

Evaluation Report on the **HDEF-ECDH Cryptosystem**

1 Introduction

This document is an evaluation of the **HDEF-ECDH Cryptosystem**. Our work is based on the analysis of documents [15, 16]. The present report is organized as follows: firstly, we briefly review the cryptosystem; next we discuss the security level of the cryptographic primitive which underlies the scheme and analyze its relation to the difficulty of the discrete logarithm problem on elliptic curves; finally, we evaluate the security level of the scheme itself in the light of strong security notions similar to semantic security and security against adaptive chosen-ciphertext attacks. This is as requested by IPA.

2 Brief description of the scheme

2.1 Specification review

HDEF-ECDH is based on the hardness of the discrete logarithm problem over an elliptic curve. Let k be a given integer (typically $k = 160$, or 192 or 224 .) The cryptosystem uses elliptic curve E over some prime p , $|p| = k$, such that E has a prescribed number of points $p - 2$ (also a prime). In other words, the trace of E equals three. This feature is claimed to be beneficial in terms of security. Once the curve has been chosen, a base point G is randomly chosen on E .

The basic function f on which HDEF-ECDH is based is defined by

$$\begin{aligned} f : \{0, \dots, p - 2\} &\longrightarrow E \\ r &\longmapsto r \cdot G \end{aligned}$$

where $r \cdot G$ is obtained, by means of the usual elliptic curve addition, as the sum of r times G . Inverting f is precisely the elliptic curve discrete logarithm problem (ECDLP). Clearly, f is one-to-one. The inverse function, denoted \log_G , is believed

to be hard to compute. Another function used by the scheme is the Diffie-Hellman function:

$$\begin{aligned} DH : E^2 &\longrightarrow E \\ X, Y &\longmapsto \log_G(Y) \cdot X = \log_G(X) \cdot Y \end{aligned}$$

HDEF-ECDH is a key-agreement scheme: entities A that enter the scheme are given a curve E_A , as described above, a base point G_A and an element Y_A of E_A together with its discrete logarithm $x_A = \log_{G_A} Y_A$. Key agreement is achieved in two different settings

1. if two entities A and B have common domain parameters, $E_A = E_B$ and $G_A = G_B$, each can derive a shared key from $DH(Y_A, Y_B)$; for example, they can just take the first coordinate of $DH(Y_A, Y_B)$
2. otherwise, A generates a random $r_A \in \{0, \dots, p_B - 2\}$ and sends $R_B = r_A \cdot G_B$ to B ; similarly, B generates a random $r_B \in \{0, \dots, p_A - 2\}$ and sends $R_A = r_B \cdot G_A$ to A ; once this has been done, each of A and B can derive the two values $DH_{E_A}(Y_A, R_A)$ and $DH_{E_B}(Y_B, R_B)$. From these values, they can obtain a shared key; for example, they can take the exclusive or of the respective x coordinates of $DH_{E_A}(Y_A, R_A)$ and $DH_{E_B}(Y_B, R_B)$.

At this point, it is useful to introduce a more formal framework, that will be useful when we later perform the security analysis. A key-agreement scheme on a message space \mathcal{M} consists of three algorithms $(\mathcal{K}, \mathcal{E}, \mathcal{D})$:

- the key generation algorithm $\mathcal{K}(1^k)$ outputs a random pair of secret-public keys (sk, pk) , relatively to a security parameter k
- the key agreement algorithm \mathcal{E} is an interactive algorithm between two entities A, B , each endowed with a pair $(\text{sk}_A, \text{pk}_A)$ (and $(\text{sk}_B, \text{pk}_B)$ resp.) and each using random coins $r_A, r_B \in \Omega$
- the key derivation algorithm $\mathcal{D}_{\text{sk}}(D; r)$ outputs the common key from the data Δ exchanged during the key agreement algorithm and the random coins. It should be the case that $\mathcal{D}_{\text{sk}_A}(\Delta; r_A) = \mathcal{D}_{\text{sk}_B}(\Delta; r_B)$

The key generation algorithm $\mathcal{K}(1^k)$ of the HDEF-ECDH Cryptosystem produces a prime number p , $|p| = k$, and an elliptic curve over \mathbb{Z}_p such that E has a prescribed number of points $p - 2$ (also a prime). It also generates a base point G , which lies on the curve E . The equation of the curve is

$$y^2 = x^3 + ax + b$$

The public key pk includes the triple (p, E, G) and thus defines the above function f . A final element $Y = f(x) = x \cdot G$ is generated and the full public key is the tuple

(p, E, G, Y) . The secret key \mathbf{sk} is x , the discrete logarithm of Y in base G . The key-agreement algorithm \mathcal{E} is played between two entities, each having a public/secret key pair, as explained earlier in this section. The key derivation algorithm \mathcal{D} has also been explicitly described above.

2.2 Comments on the specification

Document [15] is rather detailed on the reason why the authors chose to work with elliptic curves of trace 3 and includes a description of how to generate such curves. On the other hand, it is rather vague in the implementation of the key-agreement itself. In particular, the submission only briefly discusses the authentication problems surrounding the Diffie-Hellman key-agreement protocol, on which the scheme builds. On page 7 of the self-evaluation report (document [16]), the authors state:

Adversary may impersonate B against A and impersonate A against B , so-called intruder-in-the-middle-attack. To prevent such attack, users must confirm the authenticity of the public key of the other party.

This is correct. However, in the context of Diffie-Hellman key-agreement, the fact that the public key is authenticated - say by the signature of a certification authority - is not always sufficient. When a fresh session key is established each time the protocol is executed, a security property usually expected is that, even if some previous session keys have been misused or corrupted, the new key should still be indistinguishable for the adversary. Furthermore, the property of *forward-secrecy* is also often required. It means that, even if the long-lived secret key of a participant is compromised, this should not help the attacker to learn anything about the previous values of the agreed keys. These questions are not addressed in the self-evaluation report. Thus the proposed schemes may not achieve the corresponding security requirement.

On a more general level, one must mention that the security analysis is missing in documents [15, 16]. No attempt is made to provide a formal treatment of security and to precisely relate the scheme with a specified hard problem about elliptic curves.

3 Security level of the cryptographic primitive

In this section, we investigate the security of the underlying cryptographic primitive, both in terms of complexity-theoretic reductions and with respect to the recommended parameters.

3.1 Complexity-theoretic arguments

As previously mentioned, document [15] does not attempt to measure the security of the scheme in terms of a hard problem related to the discrete logarithm for elliptic

curves. There are several basic primitives that can be considered.

3.1.1 The elliptic curve decisional and computational Diffie-Hellman hypotheses

We keep the notations of section 2.1. Recall that the decisional Diffie-Hellman hypothesis on an elliptic curve E , with a large subgroup of prime order, asserts that it is hard to distinguish the distributions \mathbf{D}_E and \mathbf{R}_E , where

$$\mathbf{R}_E = \{(G_1, G_2, U_1, U_2)\}$$

with all four elements taken at random in the large subgroup and

$$\mathbf{D}_E = \{(G_1, G_2, U_1, U_2)\}$$

with $\log_{G_1}(U_1) = \log_{G_2}(U_2)$. A quantitative version measures the maximum advantage $\text{AdvDDH}(t)$ of a statistical test T that runs in time t . This means the maximum of the difference of the respective probabilities that T outputs 1, when probabilities are taken over \mathbf{D}_E or \mathbf{R}_E .

As is well known, there is a standard self-reducibility argument: by randomization, it is possible to transform an arbitrary tuple (G_1, G_2, U_1, U_2) such that $G_1 \neq G_2$ into a random equivalent one, i.e. the output is in \mathbf{D}_E (resp. \mathbf{R}_E), if and only if the input is. Thus, if $\text{AdvDDH}(t)$ is significant, one can use a distinguisher to decide, with probability close to one, whether a tuple is in \mathbf{D}_E . This involves performing repeated tests with the distinguisher and deciding whether the number of one outputs has a bias towards \mathbf{D}_E or \mathbf{R}_E . Based on the law of large numbers, a decision with small constant error probability requires running $O(\text{AdvDDH}^{-2})$ tests. One can decrease the error probability drastically by repeating the above computations an odd number of times and deciding based on the median of the averages. In [26], the authors claim that one can reach error probability 2^{-n} by repeating the test $O(p(n)) \cdot \text{AdvDDH}^{-1}$, where p is a polynomial, but the proof is missing. In any case, the loss in the reduction is huge. Thus, despite its elegance, the self-reducibility argument is a bit misleading in terms of exact security.

Related to the above is the elliptic curve computational Diffie-Hellman assumption (ECCDH) and the elliptic curve discrete logarithm assumption. The former states that it is hard to compute $xy \cdot G$ from G , $x \cdot G$ and $y \cdot G$, while the latter states that it is hard to compute x from G and $x \cdot G$. It is obvious that DDH is a stronger assumption than CDH, which in turn, is stronger than the discrete logarithm assumption. However, no other relation is known and the only way to solve the hard problems underlying DDH or CDH is to compute discrete logarithms.

It should be mentioned that the DDH cannot hold in groups with a small subgroup. This is why cryptographic schemes usually work with a subgroup of an elliptic curve of large prime order. Even with this proviso, there are subtle protocol attacks using

invalid keys, i.e. keys that do not belong to the prescribed large subgroup (see [25]). In the present context, such attacks are irrelevant since the curve itself is of prime order.

3.1.2 Security of the scheme

It appears that the security of the scheme is closely connected to the decisional Diffie-Hellman assumption. Consider the second setting where a fresh key is generated at each execution of the protocol. With the notations of section 2.1, this is an appropriate combination of the following data

1. $V_A = DH_{E_A}(Y_A, R_A)$, which is $x_A r_B \cdot G_A$
2. and $V_B = DH_{E_B}(Y_B, R_B)$, which is $x_B r_A \cdot G_B$.

We would like to see that the x coordinate of both V_A and V_B look like random strings to a passive adversary. However, this cannot hold in a simple-minded approach: if the equation of the elliptic curve E_A used by A is

$$y^2 = x^3 + ax + b$$

then, an integer x , $0 \leq x < p_A$ is the first coordinate of a point on the curve if and only if $x^3 + ax + b$ is a quadratic residue modulo p_A . We let \mathcal{X}_A be the set of such x and we use a similar notation \mathcal{X}_B for the curve E_B of the other party.

Theorem 1 *Based on the elliptic curve decisional Diffie-Hellman hypothesis (ECDDH), it is hard to distinguish the distribution*

$$(G_A, Y_A, R_A, x_{V_A})$$

generated by the cryptosystem, from the analogous distribution with x_{V_A} replaced by a random elements of \mathcal{X}_A . More accurately, if there is an adversary \mathcal{A} that distinguishes the above distributions within time bound t , with advantage ε , then there exists a machine \mathcal{B} that solves the decisional Diffie-Hellman problem with advantage ε within time bound $t + \tau$, where τ accounts for a few extra elliptic curve operations and is bounded by $\mathcal{O}(k^3)$.

In the above, the advantage in distinguishing two distributions is the absolute value of the difference of the probabilities that the algorithm outputs 1, with inputs taken from each. Also, x_{V_A} denotes the first coordinate of V_A .

Proof. Let \mathcal{A} be an adversary that distinguishes the two distributions defined in the theorem. We show how to attack the ECDDH by distinguishing the distributions \mathbf{D}_A and \mathbf{R}_A , where

$$\mathbf{R}_A = \{(G_A, Y_A, R_A, V_A)\}$$

with four elements taken at random in E_A and

$$\mathbf{D}_A = \{(G_A, Y_A, R_A, V_A)\}$$

with $\log_{G_A}(R_A) = \log_{Y_A}(V_A)$. We run the key generation algorithm and generate a prime number p_A together with an elliptic curve over \mathbb{Z}_{p_A} with $p_A - 2$ elements (also a prime). We next show how to use \mathcal{A} to break the ECDDH: we take the base point of the cryptosystem to be the first coordinate G_A of the input to \mathcal{A} and we complete the public key by the second coordinate Y_A . This implicitly defines a secret key. Next we assume that a key-agreement takes place with another entity B , sending the third coordinate R_A . Finally, we submit (G_A, Y_A, R_A, x_{V_A}) to \mathcal{A} . If the original input is from \mathbf{D}_A , the last coordinate is exactly as produced by the cryptosystem. On the other hand, if it is from \mathbf{R}_A , the last coordinate is a random element of \mathcal{X}_A , unless $V_A = \mathcal{O}$, which happens with probability $\frac{1}{p_A-2}$. Thus, we have obtained a distinguisher between \mathbf{D}_A and \mathbf{R}_A , with almost exactly the same advantage as \mathcal{A} . Finally, the advantage of any algorithm \mathcal{A} that runs in time t is bounded by $\text{AdvDDH}(O(t))$, where $O(t) = t + \tau$ accounts for few extra elliptic curve operations needed to compute the data to be handled to \mathcal{A} .

Remark. It follows from the above theorem that, when the key is computed as the exclusive or of the respective x coordinates of $DH_{E_A}(Y_A, R_A)$ and $DH_{E_B}(Y_B, R_B)$, it is indistinguishable from an element built as the exclusive or of random elements respectively taken from $\mathcal{X}_A, \mathcal{X}_B$. It would be desirable to ensure that one gets a random bitstring with $|p|$ bits, which would mean that the information x_{V_A} is semantically secure in the sense of the seminal paper [18]. However, we do not see any argument that would give such guarantee. Thus, there is no proof that, considered as a bit string, the session key is semantically secure. We will return to this feature further on in the present report.

3.2 Size of the parameters

As was just observed the security of the basic scheme appears closely related to the ECDDH for the class of elliptic curves generated by the cryptosystem, even if there is a minute security loss in terms of exact security. The only method known to attack the decisional Diffie-Hellman problem on elliptic curves is to solve the underlying discrete logarithm problem (ECDLP). In order to estimate whether the specific restrictions on the curve and the suggested parameters offer a wide security margin, it is useful to review the performances of the various algorithms known for the ECDLP. We will distinguish between exponential algorithms, whose running time depend on the size of the group and subexponential algorithms, which apply to specific classes of weak curves.

Before entering into a more precise discussion, let us mention that the idea of generating curves with a prescribed number of points is not new. For example, it appears in [21]. This remark is not in terms of intellectual property but rather in terms of the novelty of the idea.

3.2.1 Exponential algorithms

The best algorithm known to date for solving the DLP in any given group G is the Pollard ρ -method from [27] which takes computing time equivalent to about $\sqrt{\pi n/2}$ group operations. In 1993, van Oorschot and Wiener in [34], showed how the Pollard ρ -method can be parallelized so that, if t processors are used, then the expected number of steps by each processor before a discrete logarithm is obtained is $\simeq \frac{\sqrt{\pi n/2}}{t}$. In order to compute the discrete logarithm of Y in base G , each processor computes a kind of random walk within elements of the form $a \cdot G + b \cdot Y$, selecting X_{i+1} through one of the three following rules

1. set $X_{i+1} = G + X_i$
2. set $X_{i+1} = 2 \cdot X_i$
3. set $X_{i+1} = Y + X_i$

Decisions on which rule to apply are made through a random-looking but deterministic computation, using e.g. hash values. “Distinguished” points X_i are stored together with their representation $X_i = a_i \cdot G + b_i \cdot Y$ in a list that is common to all processors. When a collision occurs in the list, the requested discrete logarithm becomes known.

In recent work (see [17, 35]), it was shown how to improve the above by a multiplicative factor $\sqrt{2}$. This takes advantage of the fact that one can simultaneously handle a point X and its opposite $-X$. Slightly better improvements can be obtained for specific curves with automorphisms.

The progress of such algorithms is well documented. In April 2000, the solution to the ECC2K-108 challenge from Certicom [8] led to the computation of a discrete logarithm in a group with 2^{109} elements (see [14]). This is one of the largest effort ever devoted to a public-key cryptography challenge. The amount of work required to solve the ECC2K-108 challenge was about 50 times that required to solve the 512-bit RSA cryptosystem (see [7]) and was thus close to 400000 mips-years.

It is expected that such figures will grow slowly, unless unexpected discoveries appear in the area. From the predictions in [23], one can infer that the proposed range of parameters ($|p| = 160, 192, 224$) will presumably allow for a choice that guarantees security for the foreseeable future, at least for the next 30 years.

3.3 Security against subexponential attacks

As is well known, there are two classes of elliptic curves for which non trivial attacks have been found. They are

1. the supersingular curves
2. the anomalous curves

Supersingular curves over a field \mathbb{F}_q , with q a power of p , are defined by the condition that the trace of the Frobenius map is zero modulo p . For such curves, Menezes, Okamoto and Vanstone (MOV) have shown how to reduce the discrete logarithm problem to the DLP in an extension field \mathbb{F}_{q^j} of \mathbb{F}_q , with small j . Note that, for elliptic curves over a prime field \mathbb{Z}_p , those curves have exactly $p+1$ elements and are specifically excluded by the key generation algorithm of HDEF-ECDH.

Anomalous curves are those which contain a p -torsion point other than \mathcal{O} , or, equivalently, those whose Frobenius map has trace congruent to one modulo p . For such curves, work of Semaev ([31]), Rück ([28]), Smart ([30]) and Satoh-Araki ([29]) has shown how to solve the p -part of the DLP in polynomial time. Note that, for elliptic curves over a prime field \mathbb{Z}_p , those curves have exactly p elements and are specifically excluded by the key generation algorithm of HDEF-ECDH.

The MOV reduction constructs an embedding from the curve into the multiplicative group of a suitable extension field of \mathbb{F}_q and can be applied in a more general setting than originally envisioned by the authors. However, if the base point is an element of order ℓ , ℓ is necessarily a divisor of $q^j - 1$. Recently, Balasubramanian and Koblitz have shown in [1] that this condition was sufficient to carry the MOV reduction. The key generation algorithm specifically addresses this question. Here q is a prime p and $\ell = p - 2$. Thus one gets that $p = 2 \pmod{\ell}$ and $p^j = 1 \pmod{\ell}$. From this, it follows that $2^j = 1 \pmod{\ell}$, hence the lower bound $j \geq \log \ell$. This makes j large enough to turn down subexponential algorithms in the extension field.

Another reduction similar to the MOV reduction has appeared in the literature. It is due to Frey and Rück [13] (see also [12]) and can be stated in the more general context of Jacobians on which the Tate pairing exists. Let m be an integer relatively prime to q , and let $\mu_m(\mathbb{F}_q)$ be the group of roots of unity in \mathbb{F}_q whose order divides m . Assume that the Jacobian $J(\mathbb{F}_q)$ contains a point of order m . Then there is a surjective pairing

$$\varphi_m : J_m(\mathbb{F}_q) \times J(\mathbb{F}_q)/mJ(\mathbb{F}_q) \rightarrow \mu_m(\mathbb{F}_q)$$

which is computable in $\mathcal{O}(\log q)$, where $J_m(\mathbb{F}_q)$ is the group of m -torsion points. This pairing, the so-called Tate pairing, can be used to relate the discrete logarithm in the group $J_m(\mathbb{F}_q)$ to the discrete logarithm in some extension $\mathbb{F}_{q^j}^*$. In the case of elliptic curves considered in the current context, the above is applicable only if q is p and m is the order ℓ of the group generated by the base point. Thus, ℓ has to be a divisor

of $p^j - 1$. As a consequence, the curves produced by the key generation algorithm are protected against the FR reduction, exactly due to the same argument used for MOV reduction.

3.3.1 Conclusion

Based on current estimates, it appears that the proposed parameters for HDEF-ECDH should remain secure for at least thirty years. However, in order to guarantee that the MOV and FR reductions do not apply, the key generation algorithm restricts the class to elliptic curves used by the scheme to a very specific subset of curves with complex multiplication. This is somehow contrary to the current trend, which would recommend having the curve generated at random and ensuring that there is a point of large prime order by counting the number of elements of the curve by means of the SEA algorithm [24]. The advantages of the strategy adopted in the current scheme are not argued in a convincing manner.

4 Security Analysis

The self-evaluation report [16] does not include a serious security analysis, and does even not consider the formal security notions that a key agreement scheme should satisfy. Thus, we have found necessary to undertake our own security analysis. We first review the security notions which have been defined in the literature. Then we consider the proposed schemes in view of these notions.

4.1 Formal framework

4.1.1 Key agreement

A key agreement scheme (without TTP) involves two participants, a *client* and a *server*, who want to share a secret session key in order to thereafter possess a secure and virtually private channel. They communicate on a public channel and eventually compute a value that they both know but which nobody else knows. Many security models have been defined to cover this kind of schemes. Of these, the following two models have received more attention:

- The first model was proposed by Bellare and Rogaway [5, 6], and refined in [3]. Here, the adversary can interact with all participants and aims at learning some information about one session key. Therefore, the proper approach is to ensure indistinguishability of the session key (from a random key) for the adversary. In other words, any session key should be semantically secure [18].

- The second model was proposed by Bellare, Canetti and Krawczyk [2], and is based on the multi-party simulatability technique. This means that one first defines an idealized version of a key agreement scheme. Then, in order to prove that the real-world scheme is secure, one shows that any adversary in the real world has to behave like an adversary in the ideal game.

Shoup [32] has shown that the two models (with adequate refinements) are equivalent in preventing active adversaries to break *forward-secrecy*. This property is by now a basic requirement for any key agreement scheme. *Forward-secrecy* means that an adversary, who sees all the public communication (and possibly has access to all session keys *but one*) cannot obtain any information about *that last* session key, even if he later learns the long-term secret of any party.

4.1.2 Mutual authentication

When parties have established a common secret session key, most of the key agreement protocols, such as the Diffie–Hellman [10] key agreement scheme using public keys, *implicitly* assume that each party is actually partnered (by sharing the session key) with the party he wanted. However, it can be the case that no partnership has been established. Indeed, if an adversary uses the public key of Alice and Bob runs the key exchange process, then, upon completion, he thinks that the actual session key is shared with Alice. However, there is no actual partner since the adversary cannot extract the session key from the communication.

Accordingly, one usually wants to furthermore verify the actual partnership. Such property of a key exchange scheme is called *mutual authentication*. However, as presented in [3], an *implicitly authenticated* key agreement scheme can be easily transformed, in the random oracle model [4], into a scheme that provides mutual authentication, by simply adding one more flow (see figure 1).

4.2 Security model

4.2.1 Key agreement

At the end of each execution of the protocol (see figure 1), when a party \mathcal{U} has accepted, it gets a session key, denoted by $\text{sk}_{\mathcal{U}}$, and a session ID, denoted by $\text{sid}_{\mathcal{U}}$ which is part of the flow of data. The session ID's are made public, while session keys clearly remain secret. Indeed, the session keys are the common secret shared by the two parties at the end of the protocol. The session ID's have a technical significance: they are used to define partnership. The partner of a party is an entity which has a similar session ID. Since the session ID's are public, the partnership is also public. With such a definition of partnership, one can remark that a party may have several partners, although it is quite unlikely, in general.

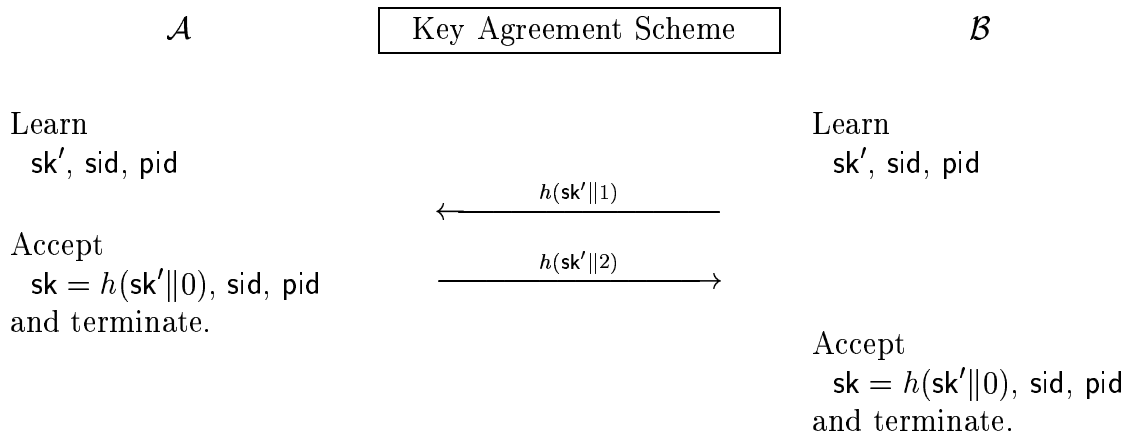


Figure 1: Key Agreement + Mutual Authentication

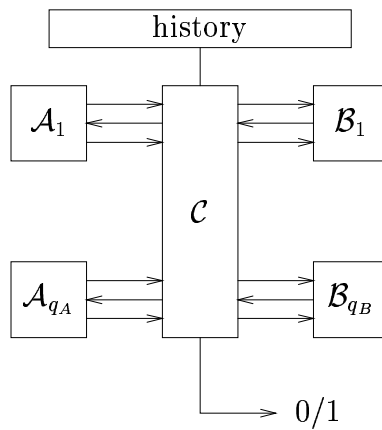


Figure 2: Security Model

In the model defined by Bellare and Rogaway [5, 6], with additional refinements in order to handle forward-secrecy (*cf.* Shoup [32]), any instance of each party, \mathcal{A} or \mathcal{B} , is seen as an oracle (see figure 2). Furthermore, it is assumed that the entire communication network is managed by the adversary \mathcal{C} , who may schedule interactions arbitrarily, and who may inject and drop messages arbitrarily as well. Thus, the adversary can interact, as a man-in-the-middle, with all parties, or more formally with several instances of them (\mathcal{A}_i for the client and \mathcal{B}_j for the server) as many times as he wants in a concurrent way. He can ask the following queries

- **Send** ($\mathcal{U}, i, \text{string}$) – which means that the adversary sends the message *string* to the oracle \mathcal{U}_i (either a server or a client). The oracle makes the requested computations according to the protocol and returns the answer.
- **Reveal** (\mathcal{U}, i) – provided oracle \mathcal{U}_i has accepted (the tag `acc` has been set to `True`), it returns the session key $\text{sk}_{\mathcal{U}}^i$. This models the misuse of a session key by the parties once this session key has been established.
- **Test** (\mathcal{U}, i) – granted that oracle \mathcal{U}_i has accepted, one tosses a coin b . If $b = 1$ then the session key $\text{sk}_{\mathcal{U}}^i$ is returned, else a random string is returned. The aim of the attack is to guess this bit b . Therefore, the usual obvious restrictions on this query apply:
 - the query is asked once;
 - no **Reveal**-query is asked to \mathcal{U}_i ;
 - no **Reveal**-query is asked to \mathcal{U}_j , where \mathcal{U}_j is partnered with \mathcal{U}_i .
- **Execute** ($\mathcal{A}, i, \mathcal{B}, j$) – in order to obtain the transcript corresponding to the communication between two parties (and build a history), the adversary may ask the parties to run the protocol. Then the transcript is returned to the adversary. Such a transcript can also be obtained using **Send**-queries.
- **Corrupt** (\mathcal{U}) – in order to deal with forward-secrecy, one allows the adversary to corrupt the parties. By this we mean that it obtains the secret key (the long-term secret key $x_{\mathcal{U}}$) of the corrupted party \mathcal{U} . This induces a further restriction on the **Test**-query, which can only be asked to an instance of a party (or one of his partners) who has accepted before it gets corrupted.

4.2.2 Mutual authentication

The above game, with the **Test**-query, just deals with the semantic security of the session key, the basic security notion of a key agreement scheme, but not with (mutual) authentication. We say that the protocol provides mutual authentication if no instance

accepts unless it has exactly one partner. Otherwise, it would mean that the adversary has impersonated a party. For example, if an instance of the server accepts without a partner, it means that the adversary has impersonated the client, and therefore broken client-to-server authentication. Similarly, if an instance of the client accepts without a partner, it means that the adversary has impersonated the server, and therefore broken server-to-client authentication. A key agreement scheme guarantees mutual authentication if, for any adversary, the probabilities of breaking the client-to-server authentication or the server-to-client authentication are both negligible.

4.2.3 Digital signature schemes

In the following, we will use digital signatures in order to strengthen the security of the proposals. Therefore, we briefly review the definition and the security notions.

Digital signature schemes are the electronic version of handwritten signatures for digital documents: a user's signature on a message M is a string which depends on M , on the secret key of the user and –possibly– on randomly chosen data, in such a way that anyone can check the validity of the signature by using the public key only.

Definitions. A signature scheme is defined by three algorithms $(\mathcal{K}, \Sigma, \mathcal{V})$:

- The *key generation algorithm* \mathcal{K} . On input 1^k , where k is the security parameter, the algorithm \mathcal{K} produces a pair $(\mathbf{pk}, \mathbf{sk})$ of matching public and secret keys. Algorithm \mathcal{K} is probabilistic.
- The *signing algorithm* Σ . Given a message M and a pair of matching public and secret keys $(\mathbf{pk}, \mathbf{sk})$, Σ produces a signature σ . The signing algorithm might be probabilistic.
- The *verification algorithm* \mathcal{V} . Given a signature σ , a message M and a public key \mathbf{pk} , \mathcal{V} tests whether σ is a valid signature of M with respect to \mathbf{pk} . In general, the verification algorithm need not be probabilistic.

Security. Various security notions have been formalized by [19, 20], based on the *goal* of the adversary, and the *means* available to the adversary to achieve this goal.

- Disclosing the secret key of the signer. It is the most serious attack. This attack is termed *total break*.
- Constructing an efficient algorithm which is able to sign messages with significant probability of success. This is called *universal forgery*.
- Providing a new message-signature pair. This is called *existential forgery*.

In many cases the is not dangerous, because the output message is likely to be meaningless. Nevertheless, a signature scheme which is not existentially unforgeable (and thus that admits existential forgeries) does not guarantee by itself the identity of the signer. For example, it cannot be used to certify randomly looking elements, such as keys, which is the context considered in this report.

Two different kinds of attacks have been considered depending on the availability of signed messages for the adversary. In a first scenario, the attacker only knows the public key of the signer. In another, the attacker has access to a list of valid message-signature pairs. The strongest model is the *adaptively chosen-message attack*, where the attacker can ask the signer to sign any message, except the message for which forgery is finally achieved.

A signature scheme is *secure* if an existential forgery is computationally impossible, even under an adaptively chosen-message attack. We denote by $\text{Succ}^{\mathcal{K}, \Sigma, \mathcal{V}}(t, q)$ the maximal success probability of any adversary in performing an existential forgery, after at most q queries to a signing oracle, within time bound t .

4.3 Analysis of the proposal: Common domain parameters

In this section, we consider the security of the key agreement scheme between two parties who share the same domain parameters, $E = E_A = E_B$ and $G = G_A = G_B$.

4.3.1 The basic protocol

A passive adversary does not make use of the $\text{Send}(\mathcal{U}, i, \text{string})$ or $\text{Corrupt}(\mathcal{U})$ facilities. It may or may not ask Reveal -queries. If it does not, then it is easily seen that it would have to break the ECDDH problem in order to get an advantage in distinguishing the real session key from a “random” one (derived from a random point on the curve), as proven in theorem 1. However, this does not guarantee that the session key k is indistinguishable from a random string of the same length (but only from a string obtained from a random curve element). Building on this remark, we see that the scheme is not secure, if the session key is computed as the first coordinate of an elliptic curve point obtained by the Diffie-Hellman function, as suggested. Let k be the answer to a Test -query. We can distinguish between an actual key and a random string by testing whether $k \in \mathcal{X}$, wher \mathcal{X} is the set of x coordinates of elements on the curve. The former passes the test in all cases, while the latter is successful with probability $\simeq 1/2$. This contradicts formal security. Thus, the scheme under review cannot be secure, unless it is enhanced by some more elaborate key derivation function that computes a random bit-string from a random element of \mathcal{X} . Such function can be provided by a suitable hash function considered as a random oracle. It is also possible to avoid the random oracle model and still obtain a random bit-string by applying a randomly keyed universal hash function, following the method described in [26] and

also used in [33]. Recall that, if H_k is a universal hash function, keyed by k , with ℓ -bit outputs, then, the leftover hash lemma of [22] implies that hashing a set of 2^λ bit-strings produces a distribution $(k, H_k(x))$ whose distance to the uniform distribution is $\leq \frac{1}{2^{(\lambda-\ell)/2}}$. Here, λ is exactly the security parameter k . Thus in order to get a bound at most $1/2^{128}$ and to obtain a 128 bit encryption key, one would need $k \geq 384$. This is beyond the parameters considered in [15]. Anyway, this is not the path followed by the submission.

In the case of a passive adversaries with **Reveal**-queries, the scheme is not secure, even with the enhancements just considered, since the session key is always the same. As soon as the adversary has asked a **Reveal**-query, he can break the semantic security of any other session key (since they are the same).

Finally, active adversaries cannot really be considered, since there are no interactions in this scheme, where only the public keys are used. Similarly, the scheme does not provide forward-secrecy, since the session key is completely and easily determined from the public keys, and the secret key of any participant.

4.3.2 One-time keys

We now turn to the situation where the two parties refresh their keys at each execution of the key agreement scheme. Notice that this scenario is not explicitly envisioned in documents [15, 16]. Still, it is needed in order to withstand active attacks. In this context, it is necessary to authenticate the flow of data by means of a digital signature. Thus, the long-term key on which the protocol relies is actually a signing key. The key-agreement scheme between two entities \mathcal{A} and \mathcal{B} , with respective identities ID_A , ID_B , can be played as follows (as proposed in [32].)

1. \mathcal{A} generates a one-time key pair (x_A, Y_A) and sends Y_A , ID_B to \mathcal{B} together with a signature $sig_A(Y_A, ID_B)$ of these data and a certificate for his public signing key
2. \mathcal{B} generates a one-time key pair (x_B, Y_B) and sends Y_B , to \mathcal{A} together with a signature $sig_B(Y_A, Y_B, ID_A)$ and a certificate for his public signing key

The key material can be derived by \mathcal{A} and \mathcal{B} by computing

$$K = DH(Y_A, Y_B) = x_A \cdot Y_B = x_B \cdot Y_A$$

Once this is done, they can obtain a shared key k .

4.3.3 Security analysis for one-time keys

We now analyze the security of the variant of the protocol, which uses one-time keys and digital signatures. Again, the scheme cannot be secure if the session key is computed as

the first coordinate of an elliptic curve point obtained by the Diffie-Hellman function, as suggested. Thus, we will consider an enhanced scheme, endowed with a more elaborate key derivation function, which can be seen as a random oracle.

Passive adversaries without Reveal-queries. It is easily seen that a passive adversary which does not ask any Reveal-query would have to break the ECDDH problem in order to get any advantage in distinguishing the real session key from a “random” one (derived from a random point on the curve), with a proof similar to the proof of theorem 1. This proves the security of the enhanced scheme in this setting.

Passive adversaries with Reveal-queries. In this case, it can be shown that the enhanced scheme is secure relative to the ECDDH problem. To see this, it is enough we prove the following result.

Theorem 2 *Let \mathcal{C} be an adversary breaking the semantic security of the enhanced key agreement scheme, with advantage ε and within time bound t , having observed q_p transcripts and asked q_r Reveal-queries. Then the ECDDH problem can be solved with advantage greater than ε/q_p and within time bound $t + 3q_p \cdot \tau$, where τ accounts for the cost of an elliptic curve operation and is bounded by $\mathcal{O}(k^3)$.*

Proof. Let \mathcal{C} be a passive adversary that guesses the bit b involved in the Test-query. We construct a distinguisher between the distributions \mathbf{D} and \mathbf{R} , where

$$\mathbf{R} = \{(G, Y, U, V)\}$$

with all four elements taken at random in E and

$$\mathbf{D} = \{(G, Y, U, V)\}$$

with $\log_G(U) = \log_Y(V)$.

We run the key generation algorithm and generate a prime number p together with an elliptic curve over \mathbb{Z}_p with $p - 2$ elements (also prime). We also run the key generation algorithm of the signature scheme. We next show how to use \mathcal{C} to break the ECDDH on E : we take the base point of the cryptosystem to be the first coordinate of the input to \mathcal{C} .

The view of the adversary \mathcal{C} can be simulated as follows:

- choose a random x_A and send $Y_A = x_A \cdot G$, with signature and certificate
- one chooses a random x_B and send $Y_B = x_B \cdot G$, with signature and certificate
- derive the shared session key k

In order to use the adversary to distinguish the above distributions, we randomly choose one index $i \in \{1, \dots, q_p\}$ and modify the i -th execution. Instead of the above simulation, one uses the second component Y of the input to \mathcal{C} in the first message and the third component U for the answer. The shared session key k is derived from the point V .

The **Reveal**-queries can easily be simulated by means of the computed shared key. Similarly, the **Test**-query can be simulated, using a random coin b . Thus, provided that the **Test**-query is asked at the i -th execution, we see that the advantage of \mathcal{C} is exactly the same as in the real game, when the input to \mathcal{C} comes from \mathbf{D} . On the other hand, when the input to \mathcal{C} comes from \mathbf{R} , the advantage of \mathcal{C} is exactly 0. Since the choice of i is independent from the view of the adversary, the advantage of our distinguisher is ε/q_p , and its running time is $t' = t + 3q_p \cdot \tau$, where τ accounts for the cost of the elliptic curve operation needed to compute the data to be handled to \mathcal{C} .

Using an improved reduction, similar to [9], one can state

Theorem 3 *Let \mathcal{C} be an adversary that breaking the semantic security of the enhanced key agreement scheme, with advantage ε and within time bound t , having observed q_p transcripts and asked q_r **Reveal**-queries, then the ECDDH problem can be solved with advantage greater than $e^{-1} \cdot \varepsilon/q_r$ and within time bound $t + 3q_p \cdot \tau$, where τ accounts for the cost of the elliptic curve operation and is bounded by $\mathcal{O}(k^3)$.*

Proof. The proof is similar to the above, but the simulation of the view of the adversary \mathcal{C} is slightly different. With probability $1 - \pi$, one performs the standard simulation, having knowledge of the discrete logarithms in use, and with probability π the simulation uses inputs derived from Y, U, V using random self-reducibility (their discrete logarithms are unknown). A similar analysis can be performed, provided that all the **Reveal**-queries have been asked to correctly simulated transcript, and that the **Test**-query has been asked on a transcript involving a key derived from (Y, U, V) . This occurs with probability $(1 - \pi)^{q_r} \pi$. If we define $\pi \approx 1/q_r$, then above probability is approximately equal to e^{-1}/q_r , which concludes the proof.

Active adversaries. We now cover the case of active adversaries, allowing the use of **Send** ($\mathcal{U}, i, \text{string}$) but not of the **Corrupt** (\mathcal{U}) facility, at this point.

Due to the signature, the scheme can be shown secure relative to the ECDDH problem. More accurately, we prove the following:

Theorem 4 *Let \mathcal{C} be an adversary breaking the semantic security of the key agreement scheme, with advantage ε and within time t , having asked q_r **Reveal**-queries, and interacted q times with \mathcal{A} and \mathcal{B} . Then*

$$\begin{aligned} \varepsilon &\leq 2 \times \text{Succ}^{\mathcal{K}, \Sigma, \mathcal{V}}(t', q) + q^2 \times \text{AdvDDH}(t') \\ \text{with } t' &\leq t + 2q \cdot \tau, \end{aligned}$$

where τ accounts for the cost of an elliptic curve operation and is bounded by $\mathcal{O}(k^3)$.

Remark. In above theorem, we assume that passive observations (the **Execute**-queries) are built from interactions with \mathcal{A} and \mathcal{B} (using **Send**-queries).

Proof. Let \mathcal{C} be an active adversary that guesses the bit b involved in the **Test**-query after q interactions with \mathcal{A} and \mathcal{B} (q_a and q_b respectively), and q_r **Reveal**-queries.

In order to prove the above security result, we will envision several games:

- game \mathcal{G}_1 , where the signatures are generated by the actual signature algorithms (by means of the secret keys)
- game \mathcal{G}_2 , where all messages to \mathcal{B} which involve a fresh signature (i.e. a signature not produced by our simulation of \mathcal{A}) are rejected
- game \mathcal{G}_3 , where all messages which involve a fresh signature (i.e. a signature not produced by our simulators) are rejected

The probability that the adversary correctly guesses b in \mathcal{G}_1 is exactly $1/2 + \varepsilon$. Indeed, game \mathcal{G}_1 provides the adversary with the real-life setting.

In the following, we bound the difference of probabilities that \mathcal{G}_1 and \mathcal{G}_3 successfully guess the query bit b , and we relate the advantage in the game \mathcal{G}_3 with the ability to break the ECDDH. This uses the simple yet useful lemma from [33]

Lemma 1 *Let E, F , and E', F' be events of two probability spaces such that both*

$$\Pr[E|\neg F] = \Pr[E'|\neg F'] \text{ and } \Pr[F] = \Pr[F'] \leq \varepsilon.$$

Then,

$$|\Pr[E] - \Pr[E']| \leq \varepsilon$$

Proof: We write

$$\begin{aligned} \Pr[E] &= \Pr[E|\neg F] \Pr[\neg F] + \Pr[E|F] \Pr[F] \\ \Pr[E'] &= \Pr[E'|\neg F'] \Pr[\neg F'] + \Pr[E|F'] \Pr[F'] \end{aligned}$$

Hence

$$\Pr[E] - \Pr[E'] = \Pr[E|\neg F](\Pr[\neg F] - \Pr[\neg F']) + (\Pr[E|F] \Pr[F] - \Pr[E|F'] \Pr[F'])$$

The right hand side becomes $\Pr[E|F] \Pr[F] - \Pr[E|F'] \Pr[F']$, which is bounded by ε .

Going from game \mathcal{G}_1 to \mathcal{G}_2 produces a difference if and only if a message produced by the adversary involves a fresh signature, accepted under the public key \mathbf{pk}_a . In order to estimate the probability that such event happens, we define a simulation. This simulation runs the key generation algorithm anew and generates a prime number p

together with an elliptic curve over \mathbb{Z}_p with $p - 2$ elements (also prime). It also runs the key generation algorithm of the signature scheme to get a public key \mathbf{pk}_a for \mathcal{A} , and a pair of keys $(\mathbf{sk}_b, \mathbf{pk}_b)$ for \mathcal{B} , as well as the certificates for the public keys. At this point, one can simulate the view of the adversary \mathcal{C} , with the help of a signing oracle for \mathcal{A} :

- to simulate \mathcal{A} , choose a random x_A and compute $Y_A = x_A \cdot G$, then ask the signing oracle the signature $sig_A(Y_A, ID_B)$. Return Y_A , the signature and a certificate for \mathbf{pk}_a
- to simulate \mathcal{B} , perform similarly, using the secret signing key in place of the oracle. More explicitly, upon receiving data produced by the simulation of \mathcal{A} , abort if the received signature is fresh (i.e. has not been created by the above \mathcal{A} -simulator.) Otherwise, choose a random x_B and compute $Y_B = x_B \cdot G$; next, produce the signature $sig_B(Y_A, Y_B, ID_B)$. Return Y_B , the signature and a certificate for \mathbf{pk}_b .
- \mathcal{A} accepts if and only if the received signature is correct. Then each of \mathcal{A} , \mathcal{B} derives the shared session key \mathbf{k}

This also simulates *Execute*-queries in \mathcal{G}_2 . The case of *Reveal*-queries is easily handled, using the computed shared key. Similarly, *Test*-queries can be simulated, by means of a random coin b . Now, game \mathcal{G}_2 differs from game \mathcal{G}_1 if a fresh signature is valid, but this is exactly the probability that the simulation provides an existential forgery, with less than q_a queries to the signing oracle and within time bound $t' = t + 2q \cdot \tau$. Using the lemma, this bounds the difference between the success probabilities by $\text{Succ}^{\mathcal{K}, \Sigma, \mathcal{V}}(t', q_a)$. A similar analysis bounds the difference between games \mathcal{G}_2 and \mathcal{G}_3 by $\text{Succ}^{\mathcal{K}, \Sigma, \mathcal{V}}(t', q_b)$.

We are thus led to study the advantage that the adversary can get in game \mathcal{G}_3 . We relate this advantage to a distinguisher between the two distributions \mathbf{D} and \mathbf{R} , that we now describe. We run the key generation algorithm and generate a prime number p together with an elliptic curve over \mathbb{Z}_p with $p - 2$ elements (also prime). We also run the key generation algorithm of the signature scheme to get the signing and verification keys, as well as the certificates. We next show how to use \mathcal{C} , on input (G, Y, U, V) , to break the ECDDH. We take the base point of the cryptosystem to be the first coordinate G of the input to \mathcal{C} . We then simulate the view of the adversary \mathcal{C} , as follows:

- to simulate \mathcal{A} , choose a random x_A and send $Y_A = x_A \cdot G$, with signature and certificate
- to simulate \mathcal{B} , choose a random x_B and send $Y_B = x_B \cdot G$, with signature and certificate
- then each of \mathcal{A} and \mathcal{B} derives the shared session key \mathbf{k} .

In above simulations of \mathcal{A} and \mathcal{B} , they only accept signatures generated by the simulator. In order to use the adversary to distinguish the above distributions, we randomly select two indices $i, j \in \{1, \dots, q\}$. The i -th simulation of \mathcal{A} is modified, using the second component Y of the input to \mathcal{C} to define Y_A . Similarly, the j -th simulation of \mathcal{B} , provided it answers a message involving Y , uses the third component U of the input to \mathcal{C} to define Y_B . The shared session key is obtained from V .

The **Reveal**-queries are simulated by answering the computed shared key. Similarly, the **Test**-query are simulated, by means of a random coin b . Thus, provided that the **Test**-query is asked at a point where Y and U are involved, we see that the advantage of \mathcal{C} is exactly the same as in the real game, when the input to \mathcal{C} comes from \mathbf{D} . On the other hand, when the input to \mathcal{C} comes from \mathbf{R} , the advantage is exactly 0. Since the choice of i and j is independent from the view of the adversary, the advantage of our distinguisher is ε/q^2 , and its running time is $t' = t + 2q \cdot \tau$. This finishes the proof.

Forward-secrecy. The above result can be extended to the case where \mathcal{C} is allowed to use the **Corrupt** (\mathcal{U}) facility. However, as pointed out in [32], its use should not be too liberal: one sees that the above proof collapses if the adversary corrupts \mathcal{B} right after a **Test**-query. The proper restriction can be expressed in terms of the session ID, as follows:

- a session ID is *established* only after two parties \mathcal{A} and \mathcal{B} have interacted to share this ID
- if a **Test**-query is asked to a party \mathcal{U} , no **Reveal**-query can be issued with the session ID involved in the **Test**-query
- if a **Test**-query is asked to a party \mathcal{A} which later gets corrupted, the corruption cannot take place before the session ID involved in the **Test**-query has been established.

As shown by Shoup [32], in order to get the most general version of forward security, it is necessary to add at least one key confirmation flow.

4.3.4 Mutual Authentication

As already studied (see figure 1), by simply adding the key confirmation flows (only one, or both), one gets the explicit unilateral or mutual authentication. Without such a key confirmation, the users may accept without having any actual partner.

4.4 Analysis of the Proposal: Different domain parameters

We now turn to the situation where the two parties own different domain parameters, (E_A, G_A) and (E_B, G_B) .

4.4.1 The basic protocol

A generates a random $r_A \in \{0, \dots, p_B - 2\}$ and sends $R_B = r_A \cdot G_B$ to B ; similarly, B generates a random $r_B \in \{0, \dots, p_A - 2\}$ and sends $R_A = r_B \cdot G_A$ to A ; once this has been done, each of A and B can derive the two values $K_A = DH_{E_A}(Y_A, R_A)$ and $K_B = DH_{E_B}(Y_B, R_B)$. From these values, they can obtain a shared key; for example, they can take the exclusive or of the respective x coordinates of K_A and K_B .

Passive adversaries without Reveal-queries. Again, it is easily seen that a passive adversary which does not ask any Reveal-query would have to break the ECDDH problem in order to get any advantage in distinguishing the real session key from a “random” one (derived from two random points on the curves), with a proof similar to the proof of theorem 1. Again, this does not guarantee that the session key k is indistinguishable from a random string of the same length (but only from a string obtained from a random curve element). Thus, a formal proof seems to require the enhanced version of the scheme. However, it is unclear whether there is an attack in this setting against This would mean distinguishing a random string from an element built as the exclusive or of random elements respectively taken from $\mathcal{X}_A, \mathcal{X}_B$.

Passive adversaries with Reveal-queries. In this case, it can be shown that the enhanced scheme is secure relative to the ECDDH problem. More accurately, we prove the following result.

Theorem 5 *Let \mathcal{C} be an adversary breaking the semantic security of the enhanced key-agreement scheme, with advantage ε and within time t , having observed q_p transcripts and asked q_r Reveal-queries. Then the ECDDH problem can be solved with advantage greater than ε/q_p and within time bound $t + 3q_p \cdot \tau$, where τ accounts for the cost of an elliptic curve operation and is bounded by $\mathcal{O}(k^3)$.*

Proof. Let \mathcal{C} be a passive adversary that can guess the bit b involved in the Test-query. We then construct a distinguisher between the distributions \mathbf{D}_A and \mathbf{R}_A , where

$$\mathbf{R}_A = \{(G_A, Y_A, U, V)\}$$

with all four elements taken at random in E_A and

$$\mathbf{D}_A = \{(G_A, Y_A, U, V)\}$$

with $\log_{G_A}(U) = \log_{Y_A}(V)$.

We run the key generation algorithm and generate two prime numbers p_A, p_B together with two elliptic curves over \mathbb{Z}_{p_A} , and \mathbb{Z}_{p_B} resp., with $p_A - 2$, and $p_B - 2$ resp., elements (also prime). We next show how to use \mathcal{C} to break the ECDDH on E_A : we take the base point of the cryptosystem to be the first coordinate G_A of the input to

\mathcal{C} and we complete the public key of \mathcal{A} by the second coordinate Y_A . This implicitly defines a secret key for \mathcal{A} . The keys of \mathcal{B} are obtained according to the key generation algorithm and thus we know G_B , Y_B and x_B such that $Y_B = x_B \cdot G_B$.

We then can easily simulate the view of the adversary \mathcal{C} :

- one chooses a random $r_A \in \{0, \dots, p_B - 2\}$ and sends $R_B = r_A \cdot G_B$
- one chooses a random $r_B \in \{0, \dots, p_A - 2\}$ and sends $R_A = r_B \cdot G_A$
- one can thus derive the two values

$$K_A = DH_{E_A}(Y_A, R_A) = r_B \cdot Y_A \text{ and } K_B = DH_{E_B}(Y_B, R_B) = r_A \cdot Y_B,$$

and then the shared session key \mathbf{k} from $x_{K_A} \oplus x_{K_B}$.

In order to use the adversary to distinguish the above distributions, we have to randomly choose one index $i \in \{1, \dots, q_p\}$. The i -th execution, instead of being simulated as above, uses the last part (U, V) of the input to \mathcal{C} : R_B and K_B are still simulated as above, but we set $R_A \leftarrow U$ and $K_A \leftarrow V$, and obtain the shared session key \mathbf{k} from $x_{K_A} \oplus x_{K_B}$.

The **Reveal**-queries can easily be simulated by answering the computed shared key. The same way, the **Test**-query can be simulated, according to a random coin b . Then, provided that the **Test**-query is asked to the i -th execution, if the input to \mathcal{C} comes from \mathbf{D}_A , the advantage of \mathcal{C} is exactly the same as in the real game. On the other hand, if the input to \mathcal{C} comes from \mathbf{R}_A , the advantage of \mathcal{C} is exactly 0. Since choice of i is independent of the view of the adversary, the advantage of our distinguisher is ε/q_p , within time $t' = t + 3q_p \cdot \tau$, where τ accounts for the cost of elliptic curve operations needed to compute the data to be handled to \mathcal{C} .

Using an improved reduction, similar to [9] which uses the random self-reducibility of the ECDDH problem, one can state

Theorem 6 *Let \mathcal{C} be an adversary that can break the semantic security of the enhanced key-agreement scheme, with advantage ε and within time t , having observed q_p transcripts and asked q_r **Reveal**-queries, with advantage ε and within time bound t . Then the ECDDH problem can be solved with advantage greater than $e^{-1} \cdot \varepsilon/q_r$ and within time bound $t + 3q_p \cdot \tau$, where τ accounts for the cost of an elliptic curve operation and is bounded by $\mathcal{O}(k^3)$.*

Proof. The proof is similar to the above, but the simulation of the view of the adversary \mathcal{C} is a bit different. With probability $1 - \pi$, one performs the standard simulation, having knowledge of the discrete logarithms in use, and with probability π the simulation uses inputs derived from U and V using random self-reducibility (their discrete logarithms are unknown). A similar analysis can be performed, provided that

all the **Reveal**-queries have been asked to correctly simulated transcript, and that the **Test**-query has been asked on a transcript involving a key derived from (U, V) . This occurs with probability $(1 - \pi)^{q_r} \pi$. If we define $\pi \approx 1/q_r$, then above probability is approximately equal to e^{-1}/q_r , which concludes the proof.

Remark. As repeatedly observed one cannot guarantee that $x_{K_A} \oplus x_{K_B}$ is indistinguishable from a random string of the same length (but only from a string obtained by the same formula from random curve elements K_A and K_B). This is why we need the enhanced version of the scheme.

Active adversaries. In the case of active adversaries, the basic version of the scheme, where the session key is computed as the exclusive or of x_{K_A} and x_{K_B} is not secure. To see this, consider an attacker \mathcal{C} impersonating \mathcal{B} and asking a **Test**-query to \mathcal{A} , once \mathcal{A} has sent his answer. Such attacker can compute x_{K_A} and $x_{K_A} \oplus \mathbf{k}$, where \mathbf{k} has been returned by the **Test**-query. By testing whether $x_{K_A} \oplus \mathbf{k}$ is in \mathcal{X}_B , \mathcal{C} distinguishes a correct key from a random string. Indeed, the former passes the test in all cases, while the latter is successful with probability $\simeq 1/2$. This contradicts formal security. Thus, only the enhanced version of the scheme under review can be secure.

Even for the enhanced scheme, we have been unable to prove anything about the security of the scheme against active attacks. Even if we have not found any attack either, we believe that one cannot claim any security result.

Forward-secrecy. As for the previous scheme, forward-secrecy is not achieved, since the session keys are completely and easily determined from the view of the adversary and the secret keys of the participants. Indeed, an adversary that has stored all the R_A and R_B , can easily derive K_A and K_B , and thus \mathbf{k} , if he later learns x_A and x_B .

4.4.2 One-time keys

Note that one can build a more elaborate protocol, which uses a one-time secret key/public key pair at each execution of the protocol (as presented for the previous proposal). Such extensions are not proposed in documents [15, 16], however, the security proofs would be very similar.

4.4.3 Mutual Authentication

As explained above, in the random oracle model [4], any key agreement scheme can be extended so as to provide mutual authentication by adding one more flow (see figure 1). However this question is not addressed at all in documents [15, 16].

5 Conclusion

Based on our analysis, we find that the basic version of the HDEF-ECDH cannot be proven secure, if the session key is derived as the first coordinate of an elliptic curve point obtained by the Diffie-Hellman function, or even as the exclusive or of two such points. With a more elaborate key derivation function, the scheme is presumably secure, with the proposed parameters, for the foreseeable future. However, based on the submission, we would not recommend the scheme as it is, for the following reasons:

- The submission does not explicitly mention a key derivation function that avoids the direct use of the first coordinate of an elliptic curve point (or the exclusive or of two such points) as the session key.
- The key generation algorithm restricts the class to elliptic curves used by the scheme to a very specific subset of curves with complex multiplication. This is somehow contrary to the current trend, which would recommend having the curve generated at random as much as possible.
- The specification does not include any formal discussion of security. Although we have been able to prove the security of an enhanced version of the scheme against a class of passive adversaries, it does not seem possible to cover the case of active adversaries.
- The scheme does not, as proposed, provide forward secrecy.

We wish to note additionally that we have been able to show that the enhanced version of the scheme could be used securely and provide forward secrecy, in conjunction with digital signatures. Nevertheless, no indication on how to use such setting is included in the specification.

References

- [1] R. Balasubramanian and N. Koblitz. The improbability than an elliptic curve has subexponential discrete log problem under the Menezes-Okamoto-Vanstone algorithm, *J. Cryptology*, 111, (1998), 141–145.
- [2] M. Bellare, R. Canetti, and H. Krawczyk. A Modular Approach to the Design and Analysis of Authentication and Key Exchange Protocols. In *Proc. of the 30th STOC*. ACM Press, New York, 1998.
- [3] M. Bellare, D. Pointcheval, and P. Rogaway. Authenticated Key Exchange Secure Against Dictionary Attacks. In *Eurocrypt '2000*, LNCS 1807, pages 139–155. Springer-Verlag, Berlin, 2000.

- [4] M. Bellare and P. Rogaway. Random Oracles Are Practical: a Paradigm for Designing Efficient Protocols. In *Proc. of the 1st CCS*, pages 62–73. ACM Press, New York, 1993.
- [5] M. Bellare and P. Rogaway. Entity Authentication and Key Distribution. In *Crypto '93*, LNCS 773, pages 232–249. Springer-Verlag, Berlin, 1994.
- [6] M. Bellare and P. Rogaway. Provably Secure Session Key Distribution: the Three Party Case. In *Proc. of the 27th STOC*. ACM Press, New York, 1995.
- [7] S. Cavallar, B. Dodson, A. K. Lenstra, W. Lioen, P. L. Montgomery, B. Murphy, H. te Riele, K. Aardal, J. Gilchrist, G. Guillerm, P. C. Leyland, J. Marchand, F. Morain, A. Muffett, C. Putnam, C. Putnam, P. Zimmermann. Factorization of a 512-Bit RSA Modulus. Eurocrypt'2000, Lecture Notes in Computer Science 1807,(2000), 1–18
- [8] Certicom, Information on the Certicom ECC challenge,
http://www.certicom.com/research/ecc_challenge.html
- [9] J.-S. Coron. On the Exact Security of Full-Domain-Hash. In *Crypto '2000*, LNCS 1880, pages 229–235. Springer-Verlag, Berlin, 2000.
- [10] W. Diffie and M. E. Hellman. New Directions in Cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, November 1976.
- [11] D. Dolev, C. Dwork, and M. Naor. Non-Malleable Cryptography. In *Proc. of the 23rd STOC*. ACM Press, New York, 1991.
- [12] G. Frey, M. Müller, and H. G. Rück. The Tate-Pairing and the Discrete Logarithm Applied to Elliptic Curve Cryptosystems. *IEEE Transactions on Information Theory*, 45:1717–1719, 1999.
- [13] G. Frey and H. G. Rück. A Remark Concerning m -Divisibility and the Discrete Logarithm in the Divisor Class Group of Curves. *Mathematics of Computation*, 62:865–874, 1994.
- [14] R. Harley, D. Doligez, D. de Rauglaudre, X. Leroy. Elliptic Curve Discrete Logarithms: ECC2K-108,
<http://cristal.inria.fr/~harley/ecdl7/>
- [15] HDEF-ECDH. Cryptographic techniques specifications
- [16] HDEF-ECDH. Self-evaluation report

- [17] R. Gallant, R. Lambert and S.A. Vanstone. Improving the parallelized Pollard lambda search on binary anomalous elliptic curves, *Mathematics of Computation*, 69, (2000), 1699–1705.
- [18] S. Goldwasser and S. Micali. Probabilistic encryption, *Journal of Computer and System Science* 28, (1984), 270–299.
- [19] S. Goldwasser, S. Micali, and R. Rivest. A “Paradoxical” Solution to the Signature Problem. In *Proc. of the 25th FOCS*, pages 441–448. IEEE, New York, 1984.
- [20] S. Goldwasser, S. Micali, and R. Rivest. A Digital Signature Scheme Secure Against Adaptative Chosen-Message Attacks. *SIAM Journal of Computing*, 17(2):281–308, April 1988.
- [21] IEEE standard specifications for public-key cryptography. IEEE Computer Society (2000).
- [22] R. Impagliazzo and D. Zuckermann, How to rectle random bits, *30th annual symposium on foundations of computer science*, (1989), 248–253.
- [23] A.K. Lenstra and E. Verheul. Selecting cryptographic key sizes, PKC’2000, Lecture Notes in Computer Science 1751,(2000), 446–465.
- [24] R. Lercier and F. Morain, Counting the number of points on elliptic curves over finite fields: strategies and perormances, Eurocrypt’ 95, Lecture Notes in Computer Science 921, (1995), 79–94.
- [25] C.H Lim and P.J. Lee. A key recovery attack on discrete log based schemes using a prime order subgroup, Crypto ’97, Lecture Notes in Computer Science 1294, (1997), 249–263.
- [26] M. Naor and O. Reingold, Number-theoretic Constructions of Efficient Pseudo-random Functions, *38-th annual symposium on foundations of computer science*, (1997), 458–467.
- [27] J. Pollard, Monte Carlo methods for index computation mod p, *Mathematics of Computation*, 32, (1978), 918–924.
- [28] H.G. Rück. On the discrete logarithm in the divisor class group of curves. Preprint (1997).
- [29] T. Satoh, K. Araki. Fermat quotients and the polynomial time discrete log algorithm for anomalous elliptic curves (1997), to appear in *Commentarii Math. Univ. St Pauli*.

- [30] N.P. Smart. The discrete logarithm problem on elliptic curves of trace one. *J. Cryptology*, 12, (1999), 141–151.
- [31] I.A. Semaev. Evaluation of discrete logarithms in a group of p -torsion points of an elliptic curve of characteristic p . *Math. Comp.*, 67 (1998), 353–356.
- [32] V. Shoup. On Formal Models for Secure Key Exchange. Technical Report RZ 3120, IBM Research, April 1999.
- [33] V. Shoup and T. Schweinberger, ACE Encrypt: The Advanced Cryptographic Engine' public key encryption scheme, Manuscript, March 2000. Revised, August 14, 2000.
- [34] P.C. van Oorschot and M. J. Wiener. Parallel collision search with cryptanalytic applications, *J. Cryptology*, 12, (1999), 1–28.
- [35] M. J. Wiener and R.J. Zuccherato. Fast attacks on elliptic curve cryptosystems, SAC'98, Lecture Notes in Computer Science 1556, (1999), 190–200.