

Secure Hashed Diffie-Hellman over Non-DDH Groups*

Rosario Gennaro[†]

Hugo Krawczyk[‡]

Tal Rabin[§]

January 10, 2006

Abstract

We show that in applications that use the Diffie-Hellman (DH) transform but take care of hashing the DH output (as required, for example, for secure DH-based encryption and key exchange) the usual requirement to work over a DDH group, i.e., a group in which the Decisional Diffie-Hellman assumption holds, can be relaxed to only requiring that the DH group contains a large enough DDH subgroup. In particular, this implies the security of (hashed) Diffie-Hellman over non-prime order groups such as Z_p^* . Moreover, our results indicate that one can work directly over Z_p^* without requiring any knowledge of the prime factorization of $p-1$ and without even having to find a generator of Z_p^* . These results are obtained via a general characterization of DDH groups in terms of their DDH subgroups, and a relaxation (called t -DDH) of the DDH assumption via computational entropy. We also show that, under the short-exponent discrete-log assumption, the security of the hashed Diffie-Hellman transform is preserved when replacing full exponents with short exponents.

1 Introduction

The Diffie-Hellman Transform and DDH Assumption. The *Diffie-Hellman transform* is one of the best-known and fundamental cryptographic primitives. Its discovery by Whitfield Diffie and Martin Hellman [DH76] revolutionized the science of cryptography and marked the birth of Modern Cryptography. Even today, almost 30 years later, the Diffie-Hellman (or DH for short) transform remains the foundation of some of the most basic and widely used cryptographic techniques. In particular, it underlies the Diffie-Hellman key exchange and the ElGamal encryption scheme [ElG85], and is used over a large variety of mathematical groups. In its basic form the DH transform maps a pair of elements g^a, g^b drawn from a cyclic group G generated by the element g into the group element g^{ab} . (Here we use the exponential notation that originates with multiplicative groups but our treatment, which is generic in nature, applies equally to additive groups such as Elliptic Curves.) The usefulness of this transform was originally envisioned under the conjecture, known as the *Computational Diffie-Hellman (CDH)* assumption, that states the infeasibility of computing the value g^{ab} given only the exponentials g^a and g^b . Namely, the value g^{ab} should be computable only by those knowing one of the exponents a or b . Note that the CDH assumption

*A conference version appears in Eurocrypt'2004.

[†]IBM T.J.Watson Research Center, PO Box 704, Yorktown Heights, NY 10598, USA. E-mail: rosario@watson.ibm.com.

[‡]IBM T.J.Watson Research Center, PO Box 704, Yorktown Heights, NY 10598, USA. E-mail: hugo@ee.technion.ac.il.

[§]IBM T.J.Watson Research Center, PO Box 704, Yorktown Heights, NY 10598, USA. E-mail: talr@watson.ibm.com.

implies the difficulty of computing discrete logarithms over the group G (the converse, however, is unknown for most practical groups).

Over time it was realized that the CDH assumption is insufficient to guarantee the security of most DH applications (in particular those mentioned above). For this reason a much stronger assumption was introduced: the *Decisional Diffie-Hellman (DDH)* assumption postulates that given the values g^a and g^b not only it is computationally hard to derive the value g^{ab} but even the seemingly much easier task of distinguishing g^{ab} from random group elements is infeasible [Bra93] (see [Bon98] for a survey on the DDH assumption). On the basis of this assumption one can consider the DH transform as a good generator of pseudorandomness as required in key-exchange, encryption and other cryptographic applications. Hereafter we refer to groups in which the DDH assumption holds as *DDH groups*. The need to rely on the DDH disqualifies many natural groups where the assumption does not hold. For example, any group whose order is divisible by small factors, such as the classic groups Z_p^* of residues modulo a large prime p ; in this case the group's order, $p - 1$, is always divisible by 2 (for random p , $p - 1$ is very likely to have additional small factors). Due to the *perceived* need to work over DDH groups it is often recommended in the cryptographic literature that one work over subgroups of large prime order where the DDH assumption is believed to hold.

The Need for Hashing the Diffie-Hellman Result. Interestingly, the DDH assumption, while apparently necessary, turns out to be insufficient for guaranteeing the security of some of the most basic applications of the DH transform. Consider for example the ElGamal encryption scheme: Given a public key $y = g^a$ (for secret a), a message $m \in G$ is encrypted by the pair (g^b, my^b) where the value b is chosen randomly anew for each encryption. In this case, the DDH assumption guarantees the semantic security [GM84] of the scheme (against chosen-plaintext attacks) *provided that the plaintexts m are elements of the group G* . However, if the message space is different, e.g. the set of strings of some length smaller than $\log |G|$, then the above encryption scheme becomes problematic. First of all, you need to encode messages m as group elements in G and that could be cumbersome. If G is a subgroup of prime order of Z_p^* , a naive (and common) approach would be to trivially encode m as an integer and perform the multiplication my^b modulo p . But now the scheme is *insecure even if the group G does satisfy the DDH assumption*. A good illustration of the potential weaknesses of this straightforward (or “textbook”) application of ElGamal is presented in [BJN00]. It is shown that if the space of plaintexts consists of random strings of length shorter than $\log |G|$ (e.g., when using public key encryption to encrypt symmetric keys) the above scheme turns out to be insecure even under a ciphertext-only attack and, as said, even if the group G is DDH. For example, if the plaintexts to be encrypted are keys of length 64, an attacker that sees a ciphertext has a significant probability of finding the plaintext with a work factor in the order of 2^{32} operations and comparable memory; for encrypted keys of length 128 the complexity of finding the key is reduced to 2^{64} .

A general and practical approach to solving these serious security weaknesses is to avoid using the DH value itself to “mask” m via multiplication, but rather to *hash* the DH value g^{ab} to obtain a pseudorandom key K of suitable length which can then be used to encrypt the message m under a particular encryption function (in particular, K can be used as a one-time pad). In this case the hash function is used to *extract the (pseudo) randomness* present in the DH value. Suitable hash functions with provable extraction properties are known, for example *universal hash functions* [CW79, HILL99]. The above considerations are common to many other applications of the DH transform, including encryption schemes secure against chosen-ciphertext attacks [CS98] and, most prominently, the Diffie-Hellman key-exchange protocol (in the latter case one should not use the DH output as a cryptographic key but rather derive the agreed shared keys via a hashing of the DH

result); see Section 4.2 for a discussion on how these applications choose a random hash function out of a given family. For additional examples and justification of the need for hashing the DH output see [Bon98, NR97, CS98, ABR01]. In the sequel we refer to the combination of the DH transform with a (universal) hash function as the *hashed DH transform*.

1.1 Our Results

The Security of the Hashed DH Transform over non-DDH Groups. In light of the need to hash the DH value, some natural questions arise: when applying the hashed DH transform, is it still necessary to work over groups where the DDH assumption holds, or can this requirement be relaxed? Can one obtain a secure (hashed) DH transform over a non-DDH group, and specifically, is doing hashed DH over Z_p^* secure? In this paper we provide answers to these questions. Our main result can be informally stated as follows: *For any cyclic group G , applying the hashed DH transform over G has the same security as applying the hashed DH transform directly over the maximal DDH subgroup of G .* In particular, one can obtain secure applications of the hashed DH transform over non-DDH groups; the only requirement is that G contain a (sufficiently large) DDH subgroup (see below for the exact meaning of “sufficiently large” and other parameter size considerations). A significant point is that we are only concerned with the *existence* of such a subgroup; there is no need to know the exact size or structural properties of, nor to be able to construct, this specific (maximal) DDH subgroup.

A particularly interesting consequence of the above result is that assuming that DDH holds on large subgroups of Z_p^* (we will see later that it is sufficient to assume that DDH holds on large prime-order subgroups of Z_p^*), one can build secure (hashed) DH applications working directly over Z_p^* , where p is an unconstrained random prime. Only the length of the prime is specified, while other common requirements such as the knowledge of the partial or full factorization of $p - 1$, insisting that $p - 1$ has a prime factor of a particular size, or disqualifying primes for which $(p - 1)/2$ has a smooth part, are all avoided here. Moreover, there is no need to find a generator of Z_p^* ; instead, the plausible assumption that $p - 1$ has sufficiently large prime factors allows us to use a randomly chosen element r from Z_p^* in lieu of a generator of Z_p^* . In this case the group $G = \langle r \rangle$ is guaranteed to have a large enough DDH subgroup, and therefore the hashed DH transform over G is secure (this is true even if the order of r has small factors or if it misses some prime divisors of $p - 1$). Note that avoiding the need to find a generator for Z_p^* allows us to work with primes p with unknown factorization of $p - 1$ (which is otherwise required to find a Z_p^* generator).

The t -DDH Assumption. In order to prove our main result (i.e., that the hashed DH transform is secure over any group G , not necessarily a DDH group, that contains a large enough DDH subgroup), we introduce a relaxation of the DDH assumption which we call the *t -DDH assumption*. Informally, a group G satisfies the t -DDH assumption (where $0 \leq t \leq |G|$) if given the pair (g^a, g^b) (where g is a generator of G) the value g^{ab} contains t bits of computational entropy. The notion of computational entropy, introduced in [HILL99], captures the amount of computational hardness present in a probability distribution. In other words, we relax the “full hardness” requirement at the core of the DDH assumption, and assume partial hardness only. Moreover, we do not care about the exact subsets of bits or group elements where this hardness is contained, but only assume their existence. On this basis, and using the entropy-smoothing theorem from [HILL99] (also known as the leftover hash lemma), we obtain a way to efficiently transform (via universal hashing) DH values over groups in which the t -DDH assumption holds into shorter outputs that are computationally indistinguishable from the uniform distribution. The maximal length of (pseudorandom) strings that one can obtain as output from the hashed DH transform depends on the maximal value of t for

which the t -DDH holds in G . In particular, in order to be 2^{-k} -computationally close to uniform one can output up to $t - 2k$ pseudorandom bits (e.g., to produce 128-bit keys with a security parameter of $k = 80$ the group G should be 288-DDH, while for $k = 128$, G is to be 384-DDH).

After defining the t -DDH assumption and showing its usefulness in extracting random bits from t -DDH groups, we show that *if G contains a DDH subgroup of order m then G is $\log(m)$ -DDH*. This forms the basis for our main result as stated above. Indeed, it suffices that G has a suitably large-order DDH subgroup to ensure that hashing the DH output results in pseudorandom outputs of the required length. Again, it is important to stress that we do not need to know the specific DDH subgroup or its order, only (assume) its existence.

A Direct Product Characterization of the DDH Assumption. A further contribution of our work is in providing a characterization of the DDH assumption in a given group in terms of its DDH subgroups. Specifically, we show that *a group is DDH if and only if it is the direct product of (disjoint) prime power DDH groups*. In other words, a group G is DDH if and only if all its prime power subgroups are DDH. Moreover, for any cyclic group G , the maximal DDH group in G is obtained as the product of all prime power DDH subgroups in G . Beyond its independent interest, this result plays a central role in our proof that the hashed DH transform over Z_p^* is secure as long as the DDH assumption holds in the subgroups of Z_p^* of large prime order. In particular, this allows us to expand significantly the groups in which one can work securely with the hashed DH transform without having to strengthen the usual assumption that DDH holds in large prime order subgroups.

Some Practical Considerations. Beyond the theoretical interest in understanding the role of the DDH assumption and proving the usefulness of relaxed assumptions, our results provide a justification of the use of non-DDH groups in practical applications of the DH transform as long as these groups contain a large enough prime-order subgroup *and the application takes care of hashing the DH output*. One interesting practical example is the IPsec’s Key Exchange (IKE) protocol [RFC2409, IKEv2] that uses a Diffie-Hellman exchange to negotiate shared keys but is careful to first hash the DH value (see [Kra03]).¹ In addition, and as pointed out before, our results also show that under the sole assumption that the DDH holds in groups of large prime order one can work directly over Z_p^* for a random prime p , without having to know the factorization of $p - 1$ and without having to find a generator of Z_p^* . Moreover, the ability to work over non-prime order groups has the benefit of eliminating the attacks on the DH transform described in [LL97], without having to search for primes of a special form (and without necessitating special parameter checks when certifying public keys [LL97]).

Short-Exponent Diffie-Hellman. One important practical consideration is the length of exponents used when applying the DH transform. Full exponents when working over Z_p^* are, typically, of size 1024 or more. Even if one works over a prime-order subgroup, one still needs to use relatively large orders (e.g. 288-bit long primes), with their correspondingly large exponents, to ensure a hashed output (say of 128 bits) that is indistinguishable from uniform. (This requirement for sufficiently large computational entropy is often overlooked; indeed, the usual practice of using 160-bit prime-order groups, which originates with Schnorr’s signatures, may not be appropriate for hashed DH-type applications.)

Motivated by the significant cost of exponentiation using long exponents, we investigate whether one can use short exponents (e.g. as in [RFC2409]) and still preserve the security of the hashed

¹ In IKE, the family of hash functions used for extracting a pseudorandom key from the DH value are implemented using common pseudorandom function families keyed with random, but known, keys. The randomness extraction properties of the latter families are studied in [DGHKR04].

DH transform. An obviously necessary requirement for the short exponent practice to be secure is the assumption that the discrete log problem is hard when exponents are restricted to a short length (say of s bits). We show that this requirement (referred to as the s -DLSE assumption) is sufficient for the secure use of short exponents in the setting of the DH transform; more precisely, we prove (based on [Gen00]) that if the s -DLSE assumption holds in a group G , then the hashed DH transform in G is as secure with full exponents as with s -bit exponents. As a consequence, one can analyze the security of the hashed DH transform in the group G with full exponents and later replace the full exponents with much shorter ones without sacrificing security. In this case the important parameter is the length s ; we note that the appropriate value of s depends on the underlying group. See [vOW96] for an extensive study of the plausible value of s for different groups.

Paper’s Organization. Section 2 introduces some of the notation and basic notions underlying our technical treatment (including the definition of DDH groups and a discussion of our concrete security treatment). In Section 3 we prove the Direct Product DDH Characterization Theorem. In Section 4 we introduce the t -DDH Assumption and its application to the hashed DH transform, and prove the central Max-Subgroup Theorem. In Section 5 we investigate the security of the hashed DH transform when using short exponents. We conclude in Section 6 by describing the applicability of our results to the hashed DH transform over non-DDH groups.

2 Preliminaries

Throughout the paper we use the following notation. Let \mathcal{D} be a probability distribution over a set A . By $x \in_{\mathcal{D}} A$ we mean that x is chosen in A according to the distribution \mathcal{D} , and with $x \in_R S$ we denote the choice of x with uniform distribution over the set S . When m is an integer we use the notation $|m|$ to indicate the length of m in bits.

2.1 The Decisional Diffie-Hellman Assumption

Let G be a cyclic group of order m generated by an element g . Consider the following problem: Given a pair g^a, g^b compute the value g^{ab} . If this problem is intractable for the group G then we say that the Computational Diffie-Hellman (CDH) assumption holds over G .

A much stronger, but also more useful, assumption is the following. Consider the set $G^3 = G \times G \times G$ and the following two probability distributions over it:

$$\mathcal{R}_G = \{(g^a, g^b, g^c) \text{ for } a, b, c \in_R [0..m]\}$$

and

$$\mathcal{DH}_G = \{(g^a, g^b, g^{ab}) \text{ for } a, b \in_R [0..m]\}$$

Definition 1 *We say that the (S, ϵ) Decisional Diffie-Hellman (DDH) Assumption holds over $G = \langle g \rangle$ (alternatively, that G is a (S, ϵ) DDH group) if the two distributions \mathcal{R}_G and \mathcal{DH}_G are (S, ϵ) -indistinguishable.*

The notion of (S, ϵ) indistinguishability is recalled in Appendix A.

Informally, the DDH assumption states that no feasible algorithm (a “distinguisher”) has a significant probability of deciding correctly whether the third element of the triple (g^a, g^b, g^c) is

the result of the Diffie-Hellman transform applied to g^a, g^b or a randomly chosen group element. Clearly this is a much weaker requirement than computing the value g^{ab} from g^a, g^b , and therefore, as a general hardness assumption, DDH is stronger than the CDH. We note that, in principle, the DDH assumption could hold for a group G with respect to a generator g but not with respect to another generator g' . As we will see in Section 2.2 this is *not* the case when using a concrete security formalism as in this paper.

Example 1: *A group where the DDH assumption does not hold.* Consider the group $G = Z_p^*$ for a prime p . Since testing for quadratic residuosity over Z_p^* is easy, by computing the Legendre symbol $(\frac{\cdot}{p})$, then we immediately get a distinguisher against DDH in this group: by mapping the Legendre symbol of 1 (i.e., quadratic residues) to 0, and the Legendre symbol of -1 to 1, we can simply check that $(\frac{g^a}{p})(\frac{g^b}{p}) = (\frac{g^c}{p})$, and output “ \mathcal{DH} ” if it holds and “ \mathcal{R} ” otherwise. Clearly, if the triple is a legal DH triple then the distinguisher outputs \mathcal{DH} with probability 1, while in the other case the probability is only 1/2.

Example 2: *A group where the DDH is conjectured to hold.* Let p, q be primes such that q divides $p - 1$, and p and q are “large” (say, $|q| = 1024$ and $|p| \geq 160$). Let G be the subgroup of order q in Z_p^* . In this case no efficient DDH distinguisher for G is known.

In informal statements and discussions (as in the above examples) we sometimes omit the (S, ϵ) parameters from the DDH assumption, e.g., we say that “ G is DDH”, in which case the intent is, as before, that no “feasible-size” distinguisher succeeds in the above task with “significant” probability. (See the discussion below on the (in)dependence of the DDH assumption on a specific generator g .) Also, when the group G is clear from the context we often omit the subscript G in the notation of the two distributions \mathcal{R}_G and \mathcal{DH}_G .

2.2 On Concrete Security and Non-Uniformity

In Definition 1 and throughout this paper, our treatment of computational difficulty follows the “concrete security” approach in which concrete (numerical) bounds are established on the resources (time/space) and success probability of algorithms (or “attackers”). This approach has the advantage of providing very clear quantitative results which, in particular, highlight the exact cost of security reductions (the downside is making the formal statements of results somewhat more cumbersome). Maybe more significant is the fact that concrete security allows us to talk about individual, finite objects, such as a single finite group. In contrast, when stating results in terms of polynomial-time one must resort to asymptotics and infinite families, something that is not very well suited to the typical use of Diffie-Hellman groups which are usually defined and fixed in advance for repeated use by many applications and in many sessions (e.g., IKE, SSH, etc.)

Moreover, for some of our results it is natural to speak about an individual group and its individual subgroups; trying to map this into an asymptotic presentation just obscures the results. Another fundamental aspect of concrete security that fits well into our setting is that concrete security captures a non-uniform notion of computation. For example, when solving a problem for a *particular* group, an algorithm can always include an “auxiliary input” that helps solving the problem over that particular group. To emphasize the non-uniformity behind the concrete security approach, we talk about *circuits* with concrete size S . (We note that while the concrete security literature usually states bounds in terms of time, talking of “size” rather than “time” is more accurate since even if one talks about time, the size of algorithms must be taken into account, or else some problems can be trivialized via huge pre-computed tables.)

In our context, it is interesting to consider two consequences of the non-uniform aspect of concrete security. First, an algorithm that works for a specific group G can always include (as part of the algorithm “code” or auxiliary input) the order of the group or even the factorization of $|G|$. Hence, in this setting we need to work under the assumption that an attacker may have such information. For example, in Theorem 2 we assume that the attacker may know the group’s order m and may even know the factorization of m (both of which can be included in the “code”, or “auxiliary input” of the algorithm). This is a non-trivial example in the sense that the theorem does not necessarily hold for groups in which the factorization of $|G|$ is secret, e.g., $G = Z_N^*$ where $N = pq$ is a modulus of unknown factorization.

Second, in a uniform (asymptotic) treatment one cannot talk about the DDH assumption holding in a group G ; instead the validity of DDH may depend on a specific generator g of G . In contrast, in a non-uniform setting as ours if the DDH assumption holds in G with respect to a generator g then it holds with respect to any other generator g' of G and hence the DDH property is independent of the specific generator (indeed, if DDH is easy in G with respect to a generator h then it is also easy with respect to any other generator g : just give the distinguisher with respect to g the value $\log_h(g)$ as auxiliary input). This independence from specific generators is more “elegant” and particularly useful when stating several of our results.

3 A Direct-Product DDH Characterization

The following theorem provides a full characterization of DDH groups in terms of their prime (power) order subgroups. As remarked in Section 2.2, the proof of this theorem assumes that the distinguisher is given the factorization of $\text{ord}(G)$. The theorem is first formulated informally without concrete bounds; these are stated and proved in the following Lemmas.

Theorem 2 (Direct Product Characterization Theorem – informal.) *A cyclic group G is DDH if and only if all its prime-power order subgroups are DDH.*

The precise meaning of the above Theorem is specified by the following Lemmas 3 and 4. Below we denote with exp_G the size of the circuit that computes exponentiations in G .

Lemma 3 *Let G be a group of order $m = m_1 m_2$, and let G_1 be the subgroup of G of order m_1 . If the (S, ϵ) DDH holds in G then the (S_1, ϵ) DDH holds in G_1 where $S_1 = S - 3\text{exp}_G$.*

Proof Let G be a DDH (cyclic) group of order $m = m_1 m_2$, and let G_1 be a subgroup of G of order m_1 . Let g be a generator of G and $g_1 = g^{m_2}$ be a generator of G_1 . Assume by contradiction that the (S_1, ϵ) DDH does not hold in G_1 , i.e., there is a distinguisher D_1 of size $\leq S_1$ that upon receiving a triple $(A_1 = g_1^{a_1}, B_1 = g_1^{b_1}, C_1 = g_1^{c_1}) \in G_1^3$, can distinguish whether it came from the distribution \mathcal{R}_{G_1} or \mathcal{DH}_{G_1} with advantage $> \epsilon$. We build a distinguisher D of size $\leq S$ for G which distinguishes between the distributions \mathcal{DH}_G and \mathcal{R}_G with the same probability ϵ . This contradicts the assumption that (S, ϵ) DDH holds in G .

Upon receiving a triple $(A = g^a, B = g^b, C = g^c)$, where $a, b \in_R Z_{m_1 m_2}$ and c is either the product of ab or picked uniformly at random in $Z_{m_1 m_2}$, the distinguisher D :

1. Computes (A_1, B_1, C_1) by setting $A_1 = A^{m_2}, B_1 = B^{m_2}$, and $C_1 = C^{m_2}$.
2. Passes the triple (A_1, B_1, C_1) to D_1

3. Outputs the same output bit as D_1 .

Note that by construction the values A_1, B_1, C_1 equal $g_1^{a_1}, g_1^{b_1}, g_1^{c_1}$, respectively, where $a_1 = a \bmod m_1, b_1 = b \bmod m_1, c_1 = c \bmod m_1$. Since $a, b \in_R Z_{m_1 m_2}$ then $a_1, b_1 \in_R Z_{m_1}$. Also, if $c = ab \bmod m_1 m_2$ then $c_1 = a_1 b_1 \bmod m_1$, while if $c \in_R Z_{m_1 m_2}$ then $c_1 \in_R Z_{m_1}$ (independently of a_1, b_1). In other words, whenever the triple (A, B, C) is distributed according to \mathcal{DH}_G then the triple (A_1, B_1, C_1) is distributed according to \mathcal{DH}_{G_1} , while if (A, B, C) is distributed according to \mathcal{R}_G then the triple (A_1, B_1, C_1) is distributed according to \mathcal{R}_{G_1} . Therefore, D distinguishes between the distributions \mathcal{DH}_G and \mathcal{R}_G with the same probability ϵ that D_1 distinguishes between \mathcal{DH}_{G_1} and \mathcal{R}_{G_1} . Notice that the size of D is $\leq S$ since all it does is computing three exponentiations in G and then invoke D_1 . ■

Lemma 4 *Let G be a cyclic group of order $m = m_1 m_2$, where $(m_1, m_2) = 1$, and let G_1 and G_2 be the subgroups of G of orders m_1, m_2 resp. If (S_1, ϵ_1) DDH holds in G_1 and (S_2, ϵ_2) DDH holds in G_2 then (S, ϵ) DDH holds in G where $S = \min(S_1, S_2) - 9\text{exp}_G$ and $\epsilon = \epsilon_1 + \epsilon_2$.*

Proof Let g, g_1, g_2 be generators of G, G_1 , and G_2 , respectively; in particular, $g_1 = g^{m_2}$ and $g_2 = g^{m_1}$. Given a triple $t_1 = (A_1 = g_1^{a_1}, B_1 = g_1^{b_1}, C_1 = g_1^{c_1}) \in G_1^3$ and a triple $t_2 = (A_2 = g_2^{a_2}, B_2 = g_2^{b_2}, C_2 = g_2^{c_2}) \in G_2^3$ we define the following transformation T which ‘‘lifts’’ this pair of triples into a triple in G^3 . (T is the standard isomorphism between the group G and its product group representation as determined by the Chinese Remainder Theorem.) On input t_1, t_2 , $T(t_1, t_2)$ outputs a triple $(A = g^a, B = g^b, C = g^c) \in G^3$ defined as follows:

1. Let r_1, r_2 be such that $r_1 m_1 + r_2 m_2 = 1$ (i.e., $r_1 = m_1^{-1} \bmod m_2$ and $r_2 = m_2^{-1} \bmod m_1$)
2. Set $A = A_1^{r_2} A_2^{r_1} = g^{a_1 m_2 r_2 + a_2 m_1 r_1} \in G$, i.e., $a = a_1 m_2 r_2 + a_2 m_1 r_1 \bmod m$
3. Set $B = B_1^{r_2} B_2^{r_1} = g^{b_1 m_2 r_2 + b_2 m_1 r_1} \in G$, i.e., $b = b_1 m_2 r_2 + b_2 m_1 r_1 \bmod m$
4. Set $C = C_1^{m_2 r_2^2} C_2^{m_1 r_1^2} = g^{c_1 m_2^2 r_2^2 + c_2 m_1^2 r_1^2} \in G$, i.e., $c = c_1 m_2^2 r_2^2 + c_2 m_1^2 r_1^2 \bmod m$

Note the following facts about the triple (A, B, C) which result from the above transformation:

Fact 1 If $a_1, b_1 \in_R Z_{m_1}$, and $a_2, b_2 \in_R Z_{m_2}$, then $a, b \in_R Z_m$.

Fact 2 $c - ab \equiv c_1 - a_1 b_1 \bmod m_1$ and $c - ab \equiv c_2 - a_2 b_2 \bmod m_2$

Fact 3 Following Facts 1 and 2, if the triple t_1 is chosen according to distribution \mathcal{DH}_{G_1} and t_2 according to distribution \mathcal{DH}_{G_2} , then the triple (A, B, C) is distributed according to the distribution \mathcal{DH}_G . Similarly, if t_1, t_2 are distributed according to \mathcal{R}_{G_1} and \mathcal{R}_{G_2} , respectively, then (A, B, C) is distributed according to \mathcal{R}_G .

For probability distributions $\mathcal{P}_1, \mathcal{P}_2$ we denote by $T(\mathcal{P}_1, \mathcal{P}_2)$ the probability distribution induced by the random variable $T(x_1, x_2)$ where x_1, x_2 are random variables distributed according to $\mathcal{P}_1, \mathcal{P}_2$, respectively, and T is the above defined transform. Using this notation and Fact 3 we get: $\mathcal{DH}_G = T(\mathcal{DH}_{G_1}, \mathcal{DH}_{G_2})$ and $\mathcal{R}_G = T(\mathcal{R}_{G_1}, \mathcal{R}_{G_2})$. Let us now consider the ‘‘hybrid’’ probability distribution $T(\mathcal{R}_{G_1}, \mathcal{DH}_{G_2})$.

Note that this distribution is $(S_1 - 9\text{exp}_G, \epsilon_1)$ indistinguishable from $T(\mathcal{DH}_{G_1}, \mathcal{DH}_{G_2})$. Indeed, since the distribution \mathcal{DH}_{G_2} is efficiently samplable (it costs 3 exponentiations to sample it) and

the transformation T is efficiently computable (it costs 6 exponentiations to compute it), then one can transform any (S, ϵ) distinguisher between the above two distributions into a $(S + 9\text{exp}_G, \epsilon)$ distinguisher between \mathcal{R}_{G_1} and \mathcal{DH}_{G_1} . Thus if $S < S_1 - 9\text{exp}_G$ and $\epsilon > \epsilon_1$ we have a distinguisher for \mathcal{R}_{G_1} and \mathcal{DH}_{G_1} of size $\leq S_1$ that distinguishes with probability $> \epsilon_1$, in contradiction to the hypothesis that G_1 is a (S_1, ϵ_1) DDH group.

Similarly, we have that the hybrid distribution $T(\mathcal{R}_{G_1}, \mathcal{DH}_{G_2})$ is $(S_2 - 9\text{exp}_G, \epsilon_2)$ indistinguishable from $T(\mathcal{R}_{G_1}, \mathcal{R}_{G_2})$.

By invoking the triangle inequality for computational indistinguishability (see Prop. 22 in Appendix A) we have that \mathcal{R}_G and \mathcal{DH}_G are (S, ϵ) indistinguishable where $S = \min(S_1 - 9\text{exp}_G, S_2 - 9\text{exp}_G) = \min(S_1, S_2) - 9\text{exp}_G$ and $\epsilon = \epsilon_1 + \epsilon_2$ as required. ■

Discussion (On prime-power subgroups). We note that the result summarized in Theorem 2 is actually asymmetric. In the “only if” direction (Lemma 3) all subgroups are guaranteed to be DDH, while for the “if” direction (Lemma 4) we need the DDH assumption on *prime-power order* subgroups. The reason for the latter is the condition $(m_1, m_2) = 1$ in the statement and proof of Lemma 4. A natural question is whether one can strengthen the latter lemma and prove a similar result for factors m_1, m_2 which are not necessarily co-prime. More specifically, we are interested in the following. Let G be a cyclic group of order q^2 for prime q , and let H be the subgroup of G of order q . Assume that H is DDH. Does this imply that G is DDH as well? This was posed as an open question in an earlier version of this paper. Recently, Don Coppersmith has built [Cop04] an ingenious counter-example, namely, a cyclic group G of order q^2 which contains a subgroup H of order q , such that H is believed to be DDH but G is trivially not DDH. We present Coppersmith’s example in Appendix B. It is still interesting to settle this question for specific families of groups (e.g., the subgroups of Z_p^* for prime p). In general, how plausible is it to assume the DDH assumption in prime-power order subgroups of Z_p^* ?

We end this section by mentioning a result by Maurer and Wolf (Corollary 5, [MW96]) that shows a relation between the hardness of the (computational) Diffie-Hellman problem in a cyclic group and the hardness of this problem in some of its subgroups. More specifically, they prove that if G is a cyclic group and H a subgroup such that the index $|G|/|H|$ is smooth then the CDH problem in G and H are polynomial-time equivalent.

4 The t -DDH Assumption and the Hashed DH Transform

In this section we introduce an intractability assumption that is, in general, weaker than the DDH assumption, yet it suffices for ensuring DH outputs from which a large number of pseudorandom bits can be extracted. We start by recalling the notions of computational entropy and entropy smoothing. We use the notations introduced at the end of Section 1.

4.1 Computational Entropy and Entropy Smoothing

Definition 5 Let \mathcal{X} be a probability distribution over A . The *min-entropy* of \mathcal{X} is the value

$$\text{min-ent}(\mathcal{X}) = \min_{x \in A: \text{Prob}_{\mathcal{X}}[x] \neq 0} (-\log(\text{Prob}_{\mathcal{X}}[x]))$$

Note that if \mathcal{X} has min-entropy t then for all $x \in A$, $\text{Prob}_{\mathcal{X}}[x] \leq 2^{-t}$.

The notion of min-entropy provides a measurement of the amount of randomness present in a probability distribution. Indeed, the Entropy Smoothing Theorem (see below) shows that if \mathcal{X} has

min-entropy t it is possible to construct from \mathcal{X} an (almost) uniform distribution over (almost) t bits, by simply hashing elements chosen according to \mathcal{X} . The basic hashing tool to do this uses the following notion of universal hashing.

Definition 6 Let \mathcal{H} be a family of functions, where each $H \in \mathcal{H}$ is defined as $H : A \rightarrow \{0, 1\}^m$. We say that \mathcal{H} is a family of (pairwise-independent) universal hash functions if, for all $x, x' \in A$, $x \neq x'$, and for all $a, a' \in \{0, 1\}^m$ we have

$$\text{Prob}_{H \in \mathcal{H}}[H(x) = a \text{ and } H(x') = a'] = 2^{-2m}.$$

That is, a randomly chosen H will map any pair of distinct elements independently and uniformly.

Our techniques use as a central tool the following Entropy Smoothing Theorem from [HILL99] (see also [Gol01, Lub96]), also known as the “Leftover Hash Lemma”. The definition of statistical distance used in the theorem’s statement is recalled in Appendix A.

Theorem 7 (Entropy Smoothing Theorem [HILL99]) Let t be a positive integer and let \mathcal{X} be a random variable defined on $\{0, 1\}^n$ such that $\text{min-ent}(\mathcal{X}) > t$. Let $k > 0$ be an integer parameter. Let \mathcal{H} be a family of universal hash functions such that $\forall h \in \mathcal{H}, h : \{0, 1\}^n \rightarrow \{0, 1\}^{t-2k}$. Let \mathcal{U} be the uniform distribution over $\{0, 1\}^{t-2k}$. Then, the distributions $[\langle h(\mathcal{X}), h \rangle]_{h \in_R \mathcal{H}}$ and $[\langle \mathcal{U}, h \rangle]_{h \in_R \mathcal{H}}$ have statistical distance at most $2^{-(k+1)}$.

Thus, the Entropy Smoothing Theorem guarantees that if \mathcal{X} is a probability distribution over A with min-entropy of at least t , and \mathcal{H} a family of universal hash functions from A to $\{0, 1\}^{t-2k}$, then the random variable $h(x)$, where $h \in_R \mathcal{H}$ and x is chosen according to the distribution \mathcal{X} , is “almost” uniformly distributed over $\{0, 1\}^{t-2k}$ even when the hash function h is given. Here, “almost” means a statistical distance of at most 2^{-k-1} .

The following notion represents a computational analogue of the notion of min-entropy and was introduced in [HILL99]. We recall it here (under a concrete security formulation) for completeness and because it is implicit in our definition of the t -DDH assumption in the next sub-section.

Definition 8 A probability distribution \mathcal{Y} has (S, ϵ) computational entropy t if there exists a probability distribution \mathcal{X} such that

- $\text{min-ent}(\mathcal{X}) \geq t$
- \mathcal{X} and \mathcal{Y} are (S, ϵ) indistinguishable

Using a standard hybrid argument it is easy to show that the Entropy Smoothing Theorem, as discussed above, can be generalized to probability distributions \mathcal{X} that have (S, ϵ) computational entropy t . In this case, applying a (randomly chosen) universal hash function with output in $\{0, 1\}^{t-2k}$ results in a distribution which is $(S, \epsilon + 2^{-k-1})$ indistinguishable from the uniform one.

4.2 t -DDH: A Relaxed DDH Assumption

We proceed to define the t -DDH assumption. The intuition behind this assumption is that if the Computational Diffie-Hellman Assumption holds in a group G generated by a generator g , then the DH value g^{ab} must have some degree of unpredictability (or “partial hardness”) even when g^a and g^b are given. Specifically, we say that the t -DDH Assumption holds in the group G if the Diffie-Hellman output g^{ab} has t bits of computational entropy (here $0 \leq t \leq \log(G)$). Formally:

Definition 9 We say that the (S, ϵ) t -DDH Assumption holds over a group G if there exists a family of probability distributions $\mathcal{X}(g^a, g^b)$ over G (one distribution for each pair g^a, g^b) such that

- $\text{min-ent}(\mathcal{X}(g^a, g^b)) \geq t$
- The probability distribution \mathcal{DH}_G (see Section 2) is (S, ϵ) indistinguishable from the ensemble

$$\mathcal{R}^* = \{(g^a, g^b, C) \text{ for } a, b \in_R \text{ord}(G) \text{ and } C \in_{\mathcal{X}(g^a, g^b)} G\}$$

It is important to note that the distributions $\mathcal{X}(g^a, g^b)$ in the above definition may be different for each pair of values g^a, g^b . Requiring instead a single distribution \mathcal{X} for all pairs g^a, g^b (as may seem more natural at first glance) results in a significantly stronger, and consequently less useful, assumption.

Consider Example 1 from Section 2: over Z_p^* one can break the DDH by detecting if the quadratic residuosity character of C is consistent with the one induced by g^a, g^b . Yet, Z_p^* can satisfy the t -DDH assumption even for high values of t . For example, if for all a, b for which one of a, b is even we define $\mathcal{X}(g^a, g^b)$ to be the set of quadratic residues in Z_p^* , and for all other pairs g^a, g^b we define $\mathcal{X}(g^a, g^b)$ to be the set of quadratic non-residues in Z_p^* , then the trivial break of DDH in the above example does not hold against these distributions. More generally, if we consider a prime p of the form $2^u q + 1$ where q is a prime then we can get that (given current knowledge) the t -DDH assumption holds for Z_p^* for $t = |p| - u$, while clearly the DDH assumptions does not hold over this group.

Note that the DDH assumption can also be stated in terms of computational entropy. Indeed the DDH assumption over a group G is equivalent to the t -DDH assumption over G for $t = \log(\text{ord}(G))$.

Sampling $\mathcal{X}(g^a, g^b)$. The t -DDH Assumption as stated above makes no requirement on the existence of an efficient sampling algorithm for the distribution $\mathcal{X}(g^a, g^b)$. We say that $\mathcal{X}(g^a, g^b)$ is S' -samplable if there exists a (probabilistic) circuit of size S' whose output distribution (on null input) is $\mathcal{X}(g^a, g^b)$. We say that $\mathcal{X}(g^a, g^b)$ is S' -semi-samplable if there exists a circuit of size S' which is run on input either a or b and whose output distribution is $\mathcal{X}(g^a, g^b)$.

We note that our results do not necessitate of any form of samplability of the \mathcal{X} distributions except for the results on using DDH with short exponents (Section 5). In the latter case our security proof requires $\mathcal{X}(g^a, g^b)$ to be S' -semi-samplable and the parameter S' will affect the quality of the reduction.

As a direct consequence of the Entropy Smoothing Theorem and the definition of t -DDH we have:

Lemma 10 Let G be a group in which the (S, ϵ) t -DDH Assumption holds, and let \mathcal{H} be a collection of universal hash functions such that for all $h \in \mathcal{H}$, $h : G \rightarrow \{0, 1\}^{t'}$ where $t' = t - 2k$. Then the induced distribution of $h(g^{ab})$, for $a, b \in_R [1.. \text{ord}(G)]$ and $h \in_R \mathcal{H}$, is $(S, \epsilon + 2^{-k})$ indistinguishable from the uniform distribution over $\{0, 1\}^{t'}$ even when h, g^a and g^b are given to the distinguisher.

Notice that the above lemma requires the hash function h to be chosen at random for each application. This is the case in several practical protocols (such as the case of IKE [RFC2409, IKEv2], mentioned in the Introduction, in which a key to the hash function is chosen by the communicating parties anew with each run of the protocol). However, it is also possible to fix a randomly chosen hash function and apply it repeatedly to different DH values. An example of such an application would be its use in the context of the Cramer-Shoup CCA-secure cryptosystem

[CS98] (also discussed in the Introduction) in which the specific hash function h would be chosen at random from the family \mathcal{H} by the owner of the decryption key, and published as part of the public key parameters. In this case, the security of the repeated use of the same hash function h can be proved via a standard simulation argument.

Next we show that for groups of prime order, the t -DDH Assumption is equivalent to the full DDH assumption. The proof uses a standard random self-reducibility argument [Sta96, NR97].

Lemma 11 *Let G be a group of prime order q . If the (S, ϵ) t -DDH Assumption holds in G for $t > 0$ then the (S', ϵ') DDH Assumption holds in G with $S' = S - 8\text{exp}_G$ and $\epsilon' = \frac{\epsilon}{1-2^{-t}}$.*

Proof Assume by contradiction that there exists a distinguisher D of size $\leq S'$ that distinguishes between \mathcal{R}_G and \mathcal{DH}_G with probability $> \epsilon'$. We use D to break the (S, ϵ) t -DDH assumption in G .

Let $\mathcal{X}(g^a, g^b)$ be a family of distributions with min-entropy t defined over G . We are given three values $A = g^a, B = g^b, C = g^c$ where either $c = g^{ab}$ or $C \in \mathcal{X}(g^a, g^b)$ G . We sample $r, s, u, v \in_R [1..q]$ and set $A' = A^r g^u = g^{ar+u}, B' = B^s g^v = g^{bs+v}$ and $C' = C^{rs} A^{rv} B^{us} g^{uv}$. Notice that

$$\log_g C' = crs + arv + bsu + uv = (c - ab)rs + (ar + u)(v + bs) \pmod q$$

thus if $C = g^{ab}$ then C' is the result of the DH transform over A', B' . On the other hand, since q is a prime and thus any element has an inverse mod q , if $c \neq ab$ then C' is uniformly distributed over G . Notice that if $C \in \mathcal{X}(g^a, g^b)$ G then $c = ab$ with probability at most 2^{-t} .

Thus by feeding A', B', C' to D we can distinguish the case in which $C = g^{ab}$ and $C \in \mathcal{X}(g^a, g^b)$ G with probability larger than $\epsilon'(1 - 2^{-t})$. Notice that this distinguisher has size $S' + 8\text{exp}_G$ since it costs 8 exponentiations to compute A', B', C' before running D .

Thus by setting $S' = S - 8\text{exp}_G$ and $\epsilon' = \frac{\epsilon}{1-2^{-t}}$ we contradict the assumption that the (S, ϵ) t -DDH Assumption holds in G . ■

This yields an interesting 0-1 law for prime order groups, in which either the DDH Assumption holds, and thus the DH output has $\log(q)$ bits of computational entropy, or we cannot claim that the DH output has *any* bits of computational entropy. We stress that this result, by itself, *does not imply* that over prime order groups either DDH holds or the Diffie-Hellman problem (i.e., Computational Diffie-Hellman) is easy. What the result says is that in this case (i.e. a prime-order group which is CDH but not DDH), pseudorandomness cannot be extracted from a DH value solely based on the computational min-entropy of the distribution but rather may require specialized hard core functions (such as Goldreich-Levin, etc. [Gol01]).

4.3 The Max-Subgroup Theorem

We now proceed to prove our main theorem concerning the t -DDH assumption. The significance of the theorem below is that for a cyclic group G to be t -DDH it suffices that t be the order of the maximal (or maximal disjoint) subgroup of G where the DDH holds.

Theorem 12 *Let G be a cyclic group of order $m = m_1 m_2$ where $(m_1, m_2) = 1$, and G_1 be a subgroup of order m_1 in G . If the (S, ϵ) DDH Assumption holds over G_1 then the (S', ϵ) $\log(m_1)$ -DDH Assumption holds in G , where $S' = S - 5\text{exp}_G$.*

Proof An initial intuition behind the correctness of the theorem is that the hardness hidden in G_1 could be “sampled” when applying a hash function to the DH values over G . This however is incorrect: the size of G_1 may be negligible in relation to $|G|$ and as such the probability to sample a triple (g^a, g^b, g^{ab}) from G_1 is negligible too. The actual argument, presented next, uses the observation that the “hardness” present in G_1 can be extended to its cosets in G .

Let g be a generator of G and $g_1 = g^{m_2}$ be a generator of order m_1 of G_1 . Given $g^a, g^b \in G$, we define the distribution $\mathcal{X}(g^a, g^b)$ to be the uniform distribution over $\{C = g^c \in G \text{ such that } c \in Z_m \text{ and } c \equiv ab \pmod{m_2}\}$. Thus, it is easy to see that $\mathcal{X}(g^a, g^b)$ has $\log(m_1)$ bits of min-entropy (since the above set has m_1 elements). Let \mathcal{R}^* denote the probability distribution $\{(g^a, g^b, C) : a, b \in_R Z_m \text{ and } C \in_{\mathcal{X}(g^a, g^b)} G\}$.

We assume by contradiction that the (S', ϵ) $\log(m_1)$ -DDH assumption does not hold in G , and thus we have a circuit D of size $\leq S'$ which distinguishes between the distributions \mathcal{DH}_G and \mathcal{R}^* with advantage ϵ . Using D we build a distinguisher D_1 of size $\leq S$ that distinguishes between the distributions \mathcal{DH}_{G_1} and \mathcal{R}_{G_1} with the same advantage, thus contradicting the theorem’s assumption.

Given a triple (A_1, B_1, C_1) where $A_1 = g_1^{a_1}, B_1 = g_1^{b_1}$, and C_1 either equals $g_1^{a_1 b_1}$ or $g_1^{c_1}$ for $c_1 \in_R Z_{m_1}$, the distinguisher D_1 does the following:

1. Chooses $i, j \in_R Z_m$
2. Sets $A = A_1 g^i, B = B_1 g^j$ and $C = C_1^{m_2} A_1^j B_1^i g^{ij}$ computed in G
3. Hands D the triple (A, B, C)
4. Outputs the same output bit as D .

Notice that D_1 computes 5 exponentiations and runs D , thus is of size $\leq S$.

Let’s examine the distribution of the triple (A, B, C) . The value A is set to $A = A_1 g^i = g_1^{a_1} g^i = g^{m_2 a_1 + i}$ thus $a = m_2 a_1 + i$. Since $i \in_R Z_m$ then also $a \in_R Z_m$. Similarly for $B = g^b$ we get $b \in_R Z_m$. In the case of C we have $C = C_1^{m_2} A_1^j B_1^i g^{ij} = g^{c_1 m_2 + m_2 a_1 j + m_2 b_1 i + ij}$, thus $c = c_1 m_2 + m_2 a_1 j + m_2 b_1 i + ij$. In addition, we have that $ab = (m_2 a_1 + i)(m_2 b_1 + j) = m_2^2 a_1 b_1 + m_2 a_1 j + m_2 b_1 i + ij$. Thus

$$c - ab = m_2^2 c_1 + m_2 a_1 j + m_2 b_1 i + ij - (m_2^2 a_1 b_1 + m_2 a_1 j + m_2 b_1 i + ij) = m_2^2 c_1 - m_2^2 a_1 b_1$$

which implies $c = m_2^2(c_1 - a_1 b_1) + ab \pmod{m}$. Therefore, if $c_1 = a_1 b_1$ then $c = ab$. On the other hand, if $c_1 \in_R Z_{m_1}$ then $c_1 - a_1 b_1 \in_R Z_{m_1}$, i.e., $c = ab + r m_2^2$ (for $r \in_R Z_{m_1}$). Now, using the fact that m_2 has an inverse modulo m_1 , we get that c is uniformly distributed over the set $\{ab + i m_2 : 0 \leq i < m_1\}$ or, equivalently, that C is distributed according to the distribution $\mathcal{X}(g^a, g^b)$. In other words, the triple (A, B, C) is distributed according to \mathcal{DH}_G if (A_1, B_1, C_1) came from \mathcal{DH}_{G_1} , and it is distributed according to \mathcal{R}^* if (A_1, B_1, C_1) came from \mathcal{R}_{G_1} . Therefore, D_1 distinguishes between \mathcal{DH}_{G_1} and \mathcal{R}_{G_1} with the same probability that D distinguishes between \mathcal{DH}_G and \mathcal{R}^* , that is ϵ .

Since we assumed that the (S, ϵ) DDH holds in G_1 we reached a contradiction. ■

Remark on samplability. The distributions $\mathcal{X}(g^a, g^b)$ defined in the above proof are efficiently samplable given m_1, m_2 and at least one of a, b (i.e., $\mathcal{X}(g^a, g^b)$ is semi-samplable in the terminology of Section 4.2). Indeed given, say, $a, B = g^b$ we can sample $\mathcal{X}(g^a, g^b)$ by choosing $k \in_R Z_{m_1}$ and setting $C = g^{k m_2} B^a$.

From the above theorem we get the following important corollary. Its first part follows immediately from Theorem 12 when using the following terminology: a subgroup H of G is called a **disjoint subgroup** if $(|H|, |G|/|H|) = 1$. The second part of the corollary (which does not involve the notion of disjoint subgroups) follows from Theorem 12 combined with Theorem 2. The corollary is stated without concrete bounds which can be derived from the previous theorems.

Corollary 13 *For any cyclic group G , G is $\log(m)$ -DDH where m is the order of the maximal disjoint DDH subgroup of G . If all the large prime-power subgroups of G are DDH, then G is $\log(m)$ -DDH where m is the order of the maximal DDH subgroup of G .*

The above Corollary is stated somewhat informally, in particular one has to specify the meaning of “large” subgroups. The idea is the following: let G be a cyclic group of order $m = \prod_{i=1}^{\ell} p_i^{e_i}$ where $p_1 < \dots < p_{\ell}$ is the prime decomposition of m . Thus G is the direct product of the subgroups G_i where each G_i has order $p_i^{e_i}$. Fix an (S, ϵ) security parameter and consider the subgroups $\{G_{j_1}, \dots, G_{j_{\ell'}}\}$ which are (S, ϵ) -DDH. Then we can apply Theorem 12 and Lemma 4 since the orders of the subgroups G_i are relatively prime with each other. And thus we have that G is (S', ϵ') m' -DDH where: $m' = \sum_{i=1}^{\ell'} e_{j_i} \log p_{j_i}$, $S' = S - 14 \exp_G$ and $\epsilon' = \ell' \epsilon$.

5 DDH and t -DDH with Short Exponents

In this section we investigate the use of the DDH and t -DDH assumptions in conjunction with the so called “short-exponent discrete-log” assumption.

The Short-Exponent Discrete-Log Assumption. A common practice for increasing the efficiency of exponentiation in cryptographic applications based on the hardness of computing discrete logarithms, and in particular those using the Diffie-Hellman transform, is to replace full-length exponents (i.e., of length logarithmic in the group order) with (significantly) shorter exponents. The security of this practice cannot be justified by the usual assumption that computing discrete logarithms (with full-length exponents) is hard, but rather requires a specific assumption first analyzed in [vOW96] and formalized in [PS98]. We give a concrete security formalization below.

Assumption 14 (s -DLSE [PS98]) *Let G be a cyclic group generated by g and of order $\text{ord}(G) = m$. We say that the (S, ϵ) s -DLSE Assumption holds in G if for every circuit I of size $\leq S$, we have that $\text{Prob}_{x \in_R [1..2^s]}(I(g, m, s, g^x) = x) \leq \epsilon$.*

Current knowledge points to the plausibility of the above assumption even for exponents s significantly shorter than $\log(\text{ord}(g))$. The exact values of s for which the assumption seems to hold depend on the group generated by the element g . An obvious lower bound on s , if one wants to achieve security against 2^k -complexity attacks, is $s \geq 2k$ which is necessary to thwart the usual square-root attacks such as Shanks and Pollard methods. However, as pointed out in [vOW96], there are cases where s needs to be chosen larger than $2k$. Specifically, they show how to use a Pohlig-Hellman decomposition to obtain some of the bits of the exponent. The power of the attack depends on the (relatively) small prime factors of the group order. For example, when working over Z_p^* with a random prime p , the [vOW96] results indicate the use of $s \approx 4k$ (e.g., with a security parameter of 80 one should use $s = 320$ which is much shorter than the 1024 or 2048 bits of p , yet twice as much as the bare minimum of $s = 160$). If one wants to use $s = 2k$ (i.e., assume the

2k-DLSE), it is necessary to work in special groups such as those of prime order or Z_p^* with p a safe prime (i.e., $p = 2q + 1$, and q prime).

From Hardness to Indistinguishability. Gennaro [Gen00] proves that if the s -DLSE assumption holds in $G = Z_p^*$ with p a safe prime then the distribution over G generated by g^x for $x \in_R [1..2^s]$ is computationally indistinguishable from the uniform distribution over G . The following proposition generalizes this result as needed for our purposes².

Proposition 15 *Let G be a cyclic group of order m generated by g , such that m is odd or $m/2$ is odd. If the (S, ϵ) s -DLSE Assumption holds in G , then the following two distributions $\mathcal{S}_G = \{g^x : x \in_R [1..2^s]\}$ and $\mathcal{U}_G = \{g^x : x \in_R Z_m\}$ are (S', ϵ) indistinguishable, where $S' \approx \left(\frac{\epsilon}{|m|-s}\right)^2 S$*

The proof is presented in Appendix C.

Next we show that if in a group G , both the s -DLSE and the t -DDH Assumptions hold, then performing the Diffie-Hellman transform with short exponents a and b , yields a DH output with t bits of computational entropy. In other words, the security of the hashed DH transform over such groups when using s -bit long exponents is essentially equivalent to that of using full exponents.

Theorem 16 *Let G be a cyclic group of order m generated by g , such that m is odd, or $m/2$ is odd. Let s, t be such that the (S_1, ϵ_1) s -DLSE and the (S_2, ϵ_2) t -DDH Assumptions hold in G . Denote with $\mathcal{X}(g^a, g^b)$ the family of distributions induced by the t -DDH assumption over G (see Def. 9). Assume that $\mathcal{X}(g^a, g^b)$ is S_3 -semi-samplable (see Sec. 4.2). Then the following two distributions*

$$\mathcal{SDH} = \{(g^a, g^b, g^{ab}) \text{ for } a, b \in_R [1..2^s]\}$$

and

$$\mathcal{SR}^* = \{(g^a, g^b, C) \text{ for } a, b \in_R [1..2^s] \text{ and } C \in_{\mathcal{X}(g^a, g^b)} G\}$$

are (S, ϵ) indistinguishable where $S = \min(S_2, \left(\frac{\epsilon}{|m|-s}\right)^2 S_1 - S_3)$ and $\epsilon \leq \epsilon_2 + 4\epsilon_1$.

Before proving the Theorem, we point out that for a technical reason inside the proof (an hybrid argument) we need the *semi-samplable* version of the t -DDH assumption here. We stress that the “short exponent” technique is the only case in which we need semi-samplability, and in this case it is easily seen that this condition holds (see Remark at the end of the next Section).

Proof Recall that if the (S_2, ϵ_2) t -DDH Assumption holds over the group G of order m , then there exists a family of probability distributions $\mathcal{X}(g^a, g^b)$ with min-entropy t (one distribution for each pair g^a, g^b) over G such that the distributions

$$\mathcal{DH} = \{(g^a, g^b, g^{ab}) \text{ for } a, b \in_R Z_m\}$$

and

$$\mathcal{R}^* = \{(g^a, g^b, C) \text{ for } a, b \in_R Z_m \text{ and } C \in_{\mathcal{X}(g^a, g^b)} G\}$$

are (S_2, ϵ_2) indistinguishable.

The following standard hybrid argument yields the proof of the theorem. Consider the intermediate distributions

$$\mathcal{D}_0 = \{(g^a, g^b, g^{ab}) \text{ for } a, b \in_R [1..2^s]\}$$

²A similar, but slightly weaker statement was independently stated in [KK04].

$$\begin{aligned}
\mathcal{D}_1 &= \{(g^\alpha, g^b, g^{\alpha b}) \text{ for } \alpha \in_R Z_m, b \in_R [1..2^s]\} \\
\mathcal{D}_2 &= \{(g^\alpha, g^\beta, g^{\alpha\beta}) \text{ for } \alpha, \beta \in_R Z_m\} \\
\mathcal{D}_3 &= \{(g^\alpha, g^\beta, C) \text{ for } \alpha, \beta \in_R Z_m \text{ and } C \in_{\mathcal{X}(g^\alpha, g^\beta)} G\} \\
\mathcal{D}_4 &= \{(g^\alpha, g^b, C) \text{ } b \in_R [1..2^s], \alpha \in_R Z_m \text{ and } C \in_{\mathcal{X}(g^\alpha, g^b)} G\} \\
\mathcal{D}_5 &= \{(g^a, g^b, C) : a, b \in_R [1..2^s] \text{ and } C \in_{\mathcal{X}(g^a, g^b)} G\}
\end{aligned}$$

Clearly $\mathcal{D}_0 = \mathcal{SDH}$ while $\mathcal{D}_5 = \mathcal{SR}^*$.

Under the (S_2, ϵ_2) t -DDH Assumption we know that \mathcal{D}_2 is (S_2, ϵ_2) indistinguishable from \mathcal{D}_3 .

Also, under the (S_1, ϵ_1) s -DLSE Assumption we know that \mathcal{D}_i is $(S_1/s - S_3, (|m| - s)\epsilon_1)$ indistinguishable from \mathcal{D}_{i+1} for $i = 0, 1, 3, 4$ by reduction to Proposition 15. The extra additive factor of S_3 is due to the fact that in the case $i = 3, 4$ one needs $\mathcal{X}(g^a, g^b)$ to be semi-samplable, which by assumption can be done by a circuit of size S_3 .

Thus by invoking the triangle inequality for computational indistinguishability (see Prop 22 in Appendix A) we have that \mathcal{SDH} is (S, ϵ) indistinguishable from \mathcal{SR}^* where $S = \min(S_2, S_1/s - S_3)$ and $\epsilon = \epsilon_2 + 4(|m| - s)\epsilon_1$ as desired. ■

Note that, as a particular case, when $t = \log(m)$ the theorem states that if G is a DDH group in which the s -DLSE assumption holds, then performing the DH transform over G with exponents of size s yields values that are indistinguishable from random elements in G .

6 Hashed DH over Z_p^* and its Subgroups

Here we discuss the security of the hashed DH transform over groups and subgroups of Z_p^* for random prime p . Throughout this section we assume that the DDH assumption holds over the large prime-order subgroups of Z_p^* (or the prime-power order subgroups in the unusual case that $p - 1$ is divisible by a large prime with multiplicity larger than 1). Under this assumption we immediately get that it is secure to use the hashed DH transform over a subgroup G_q of Z_p^* of order q , provided that q is a sufficiently large prime that divides $p - 1$. The meaning of “large” here is that DDH holds over G_q with parameters (S, ϵ) that make the distinguishing task infeasible; specifically, when talking of a “security parameter” k we require $S/\epsilon \geq 2^k$. Also, a large q is one for which a sufficient number of bits can be extracted from a Diffie-Hellman value. For example, if the application requires a pseudorandom output of ℓ bits then q needs to satisfy $|q| \geq \ell + 2k$ (see Theorem 7).

Very importantly, however, due to our results we can extract from a Diffie-Hellman value over Z_p^* more bits than those guaranteed by individual factors q of $p - 1$. If we want to extract ℓ bits and Z_p^* has a subgroup of order m , where m is the product of different large primes (say, each of size $\geq 2k$), then it suffices that $|m| \geq \ell + 2k$ in order to extract ℓ bits from a DH value over such subgroup. Moreover, these results show that one can securely apply the hashed DH transform also over some non-DDH groups whose order is divisible by small prime factors which, in particular, is the case of Z_p^* (the order $m = p - 1$ of this group is always divisible by small prime factors, e.g., 2). Specifically, we showed that the hashed DH is secure over Z_p^* provided that $p - 1$ has enough prime divisors (with multiplicity 1) whose product is larger than the entropy bound $2^{\ell+2k}$, and for which the subgroups of corresponding prime order are DDH in the above sense. (In particular, the fact that $p - 1$ has additional smaller prime factors does not invalidate the security of the hashed DDH in Z_p^* .)

A particularly interesting group is Z_p^* for $p = 2q + 1$ and q prime. In this case, working directly with the hashed DH over Z_p^* is secure since we are assuming that its subgroup of order q is DDH, and therefore the whole Z_p^* group is $\left\lfloor \frac{p-1}{2} \right\rfloor$ -DDH. Working over Z_p^* in this case has several important advantages: (i) one can produce a large (actually, largest) number of pseudorandom bits (specifically, $|p| - 1 - 2k$ bits); (ii) p can be chosen such that 2 is a generator of Z_p^* (which speeds up exponentiation); (iii) the $2k$ -DLSE Assumption (see Section 5) is conjectured to hold in these groups [vOW96] and therefore one can use minimal-length exponents (i.e., of length $2k$) in these groups, obtaining yet another significant exponentiation speedup without sacrificing the security of the (hashed) DH transform; and (iv) these groups are free from the potentially serious attacks described in [LL97] (that affect subgroups of prime order q where $(p-1)/q$ has a relatively large smooth factor). Note that items (i) and (iii) use our results in an essential way. The only downside of working over such a group is the cost of generating p 's of the above form; this, however is insignificant in typical applications (e.g., IKE [RFC2409, IKEv2]) in which prime generation is very rare, and usually done at the set-up of the system and used for a large period of time.

Note that in all of the above examples it is assumed that one knows the full or partial factorization of $p - 1$; in particular, the knowledge of this factorization is essential for selecting a generator of the group. It is a theoretically and practically important question to establish whether the knowledge of the factorization of $p - 1$ is essential for working securely over Z_p^* or over one of its subgroups. In the rest of this section we show that this knowledge is not essential (at least under some plausible assumptions on the distribution of the prime factors of $p - 1$). Specifically, it follows from our results that if one chooses a random prime p (of a pre-specified size such that the Discrete Logarithm Problem is hard in Z_p^*) and a random element e in Z_p^* , then performing the hashed DH transform over the group generated by e is secure.³

Let p be a random prime such that $p - 1 = p_1 p_2 \cdots p_n$ and $p_1 \leq p_2 \leq \dots \leq p_n$ are all (not necessarily different and possibly unknown) primes. Let e be an element randomly chosen from Z_p^* , and let G_e denote the subgroup of Z_p^* generated by e . We first claim that with overwhelming probability the large prime factors of $p - 1$ divide the order of G_e .

Lemma 17 *Let Z_p^* and $p - 1 = p_1 \cdots p_n$ be as described above. Then for all $1 \leq i \leq n$: $Pr_{e \in_R Z_p^*} [p_i \nmid \text{ord}(e)] \leq 1/p_i$.*

Proof Let g be a generator of Z_p^* . There are at most $(p-1)/p_i$ elements whose order is not divisible by p_i , and they are the elements of the form $g^{j p_i}$ for $1 \leq j \leq (p-1)/p_i$. When $p_i^2 | p-1$ this is a strict upper bound, otherwise this is an exact bound. Thus, the probability to choose e such that $p_i \nmid \text{ord}(e)$ is at most $\frac{(p-1)/p_i}{p-1} = \frac{1}{p_i}$. ■

Corollary 18 *For a given bound B , let $p - 1 = \prod_{i=1}^n p_i$ where $p_j, p_{j+1}, \dots, p_n > B$. Then*

$$Pr_{e \in_R Z_p^*} [\prod_{i=j}^n p_i \mid \text{ord}(e)] \geq 1 - \sum_{i=j}^n \frac{1}{p_i} \geq 1 - \frac{n-j}{B} \geq 1 - \frac{\log p}{B}.$$

Thus, for large values of B , the order of a random element e is divisible, with overwhelming probability, by all the prime factors of $p - 1$ which are larger than B . Or, equivalently, G_e has as subgroups all the prime-order subgroups of Z_p^* whose order is larger than B .

³We stress that while the legitimate users of such a scheme do not need to know the factorization of $p - 1$, the scheme remains secure even if this factorization is known to the attacker.

Now, if we set our security parameter to k , define $B = 2^{2k}$, and assume that the DDH holds in subgroups of prime order larger than B , then we have that, with overwhelming probability, G_e contains all the prime order DDH subgroups of Z_p^* . In other words, if we denote by P the product of all prime factors of $p - 1$ larger than B , we have that G_e contains, by virtue of our DDH Characterization Theorem (Theorem 2), a DDH subgroup of size P , and then by the Max-Subgroup Theorem (Theorem 12) we get that G_e is $|P|$ -DDH.

All that is left to argue is that $|P|$ is large enough. For this we use the following lemma from [vOW96] that provides an upper bound on the expected size of the product of all prime divisors of $p - 1$ that are smaller than B (and thus, it provides a lower bound on the expected size of $|P|$).

Lemma 19 ([vOW96]) *For a random prime p (as above) and a fixed bound B , the expected length of $\Pi_i p_i$ where $p_i < B$ is $\log B + 1$.*

In other words, the lemma states that the expected size of $|P|$ is $|p| - |B| = |p| - 2k$.

If, for the sake of illustration, we set $|p| = 1024$ and $k = 80$ we get that we expect G_e to be 864-DDH. However, note that this expected size may vary for specific p 's, and in particular the above result does not rule out that there could be many primes p 's for which $p - 1$ is smooth. Fortunately this is not the case: a better estimate of the probability that for a random prime p , the value $p - 1$ is smooth can be found in [PS02] from which one can state that most primes p have a large prime q dividing $p - 1$. We refer the reader to [PS02] for details.

Remark (Short exponents and semi-samplability). Notice that in order to use short exponents in the above scenario (i.e., when working over a random prime p with a random generator e), one must make sure that the order m of the group generated by e is either odd, or $m/2$ is odd (so that we can invoke Theorem 16). This can be easily achieved by choosing first a random element e in Z_p^* and then using as the group generator the element $e^{2^f} \bmod p$ where f is the maximal integer such that $2^f | (p - 1)$. In addition, for the application of Theorem 16, we need to show that the distributions $\mathcal{X}(g^a, g^b)$ in this case are semi-samplable. This is so since in the above arguments we are (implicitly) using the distributions defined in the proof of Theorem 12 which are semi-samplable when the factorization of the group order is known (see the remark following the proof of Theorem 12). Therefore, we obtain that, even though the honest parties may not know the factorization of $p - 1$, the DH transform with short exponents remains secure in this case *even if* such factorization is available to the attacker.

Acknowledgments

We thank Don Coppersmith for his elegant (counter) example presented in Appendix B, and the anonymous referees from Eurocrypt'2004 for their useful comments.

We also thank Alexey Kirichenko for some insightful comments on an earlier version of this paper, and Andrey Sidorenko for pointing out a better concrete security reduction for the results in [Gen00]. Finally we thank Carl Pomerance for discussions and pointers relative to Section 6.

A Indistinguishability of Probability Distributions

Definition 20 Let \mathcal{X}, \mathcal{Y} be two probability distributions over a set A . We say that \mathcal{X} and \mathcal{Y} have statistical distance bounded by Δ if

$$\sum_{x \in A} |\text{Prob}_{\mathcal{X}}[x] - \text{Prob}_{\mathcal{Y}}[x]| \leq \Delta$$

Next we adapt the classical notion of computational indistinguishability [GM84] to the concrete security setting (informally, two distributions \mathcal{X} and \mathcal{Y} are (S, ϵ) indistinguishable if no circuit of size S can distinguish between samples drawn according to \mathcal{X} or according to \mathcal{Y} with advantage larger than ϵ).

Definition 21 Let \mathcal{X}, \mathcal{Y} be two probability distributions over A . Given a circuit D (called the distinguisher) consider the following quantities

$$\delta_{D, \mathcal{X}} = \text{Prob}_{x \in \mathcal{X}}[D(x) = 1] \quad \text{and} \quad \delta_{D, \mathcal{Y}} = \text{Prob}_{y \in \mathcal{Y}}[D(y) = 1]$$

We say that the probability distributions \mathcal{X} and \mathcal{Y} are (S, ϵ) indistinguishable if for every circuit D of size $\leq S$ we have that

$$|\delta_{D, \mathcal{X}} - \delta_{D, \mathcal{Y}}| \leq \epsilon$$

We now state a simple “triangle inequality” for (S, ϵ) indistinguishability (a.k.a. the “hybrid argument”).

Proposition 22 Given three probability distributions $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ over a set A , such that (i) \mathcal{X} is (S_1, ϵ_1) indistinguishable from \mathcal{Y} and (ii) \mathcal{Y} is (S_2, ϵ_2) indistinguishable from \mathcal{Z} . Then \mathcal{X} is (S, ϵ) indistinguishable from \mathcal{Z} where $S = \min(S_1, S_2)$ and $\epsilon = \epsilon_1 + \epsilon_2$.

Proof Assume that \mathcal{X} is not (S, ϵ) indistinguishable from \mathcal{Z} . Then there exists a distinguisher D of size S such that

$$|\delta_{D, \mathcal{X}} - \delta_{D, \mathcal{Z}}| > \epsilon$$

Now by the triangle inequality we have that

$$\epsilon < |\delta_{D, \mathcal{X}} - \delta_{D, \mathcal{Y}}| + |\delta_{D, \mathcal{Y}} - \delta_{D, \mathcal{Z}}| \leq \epsilon_1 + \epsilon_2 = \epsilon$$

which is a contradiction. Note that the second upper bound is due to the fact that the size of D is smaller than both S_1 and S_2 . ■

B Coppersmith’s Example

As mentioned at the end of Section 3, Coppersmith [Cop04] has provided us with an example of a non-DDH cyclic group G of order q^2 (for prime q) that contains a DDH subgroup G_q of order q . Moreover, such a group G can be constructed on the basis of any given DDH group of order q . Here we present Coppersmith’s construction.

Let G_q be a cyclic DDH group of order q , for prime q , generated by an element g . We build a group G as follows. The set of elements in G is $S = \{(h, a) : h \in G_q, 0 \leq a < q\}$ and the group

operation $*$ is defined as: $(h_1, a_1) * (h_2, a_2) = (h, a)$ where (i) if $a_1 + a_2 < q$ then $h = h_1 h_2$ (with multiplication over G_q) and $a = a_1 + a_2$; and (ii) if $a_1 + a_2 \geq q$ then $h = h_1 h_2 g$ and $a = a_1 + a_2 - q$. The idea behind the construction of the group G , and its operation, is given by the following natural bijection between the set of integers between 0 and $q^2 - 1$ and the set S : for any $0 \leq b, c < q$, we map $bq + c$ into (g^b, c) . More specifically, we consider G as a cyclic group with generator $(1, 1)$ (the first 1 is the unit element in G_q , the second is the integer 1). In this case we have that for any $0 \leq b, c < q$, $(1, 1)^{bq+c} = (g^b, c)$, (or, equivalently, $\text{dlog}_{(1,1)}(g^b, c) = bq + c$).

Clearly, the element $(1, 1)^q = (g, 0)$ generates the subgroup $G_q \times \{0\}$ of order q which is (by assumption) DDH. However G is not DDH (not even CDH). Indeed, the Diffie-Hellman transform over G is (see footnote⁴) $DH((h_1, a_1), (h_2, a_2)) = (h_1^{a_2} h_2^{a_1} g^{\lfloor a_1 a_2 / q \rfloor}, a_1 a_2 \bmod q)$, and then trivial to compute given (h_1, a_1) and (h_2, a_2) . Note that in this example G is not even CDH. Yet, a similar, but somewhat more involved, example shows that one can build G of prime-power order ($q^e, e > 1$) with the following properties (i) CDH holds in G , (ii) DDH holds in a subgroup of G ; yet (iii) DDH does not hold in G .

C Proof of Proposition 15

In this section we prove the following proposition from Section 5.

Proposition 15 *Let G be a cyclic group of order m generated by g , such that m is odd or $m/2$ is odd. If the (S, ϵ) s -DLSE Assumption holds in G , then the following two distributions $\mathcal{S}_G = \{g^x : x \in_R [1..2^s]\}$ and $\mathcal{U}_G = \{g^x : x \in_R Z_m\}$ are (S', ϵ) indistinguishable, where $S' = \left(\frac{\epsilon}{|m|-s}\right)^2 S$.*

What follows is an extension of arguments that appeared first in [PS98, Gen00].

Let m be the order of cyclic group G and g a generator for G . Let $n = |m|$.

[Hugo: Rosario: check that indeed n is used below in lieu of $|m|$.]

H

Hard-Core Bits and the s -DLSE Assumption. In [PS98] Patel and Sundaram prove that under the s -DLSE Assumption the bits x_2, x_3, \dots, x_{n-s} are simultaneously hard for the function $f(x) = g^x \bmod p$, if p is congruent to 3 mod 4. It is not difficult to see that their proof can be extended in two ways:

- It holds for any cyclic group of odd order, in which case even the bit x_1 is hard.
- It holds for any cyclic group G of even order m but such that $m/2$ is odd. Notice that for these groups, computing x_1 when given $y = g^x$ is easy.

Short-Exponent Indistinguishability. Gennaro in [Gen00] builds on the above result from [PS98] as follows. An alternative way to say that the the bits x_i, \dots, x_j are simultaneously hard is to say that the two distributions:

$$[g^x, x_i, \dots, x_j] \text{ for } x \in_R Z_m$$

$$[g^x, r_i, \dots, r_j] \text{ for } x \in_R Z_m, r_i, \dots, r_j \in \{0, 1\}$$

⁴Let $h_1 = g^{b_1}, h_2 = g^{b_2}$. Then: $DH((h_1, a_1), (h_2, a_2)) = DH((g^{b_1}, a_1), (g^{b_2}, a_2)) = DH((1, 1)^{b_1 q + a_1}, (1, 1)^{b_2 q + a_2}) \stackrel{\text{def}}{=} (1, 1)^{(b_1 q + a_1)(b_2 q + a_2)} = (1, 1)^{(b_1 a_2 + a_1 b_2)q + a_1 a_2} = (g^{b_1 a_2 + a_1 b_2 + \lfloor a_1 a_2 / q \rfloor}, a_1 a_2 \bmod q) = (h_1^{a_2} h_2^{a_1} g^{\lfloor a_1 a_2 / q \rfloor}, a_1 a_2 \bmod q)$.

are computationally indistinguishable. Denote with $x(i, j)$ the value x with the bits in position from i to j zeroed out. Then a consequence of the above statement is that the two distributions

$$[g^x] \text{ and } [g^{x(i,j)}] \text{ for } x \in_R Z_m$$

are computationally indistinguishable.

Gennaro uses this to construct efficient pseudo-random generators in which the basic operation is an exponentiation with an exponent with a lot of contiguous zero's in it (the positions from i to j indeed) which is substantially faster to compute than a regular exponentiation.

Notice, however, that the above conclusion is still different from the statement of Proposition 15. But we show now that if Proposition 15 is false then we can contradict the above conclusion. We distinguish two cases.

Case 1: m is odd. In this case we have that [Gen00] implies that

$$[g^x] \text{ and } [g^{x(1,n-s)}] \text{ for } x \in_R Z_m$$

are computationally indistinguishable. Assume that we have a distinguisher D that distinguishes between $[g^x]_{x \in_R Z_m}$ and $[g^z]_{z \in_R [1..2^s]}$ then we can use D to distinguish in the case above. Given an element y we compute $y^{2^{-(n-s)} \bmod m} g^w$ with $w \in_R [1..2^s]$. A random group element y will be mapped to a random group element, while an element of the form $y = g^{x(1,n-s)}$ (i.e., with the least $n - s$ significant bits zeroed out) will be mapped to a random element of the form g^z with $z < 2^s$.

Case 2: m is even, but $m/2$ is odd. In this case we have that [Gen00] implies that

$$[g^x] \text{ and } [g^{x(2,n-s)}] \text{ for } x \in_R Z_m$$

are computationally indistinguishable. Notice also that given $y = g^x$, the bit x_1 is easily computable. Assume that we have a distinguisher D that distinguishes between $[g^x]_{x \in_R Z_m}$ and $[g^z]_{z \in_R [1..2^s]}$ then we can use D to distinguish in the case above. Given an element y we perform the following steps:

- Compute x_1 and set $y_1 = y \cdot g^{-x_1}$
- For $i = 2$ to $n - s$, compute y_i as the principal square root of y_{i-1} . The principal square root of a square y is that square root which is also a square. When $m/2$ is odd, the principal square root is unique and can be efficiently computed.
- Set $y' = y_{n-s} \cdot g^w$ with $w \in_R [1..2^s]$.

A random group element y will be mapped to a random group element, while an element of the form $y = g^{x(2,n-s)}$ will be mapped to a random element of the form g^z with $z < 2^s$.

CONCRETE COMPLEXITY. The concrete complexity bounds stated in Prop. 15 are a refinement of the ones stated in [Gen00].

References

- [ABR01] M. Abdalla, M. Bellare, and P. Rogaway. DHIES: An Encryption Scheme Based on the Diffie-hellman Problem. In *CT-RSA '01*, pages 143–158, 2001. LNCS No. 2020.

- [BJN00] D. Boneh, A. Joux, and P. Nguyen. Why Textbook ElGamal and RSA Encryption are Insecure . In *AsiaCrypt '00*, pages 30–44, 2000. LNCS No. 1976.
- [Bon98] D. Boneh. The Decision Diffie-Hellman Problem. In *Third Algorithmic Number Theory Symposium*, pages 48–63, 1998. LNCS No. 1423.
- [Bra93] S. Brands. An Efficient Off-Line Electronic Cash System Based on the Representation Problem. TR CS-R9323, CWI, Holland, 1993.
- [CW79] L. Carter and M. N. Wegman. Universal Classes of Hash Functions. *JCSS*, 18(2):143–154, April 1979.
- [Cop04] D. Coppersmith. Personal communication. March 2004.
- [CS98] R. Cramer and V. Shoup. A Practical Public Key Cryptosystem Provable Secure Against Adaptive Chosen Ciphertext Attack. In *Crypto '98*, pages 13–25, 1998. LNCS No. 1462.
- [DH76] W. Diffie and M. E. Hellman. New Directions in Cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- [DGHKR04] Y. Dodis, R. Gennaro, J. Hastad, H. Krawczyk and T. Rabin. Randomness Extraction and Key Derivation Using the CBC, Cascade and HMAC modes. In *Crypto'04*, pages 494–510, 2004. LNCS No. 3152,
- [ElG85] T. ElGamal. A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms. *IEEE Trans. Info. Theory*, IT 31:469–472, 1985.
- [Gen00] R. Gennaro. An Improved Pseudo Random Generator Based on Discrete Log. *J.Cryptology*, 18(2):91–110, Springer, April 2005. Preliminary version in *Crypto '00*, pages 469–481, 2000. LNCS No. 1880.
- [Gol01] Oded Goldreich. *Foundations of Cryptography*. Cambridge University Press, 2001.
- [GM84] S. Goldwasser and S. Micali. Probabilistic Encryption. *JCSS*, 28(2):270–299, April 1984.
- [HILL99] J. Hastad, R. Impagliazzo, L. Levin, and M. Luby. Construction of a Pseudo-random Generator from any One-way Function. *SIAM. J. Computing*, 28(4):1364–1396, 1999.
- [IKEv2] IKEv2. Internet Key Exchange (IKEv2) Protocol, draft-ietf-ipsec-ikev2-17.txt, (to be published as an RFC). Editor: C. Kaufman, September 2004.
- [KK04] T. Koshihara and K. Kurosawa. Short Exponent Diffie-Hellman Problems. *Public Key Cryptography 2004*, Springer LNCS 2947, pp.173-186
- [Kra03] H. Krawczyk. SIGMA: The ‘SiGn-and-MAc’ Approach to Authenticated Diffie-Hellman and Its Use in the IKE Protocols. In *Crypto '03*, pages 400–425, 2003. LNCS No. 2729. <http://www.ee.technion.ac.il/~hugo/sigma.html>
- [LL97] C.H. Lim and P.J. Lee. A Key Recovery Attack on Discrete Log-Based Schemes Using a Prime Order Subgroup. In *Crypto '97*, pages 249–263, 1997. LNCS No. 1294.

- [Lub96] M. Luby. *Pseudorandomness and Cryptographic Applications*. Princeton Computer Science Note, Princeton University Press, January 1996.
- [MW96] U. Maurer and S. Wolf. Diffie-Hellman Oracles. In *Crypto '96*, pages 268–282, 1996. LNCS No. 1109.
- [NR97] M. Naor and O. Reingold. Number-theoretic Constructions of Efficient Pseudorandom Functions. In *Proc. 38th FOCS*, pages 458–467. IEEE, 1997.
- [PS98] S. Patel and G. Sundaram. An Efficient Discrete Log Pseudo Random Generator. In *Crypto '98*, pages 304–317, 1998. LNCS No. 1462.
- [PS02] C. Pomerance and I.E. Shparlinski. Smooth Orders and Cryptographic Applications. In *ANTS 2002*, pages 338–348, 2002. LNCS No. 2369.
- [RFC2409] RFC2409. The Internet Key Exchange (IKE). Authors: D. Harkins and D. Carrel, Nov 1998.
- [Sta96] M. Stadler. Publicly Verifiable Secret sharing. In *Eurocrypt '96*, pages 190–199, 1996. LNCS No. 1070.
- [vOW96] P.C. van Oorschot and M. Wiener. On Diffie-Hellman Key Agreement with Short Exponents. In *Eurocrypt '96*, pages 332–343, 1996. LNCS No. 1070.