key exchange [12, 13, 14]. This model provides a formal security definition by carefully specifying the capabilities and goals of an attacker.

Krawczyk’s analysis of MQV [24, 25] resulted in his concluding that “...MQV falls to a variety of attacks in [the Canetti-Krawczyk] model that invalidate its basic security as well as many of its stated security goals” and “...raises clear concerns about the security of the protocol...” He then proceeded to modify MQV to produce HMQV, a ‘hashed variant’ that “...provides the same superb performance and functionality of the original protocol but for which all the MQV security goals can be formally proved to hold in the random oracle model under the computational Diffie-Hellman assumption.”

HMQV is no less efficient than MQV because the additional hashing performed takes negligible time compared to the group exponentiations. In fact, HMQV is

significantly faster in some cases because it dispenses with the potentially expensive public key validations that are mandated in MQV. An important result stated in [24, 25] is that this performance improvement of HMQV is *proven* not to affect security in any way.

We show in §2 of this paper that the omission of public key validation in HMQV is a fatal security weakness. We accomplish this by presenting attacks that let an attacker recover a victim’s static private key. Our attacks on HMQV are more realistic and have much more damaging consequences than any of the proposed attacks on MQV in [24, 25]. Our attacks can be circumvented by requiring that all public keys be validated, leading to a new variant of HMQV that we call Patched HMQV (HMQV-1). As a consequence of reinstating public key validation, HMQV-1 no longer has any performance advantages over MQV.

In §3, we identify some flaws in the security proofs given in [25]. We critique some aspects of the Canetti-Krawczyk model, and discuss the difficulty of obtaining security reductions in this model that are meaningful in practice.

In §4 the purported attacks on MQV that were listed in [24, 25] are analyzed and critiqued. The paper concludes with some remarks in §5.

2. THE HMQV KEY AGREEMENT PROTOCOLS

Let G' be a multiplicatively-written abelian group of order N , and let $G = \langle g \rangle$ be a subgroup of G' of prime order q . We call G the *main group* and G' the *supergroup*. The integer $h = N/q$ is called the *cofactor*. We assume that $\gcd(h, q) = 1$, from which it follows that G is the unique subgroup of G' having order q . The fundamental security requirement that determines the suitability of a group for implementing Diffie-Hellman key agreement protocols is that the discrete logarithm problem (DLP) in G should be intractable. Groups arising from finite fields and from elliptic curves are the only ones presently in widespread use.

In finite field systems, G' is the multiplicative group \mathbb{Z}_p^* of a prime field \mathbb{Z}_p and G is the subgroup of order q for some prime divisor q of $p - 1$. As in DSA, the primes p and q are usually chosen so that the bitlength of q is considerably smaller than that of p . This is because the DLP in \mathbb{Z}_p^* can be solved using index-calculus methods whose running times are subexponential in p , while the best algorithms known for solving the DLP directly in G have running time $O(\sqrt{q})$. If p and q are selected in this way, then the cofactor h is much larger than q . For example, to achieve an 80-bit security level¹ against discrete logarithm attacks, one might take the bitlengths of p and q to be 1024 and 160, respectively; in this case the cofactor h has bitlength about 864.

In elliptic curve systems, G' is the group of points $E(L)$ on an elliptic curve E defined over a finite field L , and G is a subgroup of $E(L)$ of prime order q .² Elliptic curves of (almost) prime order can be easily found, so in practice the cofactor h is very small. For example, the 15 NIST-recommended elliptic curves [16] have $h = 1, 2$ or 4 .

2.1. Protocol descriptions. The following notation is from [25] and will be used throughout this paper.

¹A *k-bit security level* means that an attacker is expected to expend at least 2^k effort before it can solve a computational problem or defeat the security objectives of a protocol.

²The elliptic curve group law is usually written additively, but in this paper we shall use multiplicative notation for the sake of consistency.

\hat{A}, \hat{B}	Identities of two communicating parties.
a, b	Static private keys of \hat{A} and \hat{B} ; $a, b \in_R [1, q - 1]$.
A, B	Static public keys of \hat{A} and \hat{B} ; $A = g^a, B = g^b$.
x, y	Ephemeral private keys of \hat{A} and \hat{B} ; $x, y \in_R [1, q - 1]$.
X, Y	Ephemeral public keys of \hat{A} and \hat{B} ; $X = g^x, Y = g^y$.
H	A hash function.
\overline{H}	An l -bit hash function, where $l = (\lfloor \log_2 q \rfloor + 1)/2$.
MAC	Message authentication code algorithm.
K_m	MAC key.
K	Session key.

The identities \hat{A} and \hat{B} may also include certificates for their respective static public keys.

Description of HMQV. In (two-pass) HMQV, the parties exchange their identifiers and ephemeral public keys. That is, \hat{A} sends (\hat{A}, \hat{B}, X) to \hat{B} , while \hat{B} sends (\hat{B}, \hat{A}, Y) to \hat{A} . Then \hat{A} checks³ that $Y \neq 0$ and computes the *shared secret*

$$(1) \quad \sigma_{\hat{A}} = (YB^e)^{x+da},$$

while \hat{B} checks that $X \neq 0$ and computes

$$(2) \quad \sigma_{\hat{B}} = (XA^d)^{y+eb},$$

where

$$(3) \quad d = \overline{H}(X, \hat{B}) \quad \text{and} \quad e = \overline{H}(Y, \hat{A}).$$

Notice that $\sigma_{\hat{A}} = \sigma_{\hat{B}} = g^{(x+da)(y+eb)}$, and so both parties can compute the *session key* $K = H(\sigma_{\hat{A}}) = H(\sigma_{\hat{B}})$.

Description of three-pass HMQV. Three-pass HMQV adds key confirmation to the two-pass HMQV protocol. After receiving (\hat{A}, \hat{B}, X) from the initiator \hat{A} , the responder \hat{B} checks that $X \neq 0$, computes $\sigma_{\hat{B}}, K_m = H(\sigma_{\hat{B}}, 0), K = H(\sigma_{\hat{B}}, 1)$, and sends $(\hat{B}, \hat{A}, Y, \text{MAC}_{K_m}(\text{"1"}))$ to \hat{A} . Upon receipt of this message, \hat{A} checks that $Y \neq 0$, computes $\sigma_{\hat{A}}, K_m = H(\sigma_{\hat{A}}, 0), K = H(\sigma_{\hat{A}}, 1)$. \hat{A} verifies the MAC tag received, and sends $\text{MAC}_{K_m}(\text{"0"})$ to \hat{B} . Finally, \hat{B} verifies the latter MAC tag.

Description of one-pass HMQV. In one-pass key agreement protocols, only the sender can contribute an ephemeral public key. In one-pass HMQV, the sender \hat{A} transmits the single message (\hat{A}, \hat{B}, X) to \hat{B} . Let $d = \overline{H}(X, \hat{A}, \hat{B})$. The session key computed by \hat{A} is $K = H(B^{x+da})$, while \hat{B} computes the same key (after checking that $X \neq 0$) as $K = H((XA^d)^b)$.

2.2. The insecurity of HMQV. We present attacks on all three versions of HMQV described in §2.1. The adversary's actions in all these attacks are legitimate in that they are allowable within the Canetti-Krawczyk security model (cf. §3.1). The attacks have the most damaging consequences possible since the attacker is able to learn the victim's static private key.

Our attacks exploit the omission in HMQV of public key validation of static and ephemeral public keys. Namely, the receiver of an element X does not take any

³If G' is the multiplicative group of a finite field, then the element 0 in checks such as " $Y \neq 0$ " refers to the zero element of the finite field. However, since the HMQV protocol is described for generic cyclic groups G' , such checks are not well defined since it is not clear what is meant by the element "0".

steps to verify that X is indeed an element of order q in the main group G – the only required check is $X \neq 0$. Also, a certification authority who issues a certificate to party \hat{A} for its static public key A does not take any steps to verify that A is an element of order q in G . These omissions are by design rather than by accident. They are desirable because the parties and certification authority do not have to perform costly exponentiations by q to verify that a supergroup element is also in G , and also because some certification authorities may not be configured to do such tests [25]. In [25] proofs are given that HMQV is secure without these checks. Our attacks demonstrate that the security proofs must be fallacious. The errors in the proofs will be identified in §3.2.

Attack on one-pass HMQV. We first describe the attack on one-pass HMQV. The attacker \hat{A} , who wishes to learn \hat{B} 's static private key, selects an element $\gamma \in G'$ of order t for some small prime divisor $t \neq q$ of N . (Recall that N is the order of the supergroup G' .) Here by ‘small’ we mean that \hat{A} can feasibly perform t group operations and hash function evaluations. The attacker selects $a \in [1, t-1]$, and obtains a certificate for its static public key $A = \gamma^a$. She then selects $x \in [1, t-1]$, computes $X = \gamma^x$, and sends (\hat{A}, \hat{B}, X) to \hat{B} . Upon receipt, \hat{B} verifies that $X \neq 0$ and computes the session key $K = H((XA^d)^b) = H((\gamma^{x+ad})^b)$ where $d = \overline{H}(X, \hat{A}, \hat{B})$. Next, \hat{A} learns K by issuing a session-key query⁴ and computes $\beta = \gamma^{x+ad}$. \hat{A} then computes $K' = H(\beta^s)$ for $s = 0, 1, 2, \dots, t-1$ until $K' = K$. When the latter condition is met, \hat{A} knows that $s = b \bmod t$. After repeating this procedure for several different primes t , the value b can be easily determined by the Chinese Remainder Theorem.

Feasibility of the attack. We first note that session-key queries are allowable in the Canetti-Krawczyk model. However, in practice \hat{A} may not even have to issue a session-key query. For example, if \hat{B} were to subsequently send \hat{A} an authenticated message $(m, T = \text{MAC}_K(m))$, then \hat{A} could determine $b \bmod t$ by iteratively computing $K' = H(\beta^s)$ and $T' = \text{MAC}_{K'}(m)$ for $s = 0, 1, 2, \dots$ until $T' = T$. As before, \hat{A} then deduces that $s = b \bmod t$.

Is it reasonable to expect that N/q has several small prime factors? In the case of DSA-like parameters, the cofactor $h = (p-1)/q$ is large and thus one can expect it to have several small prime factors if p and q were randomly generated. However, if it happens that $(p-1)/q$ does not have several small prime factors, the extreme case occurring when $(p-1)/q$ is twice a prime, then the attack cannot be mounted.

In the case of elliptic curve groups, we have noted that the cofactor h is very small, and sometimes $h = 1$, so the attack as described above will generally fail. However, \hat{A} has the option of mounting an *invalid-curve attack* [4] which we describe next. Suppose that $G' = E(L)$ where E is an elliptic curve defined over a finite field L . The attacker selects an elliptic curve E' defined over L whose Weierstrass equation $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$ differs from E 's only in the coefficient a_6 , and such that the order of $E'(L)$ contains a prime factor t of the desired size. Then, as observed by Biehl, Meyer and Müller [9], the usual formulas for adding points in $E(L)$ do not involve a_6 and thus are identical to the formulas for adding

⁴A well-designed key agreement protocol should achieve its security goals even if an attacker is able to learn some session keys. This is because a key agreement protocol cannot guarantee that session keys won't be subsequently used in insecure applications (e.g., to encrypt messages with a weak symmetric-key encryption scheme).

points in $E'(L)$. If \hat{A} selects order- t points $X, A \in E'(L)$ and \hat{B} does not check whether $X, A \in E(L)$, then \hat{B} will successfully compute $(XA^d)^b \in E'(L)$ and the attack can proceed as before. Unlike in the case of DSA-like groups, the invalid-curve attack with elliptic curves is certain to succeed since there is an abundance of elliptic curves over L with group orders divisible by small primes.

We note that $E'(L)$ is not a supergroup of the main group of $E(L)$. One could perhaps argue that the HMQV specification implicitly required a recipient of an element X to check that it belongs to the supergroup $E(L)$, thus thwarting the attack. However, we maintain that it is reasonable to suppose that this check is *not* likely to be performed since points in $E'(L)$ are represented in the same way as points in $E(L)$,⁵ and since the HMQV specification [25] explicitly states that no checks are required other than $X \neq 0$.

Attack on two-pass HMQV. As before, the attacker \hat{A} selects an element $\gamma \in G'$ of order t , selects a static public key $A = \gamma^a$, and sends an ephemeral public key $X = \gamma^x$ to \hat{B} . While \hat{B} is computing the session key, \hat{A} issues a state-reveal query⁶ and learns \hat{B} 's ephemeral private key y . After \hat{B} has computed the session key $K = H((XA^d)^{y+eb})$, \hat{A} learns K via a session-key query. Now \hat{A} can obtain $s = b \bmod t$ by iteratively computing $K' = H((\gamma^{x+da})^{y+es})$ for $s = 0, 1, 2, \dots$ until $K' = K$.

This attack is perhaps not as realistic as the attack on one-pass HMQV since a state-reveal query is required. Nonetheless, it is allowable in the Canetti-Krawczyk model, and it is relevant to our analysis of the claims in [25], which include the author's claim to have *proven* resistance of two-pass HMQV to damage from leakage of ephemeral private keys.

Attack on three-pass HMQV. The attack on three-pass HMQV is similar to the one on two-pass HMQV except that a session-key query is not needed. This is because \hat{A} learns $\text{MAC}_{K_m}(\text{"1"})$ from \hat{B} 's response, and can use this tag to recognize a correct guess for $b \bmod t$.

2.3. Patched HMQV (HMQV-1). Our attacks on HMQV are thwarted if public keys are validated as mandated in MQV [26]. That is, the certification authority should verify that a party's static public key A is indeed an element of order q in the main group G before issuing a certificate. Similarly, the recipient of an ephemeral public key X should verify that X is an element of order q in G before proceeding with the next step of the protocol.

If the cofactor h is small, as is the case with elliptic curve groups, the computational cost of ephemeral public key validation can be reduced significantly by requiring only that the recipient \hat{B} of X verify that $X \in G'$ and that $\sigma_{\hat{B}}^h \neq 1$ before computing the session key $H(\sigma_{\hat{B}}^h)$.

3. CRITIQUE OF THE FORMAL ANALYSIS IN [25]

Formal analysis of a cryptographic protocol entails the following steps.

⁵For example, if $L = \mathbb{F}_{2^m}$, then a point in $E(L)$ or $E'(L)$ is represented by a pair of bitstrings of length m .

⁶State-reveal queries capture the possibility of an attacker learning some secret information such as ephemeral private keys that are specific to a particular session. Such queries do not reveal static private keys.

- (i) (*security model*) Providing a clear specification of the capabilities of an adversary and how it interacts with honest parties.
- (ii) (*security definition*) Providing a clear description of the goals of the adversary.
- (iii) (*security proof*) Proving that the only way an adversary can meet its goals is by solving a computational problem that is widely believed to be intractable.

The proof in step (iii) typically takes the form of a *reduction*, whereby one shows how a (hypothetical) adversary who succeeds in its task can be used to solve the hard computational problem with little additional effort.⁷

The declaration that a cryptographic protocol is ‘provably secure’ should always be met with a healthy dose of skepticism. First, one should verify that the security model and definition are the ‘right’ ones, i.e., they accurately capture the capabilities of attackers that one could expect in the environments in which the protocol will eventually be deployed. Second, the proof should be checked for correctness. This step is notoriously difficult since the proofs are typically long, complicated, and often poorly written. History has shown that subtle flaws may take years to be discovered. Indeed, Stern [33] has proposed adding a validation step to any security proof:

Also, the fact that proofs themselves need time to be validated through public discussion was somehow overlooked.

Finally, one should carefully consider what the proof actually means, i.e., what assurances are actually provided when the protocol is deployed. This step involves examining the reasonableness of the stated assumptions, and performing a *concrete security analysis*. The next paragraph elaborates on this last point.

About ten years ago, Bellare and Rogaway (see [5]) developed the notion of “practice-oriented provable security”. As a result of their work there was a greater awareness of the need to quantify the aforementioned ‘additional effort’ expended in a security reduction. The measure of this additional effort, obtained through a concrete security analysis, could then be used to provide concrete recommendations about keylengths. If the additional effort is relatively small, then the reduction is said to be *tight* and the security of the protocol is closely linked to the hardness of the related computational problem. For example, if there is a tight reduction from the DLP in a group G to the breaking of a protocol, then the protocol attains an 80-bit security level if the parameters for G are selected so that the best algorithms known for solving the DLP in G take about 2^{80} operations. On the other hand, a non-tight reduction only assures us that the protocol has an 80-bit security level if the parameters for G are selected so that the best DLP solvers require significantly more than 2^{80} operations. The latter assurance can be achieved only by using significantly larger group parameters, a direct consequence of which is slower performance.

In the remainder of this section we provide a critique of various aspects of the HMQV security proof given in [25].

⁷The effort includes the number of times the adversary is invoked; here we are thinking of the adversary as a computer program.

3.1. Security model and definition. In a series of papers, Canetti and Krawczyk [12, 13, 14] carefully developed security models and definitions for key establishment. Their model corrected shortcomings in previous attempts (e.g., [8, 10, 6, 31]), and at present is widely accepted as being the ‘right’ one.

In the Canetti-Krawczyk model, the attacker controls all communications between honest parties. Through ‘corrupt’, ‘session-key’ and ‘state-reveal’ queries, she is even able to learn the long-term static keys of some parties, learn session keys, and learn certain secret information that is associated with a particular session. Her goal is to learn something about a session key that she cannot deduce by trivial means (e.g. through a session-key query). The thesis is that the Canetti-Krawczyk model captures all realistic attacks in a open network such as the Internet.

However, one should be careful not to be lulled into a false sense of security since there may be desirable security attributes which are not adequately addressed by the model.

As an example, consider the case of key-compromise impersonation (KCI) attacks. In these attacks, first studied by Just and Vaudenay [19], the adversary learns a party \hat{A} ’s static private key and then tries to impersonate another honest party \hat{B} to \hat{A} . Resistance to KCI attacks is important in situations where an attacker wishes to obtain some information possessed by \hat{A} , who is only willing to divulge this information to \hat{B} (and where the attacker is not able to obtain \hat{B} ’s static private key). The Canetti-Krawczyk model did not cover resistance to KCI attacks, and hence protocols proved secure in the model have to be examined on a case-by-case basis for KCI resistance. For example, the ISO [12] and SIGMA [23] protocols would appear to resist KCI attacks, while the Boyd-Mao-Paterson [11] protocol does not. This deficiency in the Canetti-Krawczyk model has apparently been addressed by Krawczyk [25].

One should also be aware that the Canetti-Krawczyk model does not account for most kinds of secret-information leakage, e.g., partial information about the static private key that may be obtainable through side-channel attacks that analyze power consumption [22] or electromagnetic radiation [1]. An example of an attack on MQV (and HMQV-1) that is not considered in the Canetti-Krawczyk model is the side-channel attack of Leadbitter and Smart [27] whereby a static private key a can be deduced from the most significant bits of some ephemeral private keys x and associated secret values $(x + da) \bmod q$.

The state-reveal query is indeed useful for modeling some realistic scenarios such as the loss of ephemeral private keys. However, in the Canetti-Krawczyk model each item that is the result of some intermediate calculation involving secret data must be specified as either being in highly secure memory (whose contents are only accessible via a corrupt query), or belonging to less-secure memory (whose contents are accessible via either a state-reveal query or a corrupt query). This distinction is potentially misleading because implementors may be lulled into thinking that there is no possible risk in performing certain calculations in less-secure memory, and this may increase their exposure to side-channel and other attacks.

3.2. The proofs. The HMQV security proof in [25] has two steps. First, an ‘exponential challenge-response’ signature scheme XCR is defined and proven secure under the computational Diffie-Hellman (CDH) assumption. Second, the security of XCR (actually a ‘dual’ version of XCR) is proven to imply the security of HMQV.

The reduction in the first step is relatively straightforward, but the second reduction is extremely long and complicated.

In the XCR scheme, a verifier \hat{A} selects $x \in_R [0, q - 1]$ and sends the challenge $X = g^x$ and a message m to the signer \hat{B} . \hat{B} responds by selecting $y \in_R [0, q - 1]$ and sending the signature $(Y = g^y, \sigma = X^{y+eb})$ to \hat{A} where $e = \overline{H}(Y, m)$ and (B, b) is \hat{B} 's static key pair. The signature is accepted by \hat{A} provided that $Y \neq 0$ and

$$(4) \quad \sigma = (YB^e)^x.$$

XCR signatures are different from ordinary digital signatures. In particular, \hat{A} cannot convince a third party that \hat{B} generated a signature (Y, σ) for message m and challenge X because \hat{A} could have generated this signature herself.

The HCR signature scheme introduced in [25] is a slight variant of XCR where the second signature component is hashed; i.e., \hat{B} 's HCR signature is $(Y, \overline{H}(\sigma))$ instead of (Y, σ) . The HCR scheme is proven secure under the Gap Diffie-Hellman (GDH) and 'KEA1' assumptions even if the attacker learns the ephemeral secrets y for each Y . The security of HCR is then proven to imply the security of HMQV even if the attacker is able to deduce ephemeral private keys x and y via state-reveal queries.

The security proofs for XCR and HCR both have flaws, which explains why our attacks in §2.2 could be successfully mounted on 'provably secure' protocols. The flaws arise because neither XCR nor HCR perform any validation of static or ephemeral public keys (other than the check $Y \neq 0$). We identify these flaws next and then present attacks on XCR and HCR.

Flaw in the XCR proof. In the reduction of CDH to XCR forgeries (Figure 4 of [25]), the simulator \mathcal{C} responds to an XCR-forgery \mathcal{F} 's query (X, m) with $(\overline{Y} = g^s/B^e, \overline{\sigma} = X^s)$ where $s \in_R [0, q - 1]$ and $e \in_R \{0, 1\}^l$. However, this response is in general *not* a valid one if $X \notin G$. To see this, suppose that the order of X is t , where $t \neq q$ is prime. Then $\overline{\sigma}$ is in a subgroup of order t , whereas the right side of (4) is in a subgroup of order q . Since t and q are distinct primes, the verification condition (4) fails unless both sides are the identity. Hence \mathcal{C} 's response to \mathcal{F} 's query is not correct, and consequently the subsequent behaviour of \mathcal{F} cannot be predicted.

Flaw in the HCR proof. Similarly, in the reduction of GDH to HCR forgeries (Figure 6 of [25]), the simulator \mathcal{C} cannot respond to a challenge (X, m) if $X \notin G$. This is because it needs to decide whether $\text{CDH}(X, B) = (\sigma/X^y)^{1/e}$ for certain $\sigma \in G$, but its DDH-oracle assumes that the inputs X, B are in G – indeed $\text{CDH}(X, B)$ is not even defined if $X \notin G$ and $B \in G$.

Attacks on XCR and HCR. We observe that the XCR and HCR schemes are insecure if the signer does not verify that challenges X are in the main group G and if the attacker can learn ephemeral private keys y .⁸ The attacker would select a challenge $X \in G'$ of order $t \neq q$. Since \hat{B} responds with (Y, X^{y+eb}) or $(Y, \overline{H}(X^{y+eb}))$, the attacker knowing y can compute $b \bmod t$ by exhaustive search. These particular attacks on XCR and HCR can be defeated by performing validation on challenges X .

⁸This answers the question posed in Remark 4.5 of [25] about the security of XCR in the face of ephemeral private key revelations.

Security of HMQV-1. As noted in §2.3, the patched HMQV-1 is resistant to our attacks because key validation is performed. Furthermore, the flaws in the HMQV proof in [25] that were identified above are no longer pertinent. It remains to be seen whether the remainder of the proof is correct, and whether new reductions can be found that are tight enough to provide meaningful assurances about the security of HMQV-1.

3.3. Interpreting the proofs. Previous security proofs for key agreement protocols in the Canetti-Krawczyk model (e.g., see [12, 13, 14]) are all missing a concrete security analysis. This omission is partly because the analysis would be quite tedious since it would have to track several parameters including the number of honest parties, the number of sessions a party can be engaged in, and the number of hash function queries if the random oracle assumption is being invoked. Another likely reason for the omission is that the reductions are not at all tight, and revealing this deficiency would call into question the practical value of the security proofs.

The HMQV security proof in [25] does not include a concrete security analysis. However, such an analysis is presented for the security proof of the XCR signature scheme. (Recall that this is the first step of the HMQV security proof.) The XCR security proof takes place in the random oracle model, and the forking lemma [29] is invoked. The result in [25], based on the analysis of the forking lemma in [29], is that the existence of an attacker who works in time at most T , presents at most Q queries to an l -bit hash function \overline{H} , and produces a forgery with probability at least ϵ , implies the existence of a CDH solver that runs in expected time

$$(5) \quad T' \leq 2^{17}QT/\epsilon$$

provided that $\epsilon \geq 7Q/2^l$.

The following analysis is from [21, Section 5.5]. Suppose for simplicity that $\epsilon \approx 1$. There is no justification for assuming that the number of hash function queries the forger makes is bounded by anything other than the forger's running time T . Thus (5) simplifies to $T' \leq 2^{17}T^2$. For the remaining analysis, we take $T' \approx 2^{17}T^2$. Since Pollard's rho method for finding discrete logarithms in the main group G has expected running time \sqrt{q} , the security reduction is only meaningful if $T' \approx 2^{17}T^2 \ll \sqrt{q}$. Suppose now that the main group G is an order- q subgroup of the multiplicative group \mathbb{Z}_p^* . Then for XCR to be assured an 80-bit security level one would have to select roughly a 354-bit q and a 7000-bit p (see [28]). Such parameter sizes would be too inefficient to be used in practice. On the other hand, if we insist on efficiency and use 1024-bit p and 160-bit q , then the security proof only guarantees that forgers whose running times are less than $2^{31.5}$ do not exist; this is a totally useless level of security.

If we analyze the extremely non-tight HMQV security proof, we find that the assurances it provides are even worse than those that follow from the non-tight XCR proof.

We conclude this section with a quote from [21, Section 5.5] that is very relevant to our discussion of concrete security analysis.

Unfortunately, this type of analysis is generally missing from papers that argue for a new protocol on the basis of a “proof” of its security. Typically, authors of such papers trumpet the advantage that their protocol has over competing ones that lack a proof of security (or that have a proof of security only in the random oracle model),

then give a non-tight reductionist argument, and at the end give key-length recommendations that would make sense if their proof had been tight. They fail to inform the potential users of their protocol of the true security level that is guaranteed by the “proof” if, say, a 1024-bit prime is used. It seems to us that cryptographers should be consistent. If one really believes that reductionist security arguments are very important, then one should give recommendations for parameter sizes based on an honest analysis of the security argument, even if it means admitting that efficiency must be sacrificed.

4. THE MQV KEY AGREEMENT PROTOCOLS

Section 3.2 of [25], entitled ‘Insecurity of MQV’, examined the MQV protocol in the Canetti-Krawczyk model to determine whether the protocol possesses the security properties that were informally stated in [26]. Seven weaknesses and explicit attacks were identified to demonstrate that some of the most essential security properties were not satisfied in general. We briefly revisit the seven attacks and explain why they are irrelevant in practice. We begin by reviewing the MQV protocols.

The (two-pass) MQV protocol differs from HMQV in the following ways. For simplicity, we only describe MQV in the case where G' is the group of points on an elliptic curve and the cofactor h is small; some details are omitted for the sake of brevity.

- (i) The certification authority must verify that static public keys belong to the main group G .
- (ii) The recipient of an ephemeral public key X must verify that $X \in G'$, check that $\sigma^h \neq 1$, and compute the session key as $K = H(\sigma^h)$.
- (iii) The integers d and e are derived from the l least significant bits of the x -coordinates of X and Y respectively.

In three-pass MQV, the MAC tags transmitted are $\text{MAC}_{K_m}(\text{“2”}, \hat{B}, \hat{A}, Y, X)$ and $\text{MAC}_{K_m}(\text{“3”}, \hat{A}, \hat{B}, X, Y)$.

The seven weaknesses and explicit attacks identified in [24, 25] are the following:

(1) *Resistance to basic impersonation attacks.* The observation is made that since the calculation of $\sigma = g^{(x+da)(y+eb)}$ in MQV depends on the representation chosen for elements of G , some implementations of MQV may be weak because of a poor selection of representation. For example, if a representation is chosen where the l least significant bits of all x -coordinates of elliptic curve points are constant, then MQV falls to impersonation attacks – an attacker who selects x and sends $(\hat{A}, \hat{B}, g^x/g^{da})$ to \hat{B} , can easily compute the resulting session key $K = H((g^x)^{y+eb})$.

Of course, such contrived group representations would never arise in practice (and are certainly not allowed in standards such as [2, 3, 32] that specify the bit representations of group elements). The observation does illustrate that MQV cannot be proven secure for generic prime-order groups, but this can hardly be called an ‘attack’.

(2) *Prime-order testing and exponentiation performance.* This is not an attack, but rather the observation that the validation of ephemeral public keys may require an expensive exponentiation.

(3) *Validation of long-term public keys.* This too is not an attack, but the observation that validation of static public keys imposes an extra requirement on certification authorities.

HMQV was designed to address the perceived weaknesses in MQV raised in the first three points. That is, HMQV was described using generic prime-order groups (and so it isn't clear whether checks $X \in G'$ have to be performed), and validation of ephemeral and static public keys was omitted. Ironically, these three changes are *precisely* what led to the attacks on HMQV that were described in §2.2.

(4) *Resistance to key-compromise impersonation attacks.* Both MQV and HMQV are shown to fall to KCI attacks provided that the attacker, in addition to knowing \hat{A} 's static private key, can also learn the shared secret $\sigma_{\hat{B}}$ (by issuing a state-reveal query). This observation is used to justify the design principle of deriving the session key by hashing the shared secret (which both MQV and HMQV do), and for advising implementors to carefully protect and delete ephemeral information such as the σ values (i.e., the attacker should not be able to get any σ values via state-reveal queries).

(5) *Resistance to disclosure of Diffie-Hellman exponents.* Another reason for protecting σ values is that both MQV and HMQV fall to certain attacks if the adversary can obtain one of the ephemeral private keys x or y in addition to σ . No attacks are known on MQV or HMQV if only ephemeral private keys (and not σ values) are divulged.

(6) *Resistance to unknown-key share (UKS) attacks.* In these attacks, first discovered by Diffie, van Oorschot and Wiener [15], an adversary interferes with \hat{A} 's and \hat{B} 's communications. The result is that \hat{A} and \hat{B} still compute the same session key. However, even though \hat{A} correctly believes the key is shared with \hat{B} , \hat{B} mistakenly believes that the key is shared with a third party $\hat{C} \notin \{\hat{A}, \hat{B}\}$. Kaliski [20] observed that the two-pass MQV protocol is vulnerable to a UKS attack. One countermeasure, which is mandated in [32], is to include the identities of parties in the key derivation function, i.e., $K = H(\sigma, \hat{A}, \hat{B})$. Another countermeasure that was noted in [26] is to use the three-pass MQV protocol. The attack is thwarted in this case because party \hat{B} would send a tag $\text{MAC}_{K_m}(z, \hat{B}, \hat{C}, Y, X)$, whereas \hat{A} would expect a tag $\text{MAC}_{K_m}(z', \hat{B}, \hat{A}, Y', X')$. The trivial observation made in [24, 25] is that if an adversary could learn the MAC key K_m via a state-reveal query, then she could compute the appropriate tags and complete Kaliski's UKS attack.

(7) *Perfect forward secrecy.* It was observed in [24, 25] that no two-pass key agreement protocol can achieve 'full' forward secrecy. This observation is thus not specific to MQV, and in any case was already well known (e.g. see [7, 30]).

Summary. Of the seven weaknesses and explicit attacks on MQV reported in [24, 25], the first three are non-attacks and in fact were incorrectly addressed in HMQV, leading to a fatal flaw. Attacks (4), (5) and (6) rely on the adversary's ability to learn secret information (other than ephemeral private keys) through a state-reveal query; if this is a realistic threat in practice, then even HMQV-1 is insecure. Finally, attack (7) is a well-known observation about the limited form of perfect forward secrecy that can be attained in any two-pass protocol. We conclude that the analysis performed in [24, 25] did not uncover any new weaknesses in MQV.

5. CONCLUDING REMARKS

The work of Canetti and Krawczyk [12, 13, 14, 23, 24] is widely acclaimed for its development of a formal model and definition of secure key exchange. However, the Canetti-Krawczyk model does have some deficiencies. On the one hand, an attack that is legitimate in the model may not be a realistic threat in practice. On the other hand, the model does not account for many realistic threats and often yields proofs that are lengthy, open to misinterpretation, and of questionable value in practice.

Researchers who use the Canetti-Krawczyk model to analyze the security of key establishment protocols should be careful not to rely exclusively on this analysis, and in particular should not abandon old-fashioned cryptanalysis and prudent security engineering practices. Standards committees and security engineers should not be mesmerized by claims of provable security, nor should they be intimidated by technical proofs.

Above all, the results of the present paper validate the main thesis of [21] that the field of provable security is as much an art as a science.

ACKNOWLEDGEMENTS

I would like to thank Darrel Hankerson, Ann Hibner Koblitz, Neal Koblitz and Scott Vanstone for their valuable comments.

REFERENCES

- [1] D. Agrawal, B. Archambeault, J. Rao and P. Rohatgi, “The EM side-channel(s)”, *Cryptographic Hardware and Embedded Systems – CHES 2002*, Lecture Notes in Computer Science, 2523 (2002), 29-45.
- [2] ANSI X9.42, *Public Key Cryptography for the Financial Services Industry: Agreement of Symmetric Keys Using Discrete Logarithm Cryptography*, American National Standards Institute, 2003.
- [3] ANSI X9.63, *Public Key Cryptography for the Financial Services Industry: Key Agreement and Key Transport Using Elliptic Curve Cryptography*, American National Standards Institute, 2001.
- [4] A. Antipa, D. Brown, A. Menezes, R. Struik and S. Vanstone, “Validation of elliptic curve public keys”, *Public Key Cryptography – PKC 2003*, Lecture Notes in Computer Science, 2567 (2003), 211-223.
- [5] M. Bellare, “Practice-oriented provable-security”, *Proceedings of the First International Workshop on Information Security – ISW ’97*, Lecture Notes in Computer Science, 1396 (1998), 781-793.
- [6] M. Bellare, R. Canetti and H. Krawczyk, “A modular approach to the design and analysis of authentication and key exchange protocols”, *Proceedings of the 30th Annual ACM Symposium on the Theory of Computing*, 1998.
- [7] M. Bellare, D. Pointcheval and P. Rogaway, “Authenticated key exchange secure against dictionary attacks”, *Advances in Cryptology – EUROCRYPT 2000*, Lecture Notes in Computer Science, 1807 (2000), 139-155.
- [8] M. Bellare and P. Rogaway, “Entity authentication and key distribution”, *Advances in Cryptology – CRYPTO ’93*, Lecture Notes in Computer Science, 773 (1994), 232-249.
- [9] I. Biehl, B. Meyer and V. Müller, “Differential fault analysis on elliptic curve cryptosystems”, *Advances in Cryptology – CRYPTO 2000*, Lecture Notes in Computer Science, 1880 (2000), 131-146.
- [10] S. Blake-Wilson, D. Johnson and A. Menezes, “Key agreement protocols and their security analysis”, *Proceedings of the Sixth IMA International Conference on Cryptography and Coding*, Lecture Notes in Computer Science, 1355 (1997), 30-45.

- [11] C. Boyd, W. Mao and K. Paterson, “Key agreement using statically keyed authenticators”, *Applied Cryptography and Network Security – ACNS 2004*, Lecture Notes in Computer Science, 3089 (2004), 248-262.
- [12] R. Canetti and H. Krawczyk, “Analysis of key-exchange protocols and their use for building secure channels”, *Advances in Cryptology – EUROCRYPT 2001*, Lecture Notes in Computer Science, 2045 (2001), 453-474. Full version available at <http://eprint.iacr.org/2001/040/>.
- [13] R. Canetti and H. Krawczyk, “Universally composable notions of key exchange and secure channels”, *Advances in Cryptology – EUROCRYPT 2002*, Lecture Notes in Computer Science, 2332 (2002), 337-251. Full version available at <http://eprint.iacr.org/2002/059/>.
- [14] R. Canetti and H. Krawczyk, “Security analysis of IKE’s signature-based key-exchange protocol”, *Advances in Cryptology – CRYPTO 2002*, Lecture Notes in Computer Science, 2442 (2002), 143-161. Full version available at <http://eprint.iacr.org/2002/120/>.
- [15] W. Diffie, P. van Oorschot and M. Wiener, “Authentication and authenticated key exchanges”, *Designs, Codes and Cryptography*, 2 (1992), 107-125.
- [16] FIPS 186-2, *Digital Signature Standard (DSS)*, Federal Information Processing Standards Publication 186-2, National Institute of Standards and Technology, 2000.
- [17] IEEE Std 1363-2000, *Standard Specifications for Public-Key Cryptography*, 2000.
- [18] ISO/IEC 15946-3, *Information Technology – Security Techniques – Cryptographic Techniques Based on Elliptic Curves – Part 3: Key Establishment*, 2002.
- [19] M. Just and S. Vaudenay, “Authenticated multi-party key agreement”, *Advances in Cryptology – ASIACRYPT ’96*, Lecture Notes in Computer Science, 1163 (1996), 36-49.
- [20] B. Kaliski, “An unknown key-share attack on the MQV key agreement protocol”, *ACM Transactions on Information and System Security*, 4 (2001), 275-288.
- [21] N. Koblitz and A. Menezes, “Another look at ‘provable security’”, *Journal of Cryptology*, to appear. Available at <http://eprint.iacr.org/2004/152/>.
- [22] P. Kocher, J. Jaffe and B. Jun, “Differential power analysis”, *Advances in Cryptology – CRYPTO ’99*, Lecture Notes in Computer Science, 1666 (1999), 388-397.
- [23] H. Krawczyk, “SIGMA: The ‘SiGn-and-MAC’ approach to authenticated Diffie-Hellman and its use in the IKE protocols”, *Advances in Cryptology – CRYPTO 2003*, Lecture Notes in Computer Science, 2729 (2003), 400-425.
- [24] H. Krawczyk, “HMQV: A high-performance secure Diffie-Hellman protocol”, *Advances in Cryptology – CRYPTO 2005*, Lecture Notes in Computer Science, to appear.
- [25] H. Krawczyk, “HMQV: A high-performance secure Diffie-Hellman protocol”, Full version of [24], available at <http://eprint.iacr.org/2005/176/>.
- [26] L. Law, A. Menezes, M. Qu, J. Solinas and S. Vanstone, “An efficient protocol for authenticated key agreement”, *Designs, Codes and Cryptography*, 28 (2003), 119-134.
- [27] P. Leadbitter and N. Smart, “Analysis of the insecurity of ECMQV with partially known nonces”, *Information Security – ISC 2003*, Lecture Notes in Computer Science, 2851 (2003), 240-251.
- [28] A. Lenstra and E. Verheul, “Selecting cryptographic key sizes”, *Journal of Cryptology*, 14 (2001), 255-293.
- [29] D. Pointcheval and J. Stern, “Security arguments for digital signatures and blind signatures”, *Journal of Cryptology*, 13 (2000), 361-396.
- [30] P. Rogaway, M. Bellare and D. Boneh, “Evaluation of security level of cryptography: ECMQVS (from SEC 1)”, CRYPTREC report, Information-technology Promotion Agency, Japan, January 2001.
- [31] V. Shoup, “On formal models for secure key exchange”, version 4, 1999. Available at <http://shoup.net/papers/skey.pdf>.
- [32] SP 800-56 (DRAFT), *Special Publication 800-56, Recommendation on Key Establishment Schemes*, National Institute of Standards and Technology, January 2003.
- [33] J. Stern, “Why provable security matters”, *Advances in Cryptology – EUROCRYPT 2003*, Lecture Notes in Computer Science, 2656 (2003), 449-461.

DEPARTMENT OF COMBINATORICS & OPTIMIZATION, UNIVERSITY OF WATERLOO, WATERLOO, ONTARIO N2L 3G1 CANADA

E-mail address: ajmeneze@uwaterloo.ca