

New Results on the Hardness of Diffie-Hellman Bits

María Isabel González Vasco¹, Mats Näslund², and Igor E. Shparlinski³

¹ Dept. of Mathematics, University of Oviedo
Oviedo 33007, Spain

`mvasco@orion.ciencias.uniovi.es`

² Ericsson Research

SE-16480 Stockholm, Sweden

`mats.naslund@ericsson.com`

³ Dept. of Computing, Macquarie University

Sydney, NSW 2109, Australia

`igor@ics.mq.edu.au`

Abstract. We generalize and extend results obtained by Boneh and Venkatesan in 1996 and by González Vasco and Shparlinski in 2000 on the hardness of computing bits of the Diffie-Hellman key, given the public values. Specifically, while these results could only exclude (essentially) error-free predictions, we here exclude any non-negligible advantage, though for larger fractions of the bits. We can also demonstrate a trade-off between the tolerated error rate and the number of unpredictable bits.

Moreover, by changing computational model, we show that even a very small proportion of the most significant bits of the Diffie-Hellman secret key cannot be retrieved from the public information by means of a Las Vegas type algorithm, unless the corresponding scheme is weak itself.

1 Introduction

So called “provable” security models, in which the robustness of a cryptographic tool can be justified by means of a formal proof, are gaining more and more attention. In such models, the security of a scheme or protocol is measured in terms of the chances (non-negligible advantage over a random guess) a malicious adversary has of retrieving information he is not supposed to have access to. There are already several proposals for schemes that are robust in this sense, some of them merely theoretical but others already deployed in practice. Much research has been devoted to this topic, and there is indeed a large battery of results for various schemes in different computational models (e.g. [2, 3, 9, 10]).

Since the early days of cryptography, one security property that has been extensively studied is the security with respect to “approximate cracking”. Robustness in this sense is stated by proving that single bits in an encrypted message are no easier to obtain than the whole message itself (or other information

close to the secret in some metric). General frameworks for such studies are for example the *hard-core bit problem*, first formalized in [4], and the *hidden number problem*, introduced by Boneh and Venkatesan [6, 7].

In addition, any security property can be studied in different computational models, ranging from the classical Turing machine model, passive/active adversaries, restricted algebraic models, up to the more recent quantum and side-channel attack models, the latter two being more “physical” in nature. Indeed, models that might seem unrealistic today could become a reality in 20 years, a time-span which may be required to cryptographically guard secrets in many applications. We therefore believe it is important to keep an open mind to various models and investigate which implications they have.

In this paper we extend the area of application of algorithms for the *hidden number problem*, deriving new bit security properties of the Diffie–Hellman key exchange scheme. Detailed surveys of bit security results for various cryptographic schemes are given in [14]; several more recent results can be found in [5–7, 15–17, 20, 22, 26, 27, 32, 34, 35].

We show that making some adjustments to the scheme proposed in [6] and refined in [16], one can obtain bit security results for the Diffie–Hellman secret key of the same strength as in [6, 16], but in a much more restricted computational model of unreliable oracles, which represent adversaries that only retrieve correct guesses for the target bits with a certain probability. We perform the study with two types of unreliable oracles, roughly corresponding to the classical “Monte Carlo” type algorithms, as well as the in cryptography less conventional “Las Vegas” type of algorithm. In fact, we also obtain an improvement of the result of [16] for “error-free” oracles, as we use the recent bound of exponential sums over small subgroups from [8] instead of the bound from [21] that lead to the result in [16]. Also, our Lemma 3 is based on a recent improvement [1] in lattice reduction algorithms, whereas in [16] older results were applied.

2 Notation

As usual we assume that for a prime p the field \mathbb{F}_p of p elements is represented by the set $\{0, 1, \dots, p - 1\}$. Accordingly, sometimes, where obvious, we treat elements of \mathbb{F}_p as integer numbers in the above range. Also, for an integer s we denote by $[s]_p$ the remainder of s on division by p .

For a real $\eta > 0$ and $t \in \mathbb{F}_p$ we denote by $\text{MSB}_{\eta,p}(t)$ any integer which satisfies the inequalities

$$\frac{p}{2^\eta}(\text{MSB}_{\eta,p}(t) - 1) \leq t < \frac{p}{2^\eta}(\text{MSB}_{\eta,p}(t)). \quad (1)$$

Thus, roughly speaking, $\text{MSB}_{\eta,p}(t)$ is the integer defined by the η most significant bits of t . However, this definition is more flexible and better suited to our purposes. In particular note that η in the inequality (1) need not be an integer.

Throughout the paper $\log x$ denotes the binary logarithm of $x \geq 1$. The implied constants in the symbol “ O ” may occasionally, where obvious, depend on a real parameter $\varepsilon > 0$ and are absolute otherwise.

We denote by $\mathbb{E}[\xi]$ the expected value of a random variable ξ . Accordingly, $\mathbb{E}_\xi[g(\xi)]$ denotes the expected value of a random variable $g(\xi)$, which, for a given function g , only depends on the distribution of ξ . We make use of the following variant of the Markov inequality: for positive c and a random variable ξ upper bounded by M ,

$$\Pr[\xi \geq \mathbb{E}_\xi[\xi]/c] \geq M^{-1}(1 - 1/c)\mathbb{E}_\xi[\xi]. \quad (2)$$

3 Preparations

Reconstructing g^{ab} from “noisy” approximations of g^{ab} can be formulated as a *hidden number problem*, [6, 7]. We review important ingredients for this problem. In particular, we collect several useful results about the hidden number problem, lattices and exponential sums and establish some links between these techniques.

3.1 Hidden Number Problem and Uniform Distribution mod p

One of many possible variations of the hidden number problem is:

Given a finite sequence \mathcal{T} of elements of \mathbb{F}_p^ , recover $\alpha \in \mathbb{F}_p^*$ for which for polynomially many known random $t \in \mathcal{T}$ we are given $\text{MSB}_{\eta,p}(\alpha t)$ for some $\eta > 0$.*

The case of $\mathcal{T} = \mathbb{F}_p^*$ is exactly the one considered in [6, 7]. However, it has been noticed for the first time in [16], and exploited in a series of works, see [11, 17, 26, 28, 29], that in fact for cryptographic applications one has to consider more general sequences \mathcal{T} . An important issue is the uniformity of distribution of these sequences.

For a sequence of N points $0 \leq \vartheta_1, \dots, \vartheta_N < 1$ define its *discrepancy* D by

$$D = \sup_{0 \leq \gamma < 1} \left| \frac{T(\gamma)}{N} - \gamma \right|,$$

where $T(\gamma)$ is the number of points of this sequence in the interval $[0, \gamma]$.

We say that a finite sequence \mathcal{T} of integers is Δ -homogeneously distributed modulo a prime p if for any integer a with $\gcd(a, p) = 1$, the discrepancy $\mathcal{D}_a(\mathcal{T})$ of the sequence of fractional parts $\{at/p\}$, $t \in \mathcal{T}$, satisfies $\mathcal{D}_a(\mathcal{T}) \leq \Delta$. It has been shown in Lemma 4 of [28] that the algorithm of [6] can be modified to work for sequences that are Δ -homogeneously distributed modulo a prime p , provided Δ is small enough.

3.2 Lattices

As in the pioneering papers [6, 7], our results rely on rounding techniques in lattices. We briefly review a few results and definitions. For general references on lattice theory and its important cryptographic applications, we refer to [18] and also to the recent surveys [30, 31].

A basic lattice problem is the *closest vector problem* (CVP): given a basis of a lattice L in \mathbb{R}^s and a target $\mathbf{u} \in \mathbb{R}^s$, find a lattice vector $\mathbf{v} \in L$ which minimizes the Euclidean norm $\|\mathbf{u} - \mathbf{v}\|$ among all lattice vectors. A modification where $\mathbf{u} = 0$ is a zero vector (thus $\mathbf{u} \in L$) is the *shortest vector problem* (SVP): find a nonzero $\mathbf{v} \in L$ of smallest Euclidean norm $\|\mathbf{v}\|$ among all lattice vectors.

Here, as in [28], we use the best CVP approximation polynomial-time result known, which follows from the recent shortest vector algorithm of [1] combined with the reduction of [23] from approximating the CVP to approximating the SVP, which leads to the following statement:

Lemma 1. *For any constant $\gamma > 0$, there exists a randomized polynomial time algorithm which, given a lattice L and a vector $\mathbf{r} \in \mathbb{Q}^s$, finds a lattice vector \mathbf{v} satisfying with probability exponentially close to 1 the inequality*

$$\|\mathbf{v} - \mathbf{r}\| \leq 2^{\gamma s \log \log s / \log s} \min \{\|\mathbf{z} - \mathbf{r}\|, \mathbf{z} \in L\}.$$

For integers t_1, \dots, t_d selected in the interval $[0, p-1]$, we denote by $L(t_1, \dots, t_d)$ the full rank $d + 1$ -dimensional lattice generated by the rows of the following $(d + 1) \times (d + 1)$ -matrix

$$\begin{pmatrix} p & 0 & \dots & 0 & 0 \\ 0 & p & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & p & 0 \\ t_1 & t_2 & \dots & t_d & 1/p \end{pmatrix}. \quad (3)$$

Our principal tool is an extension of Lemma 4 of [28], which in turn extends the algorithm of [6]. The results below are analogues of Lemmas 6.2 and 6.3 of [35]. For applications to Diffie-Hellman we deal with sequences corresponding to small finite subgroups of \mathbb{F}_p^* which satisfy the above requirement of Δ -homogeneous distribution, so Lemma 4 of [28] can be applied directly to them.

Lemma 2. *Assume that a real μ and an integer d satisfy*

$$d(\mu - \log 5) \geq 2 \log p$$

and let α be a fixed integer in the interval $[0, p - 1]$. Assume that t_1, \dots, t_d are chosen uniformly and independently at random from a finite $2^{-\mu}$ -homogeneously distributed integer sequence \mathcal{T} modulo p . Then with probability $P \geq 1 - 1/p$ for

any vector $\mathbf{s} = (s_1, \dots, s_d, 0)$ with

$$\left(\sum_{i=1}^d \left(\lfloor \alpha t_i \rfloor_p - s_i \right)^2 \right)^{1/2} \leq p2^{-\mu},$$

all vectors $\mathbf{v} = (v_1, \dots, v_d, v_{d+1}) \in L(t_1, \dots, t_d)$ satisfying

$$\left(\sum_{i=1}^d (v_i - s_i)^2 \right)^{1/2} \leq p2^{-\mu},$$

are such that

$$v_i \equiv \beta t_i \pmod{p}, \quad i = 1, \dots, d, \quad v_{d+1} = \beta/p$$

with some $\beta \equiv \alpha \pmod{p}$.

Proof. We define the *modular norm* of an integer γ modulo p as

$$\|\gamma\|_p = \min_{b \in \mathbb{Z}} |\gamma - bp|.$$

For any γ such that $\gamma \not\equiv 0 \pmod{p}$ the probability $P(\gamma)$ of

$$\|\gamma t\|_p > p2^{-\mu+1}$$

for an integer t chosen uniformly at random from the elements of a Δ -homogeneously distributed sequence modulo p is

$$P(\gamma) \geq 1 - 2^{-\mu+2} - \Delta.$$

Thus for the $2^{-\mu}$ -homogeneously distributed sequence modulo p , \mathcal{T} , we have

$$P(\gamma) \geq 1 - \frac{5}{2^\mu}.$$

Therefore, for any $\beta \not\equiv \alpha \pmod{p}$,

$$\Pr \left[\exists i \in [1, d] \mid \|\beta t_i - \alpha t_i\|_p \geq p2^{-\mu+1} \right] = 1 - (1 - P(\beta - \alpha))^d \geq 1 - \left(\frac{5}{2^\mu} \right)^d,$$

where the probability is taken over integers t_1, \dots, t_d chosen uniformly and independently at random from the elements of \mathcal{T} .

Since for $\beta \not\equiv \alpha \pmod{p}$ there are only $p-1$ possible values for the residue of β modulo p , we obtain

$$\Pr \left[\forall \beta \not\equiv \alpha \pmod{p}, \exists i \in [1, d] \mid \|\beta t_i - \alpha t_i\|_p > p2^{-\mu+1} \right] \geq 1 - (p-1) \left(\frac{5}{2^\mu} \right)^d > 1 - 1/p,$$

because of the conditions of the theorem.

The rest of the proof is identical to the proof of Theorem 5 of [6], we outline it for the sake of completeness.

Let us fix some integers t_1, \dots, t_d with

$$\min_{\beta \not\equiv \alpha \pmod{p}} \max_{i \in [1, d]} \|\beta t_i - \alpha t_i\|_p > p2^{-\mu+1}. \quad (4)$$

Let \mathbf{v} be a lattice point satisfying

$$\left(\sum_{i=1}^d (v_i - s_i)^2 \right)^{1/2} \leq p2^{-\mu}.$$

Clearly, since $\mathbf{v} \in L(t_1, \dots, t_d)$, there are integers β, z_1, \dots, z_d such that

$$\mathbf{v} = (\beta t_1 - z_1 p, \dots, \beta t_d - z_d p, \beta/p).$$

If $\beta \equiv \alpha \pmod{p}$, then we are done, so suppose that $\beta \not\equiv \alpha \pmod{p}$. In this case,

$$\begin{aligned} \left(\sum_{i=1}^d (v_i - s_i)^2 \right)^{1/2} &\geq \min_{i \in [1, d]} \|\beta t_i - s_i\|_p \\ &\geq \min_{i \in [1, d]} \left(\|\beta t_i - \alpha t_i\|_p - \|s_i - \alpha t_i\|_p \right) \\ &> p2^{-\mu+1} - p2^{-\mu} = p2^{-\mu} \end{aligned}$$

that contradicts our assumption. As we have seen, the condition (4) holds with probability exceeding $1 - 1/p$ and the result follows. \square

Lemma 3. *Let $1 > \tau > 0$ be an arbitrary absolute constant and p be a prime. Assume that a real η and an integer d satisfy*

$$\eta \geq \left\lceil \left(\tau \frac{\log p \log \log \log p}{\log \log p} \right)^{1/2} \right\rceil \quad \text{and} \quad d = \lceil 5 \log p / \eta \rceil.$$

Let \mathcal{T} be a sequence of $2^{-\eta}$ -homogeneously distributed integers modulo p . There exists a probabilistic polynomial-time algorithm \mathcal{A} such that for any fixed integer $\alpha \in \mathbb{F}_p^$, given $2d$ integers*

$$t_i \quad \text{and} \quad s_i = \text{MSB}_{\eta, p}(\alpha t_i), \quad i = 1, \dots, d,$$

its output satisfies for sufficiently large p

$$\Pr[\mathcal{A}(m, t_1, \dots, t_d; s_1, \dots, s_d) = \alpha] \geq 1 - p^{-1},$$

with probability taken over all t_1, \dots, t_d chosen uniformly and independently at random from the elements of \mathcal{T} and all coin tosses of the algorithm \mathcal{A} .

Proof. We follow the same arguments as in the proof Theorem 1 of [6] which we briefly outline here for the sake of completeness. We refer to the first d vectors in the defining matrix of $L(t_1, \dots, t_d)$ as p -vectors.

Multiplying the last row vector $(t_1, \dots, t_d, 1/p)$ of the matrix (3) by α and subtracting certain multiples of p -vectors, we obtain a lattice point

$$\mathbf{u}_\alpha = (u_1, \dots, u_d, \alpha/p) \in L(t_1, \dots, t_d)$$

such that $|u_i - s_i| < p2^{-\eta}$, $i = 1, \dots, d+1$. Therefore,

$$\min \left\{ \sum_{i=1}^{d+1} (z_i - s_i)^2, \quad \mathbf{z} = (z_1, \dots, z_d, z_{d+1}) \right\} \leq \sum_{i=1}^{d+1} (u_i - s_i)^2 \leq (d+1)p^2 2^{-2\eta}.$$

Let $\mu = \eta/2$. One can verify that under the conditions of the theorem we have,

$$0.1\tau \frac{(d+1) \log \log(d+1)}{\log(d+1)} \leq \mu - 1 \quad \text{and} \quad d(\mu - \log 5) \geq 2 \log p.$$

Now we use the algorithm of Lemma 1 with $\mathbf{s} = (s_1, \dots, s_d, 0)$ to find in probabilistic polynomial time a lattice vector

$$\mathbf{v} = (v_1, \dots, v_d, v_{d+1}) \in L(t_1, \dots, t_d)$$

such that

$$\left(\sum_{i=1}^d (v_i - s_i)^2 \right)^{1/2} \leq 2^{0.1\gamma(d+1) \log \log(d+1) / \log(d+1)} p(d+1)^{1/2} 2^{-\eta} \leq p2^{-\mu-1},$$

provided that p is sufficiently large. We also have

$$\left(\sum_{i=1}^d (u_i - s_i)^2 \right)^{1/2} \leq pd^{1/2} 2^{-\eta} \leq p2^{-\mu-1}.$$

Therefore,

$$\left(\sum_{i=1}^d (u_i - v_i)^2 \right)^{1/2} \leq p2^{-\mu}.$$

Applying Lemma 2, we see that $\mathbf{v} = \mathbf{u}_\alpha$ with probability at least $1 - 1/p$, and therefore, α can be recovered in polynomial time. \square

3.3 Distribution of Exponential Functions Modulo p

To apply the results above, we will need to establish approximate uniform distribution of sequences of form $t_i = g^{u_i}$, $i = 1, 2, \dots$. A procedure to establish such

results in general is to bound certain exponential sums, related to the sequences under consideration.

The following statement is a somewhat simplified version of Theorem 4 of [8] and greatly improve several previously known bounds from [21, 25], which have been used in [16].

Lemma 4. *For any $\varepsilon > 0$ there exists $\delta > 0$ such that for any element $g \in \mathbb{F}_p$ of multiplicative order $T \geq p^\varepsilon$ we have*

$$\max_{\gcd(c,p)=1} \left| \sum_{x=0}^{T-1} \exp(2\pi icg^x/p) \right| \leq T^{1-\delta}.$$

Using Lemma 4 and arguing as in [16], we derive the following statement.

Lemma 5. *For any $\varepsilon > 0$ there exists $\delta > 0$ such that for any element $g \in \mathbb{F}_p$ of multiplicative order $T \geq p^\varepsilon$ the sequence g^x , $x = 1, \dots, T$, is $p^{-\delta}$ -homogeneously distributed modulo p .*

4 Bit Security of the Diffie–Hellman Scheme

Let us fix an element $g \in \mathbb{F}_p^*$ of multiplicative order q , where q is prime. We recall that classically, breaking the Diffie–Hellman scheme means the ability to recover the value of the *secret key* g^{xy} from publicly known values of g^x and g^y (with unknown x and y , of course).

The attacker, however, may pursue a more modest goal of recovering only partial information about the secret g^{xy} . For instance, the Legendre symbol of g^{xy} is trivially deducible from that of g^x, g^y . If only part of g^{xy} is used to derive a key for a secret key cryptosystem, this may be harmful enough. The purpose of the *bit security* results is to show that deriving such partial information is as hard as finding the whole key, which is believed to be infeasible.

It has been shown in [6, 16] that recovering (without significant errors) about $\log^{1/2} p$ most significant bits of g^{xy} for every x and y is not possible unless the whole scheme is insecure.

However, it is already dangerous enough if the attacker possesses a *probabilistic* algorithm which recovers some bits of g^{xy} only for some, not too small, fraction of key exchanges. Here we obtain first results in this direction. We consider two types of attacking algorithms:

- more traditional Monte Carlo type algorithms where our results are weaker in terms of number of bits, but stronger in error-tolerance than those of [6, 16];

- more powerful Las Vegas type algorithms where, given such an algorithm, our results are stronger than the case of deterministic algorithms obtained in [6, 16]. Cryptographic security of other schemes in this model has been studied in [27].

In fact, it is more convenient to treat a possible attacking algorithm as an *oracle* which, given g^x and g^y returns, sometimes, some information about g^{xy} . Accordingly, our purpose is to show that having such an oracle one can recover the secret key completely.

In the sequel, to demonstrate our arguments in the simplest situation we restrict ourselves to the case of most practical interest, that is, g generating a sub-group of prime order q .

4.1 Monte Carlo Type Attacks

Given positive η and γ , we define the oracle $\text{DH}_{\eta,\gamma}^{\text{MC}}$ as a “black box” which, given $g^x, g^y \in \mathbb{F}_p^*$, outputs the value of $\text{MSB}_{\eta,p}(g^{xy})$, with probability γ , taken over random pairs $(x, y) \in \mathbb{Z}_q^2$ (and possible internal coin-flips), and outputs an arbitrary value otherwise.

That is, $\text{DH}_{\eta,\gamma}^{\text{MC}}$ is a Monte Carlo type oracle which sometimes outputs some useful information and otherwise returns a wrong answer following *any* distribution. This is qualitatively thus the same type of oracles considered in [6, 16].

Theorem 1. *For any $\varepsilon > 0$ such that the following statement holds. Let $\delta > 0$ be an arbitrary positive number and let*

$$\eta = \lceil \delta \log p \rceil .$$

For any element $g \in \mathbb{F}_p^$ of multiplicative order $q \geq p^\varepsilon$, where q is prime, there exists a probabilistic algorithm which, in time polynomial in $\log p$ and $(0.25\gamma)^{-(5\delta^{-1}+1)} \log \gamma^{-1}$, for any pair $(a, b) \in \mathbb{Z}_q^2$, given the values of $g^a, g^b \in \mathbb{F}_p$, makes the expected number of $O\left(\delta^{-1}(0.25\gamma)^{-(5\delta^{-1}+1)} \log \gamma^{-1} \log \log p\right)$ calls to the oracle $\text{DH}_{\eta,\gamma}^{\text{MC}}$ and computes g^{ab} correctly with probability $1 + O(\log^{-1} p)$.*

Proof. Put $d = \lceil 5 \log p / \eta \rceil \leq 5\delta^{-1} + 1$. Given g^x, g^y the oracle $\text{DH}_{\eta,\gamma}^{\text{MC}}$ returns $\text{MSB}_{\eta,p}(g^{xy})$ with probability γ . We define an algorithm, $\mathcal{O}(g^x, g^y)$, which uses $\text{DH}_{\eta,\gamma}^{\text{MC}}$ as a black box and retrieves g^{xy} with non-negligible probability, ρ . We then apply a result by Shoup, [33], to this \mathcal{O} , and get an algorithm which retrieves g^{xy} almost surely. In the following we define and analyze \mathcal{O} .

By randomizing the second component input to $\text{DH}_{\eta,\gamma}^{\text{MC}}, g^y$, we hope to hit a set of “good” values of y , for which we have a sufficient advantage, taken over x only. We then query by randomizing the g^x -component, keeping y fixed.

Let γ_y be the average success probability of $\text{DH}_{\eta,\gamma}^{\text{MC}}(g^x, g^y)$, taken over random x for a given y . Thus, $\mathbb{E}_y[\gamma_y] = \gamma$. Let us define $k = \lceil \log(2/\gamma) \rceil$ and say that y is

j -good if $\gamma_y \in [2^{-j}, 2^{-j+1})$, $j = 1, 2, \dots, k$, and let $S_j = \{y \mid y \text{ is } j\text{-good}\}$ (thus we do not care about y for which $\gamma_y < \gamma/2$). By the Markov inequality, (2),

$$\Pr_y[\gamma_y \geq \gamma/2] \geq \frac{\gamma}{2}. \quad (5)$$

We claim that there must exist j as above for which $\Pr_y[y \in S_j] \geq 2^{j-2}\gamma/k$. If this was not the case, by (5), we would get the following contradiction:

$$\frac{\gamma}{2} \leq \sum_{j=1}^k 2^{-j+1} \Pr_y[y \in S_j] < \sum_{j=1}^k 2^{-j+1} \frac{2^{j-2}\gamma}{k} = \frac{\gamma}{2}.$$

Now, given g^a, g^b , the algorithm \mathcal{O} starts by choosing a random $v \in \mathbb{Z}_q$. Next, choose d independent random elements $u_1, \dots, u_d \in \mathbb{Z}_q$ and query the oracle $\text{DH}_{\eta, \gamma}^{\text{MC}}$ with g^{a+u_i} and (the fixed) g^{b+v} . After that we apply the algorithm of Lemma 3 to the obtained answers, and the value returned is finally output by \mathcal{O} . To analyze this, note that if $y = v + b$ is j -good for some j , the oracle $\text{DH}_{\eta, \gamma}^{\text{MC}}$ with probability 2^{-jd} returns the correct values of

$$s_i = \text{MSB}_{\eta, p} \left(g^{(a+u_i)(b+v)} \right) = \text{MSB}_{\eta, p}(\alpha t_i)$$

for every $i = 1, \dots, d$, where $\alpha = g^{a(b+v)}$ and $t_i = g^{u_i(b+v)}$. Since q is prime, t_1, \dots, t_d are distinct and applying Lemma 5 we see that the algorithm of Lemma 3 then finds α (and thus also $g^{ab} = \alpha g^{-av}$) with probability $1 - 1/p$.

The above procedure performs as stated with probability at least $2^{-jd} \Pr_v[v + b \in S_j]$. As we have seen, there must be a j for which $\Pr_v[v + b \in S_j] \geq 2^{j-2}\gamma/k$, so \mathcal{O} succeeds with probability at least

$$\rho = (1 - p^{-1}) \frac{\gamma 2^{-j(d-1)}}{4k} \geq (1 - p^{-1}) \frac{\gamma 2^{-k(d-1)}}{4k} \geq \frac{\gamma 2^{-(\log(2/\gamma)+1)(d-1)}}{8 \log(2/\gamma)} = \frac{\gamma^d 2^{-2d-1}}{\log(2/\gamma)},$$

which is $\Omega \left((0.25\gamma)^{5\delta^{-1}+1} / \log \gamma^{-1} \right)$. The above algorithm satisfies the definition of faulty Diffie-Hellman oracle given in [33]. Therefore, applying Corollary 1 of [33] (with $\varepsilon = \rho$ and $\alpha = 1/\log p$ in the notations of [33]) we finish the proof. \square

We remark that the proof of Theorem 1 only relies on the *existence* of a “good” j , but we stress that it is also possible to efficiently find this j and a corresponding v such that $v + b$ is indeed j -good. To this end, choose a random v , and query the oracle on inputs of the form g^r, g^{b+v} for random, independent r . Since r, v , and g^b are known, so is the corresponding Diffie-Hellman secret $g^{(b+v)r}$. This means that we for each r can check if the oracle is correct on this input. Repeating this for polynomially many independent r , we get a sufficiently good approximation of γ_{v+b} , on which we can base the decision on whether $v + b$ is “good” or not, see also the proof of Theorem 2.

Obviously, the algorithm of Theorem 1 remains polynomial time under the condition $\delta^{-1} \log \gamma^{-1} = O(\log \log p)$. For example, for any fixed $\delta > 0$ (that is, when η corresponds to $\Omega(\log p)$ bits) the tolerated rate of correct oracle answers can be as low as $\gamma = \Omega(\log^{-A} p)$ with some constant $A > 0$. On the other hand, if the oracle is correct with a constant rate γ , then it is enough if it outputs $\eta = O(\log p / \log \log p)$ bits. This range can be compared to the original works in [6, 16], which apply with only $O(\log^{1/2} p)$ bits from the oracle, but on the other hand requires the rate of correct answers to be $\gamma = 1 + o(1)$.

4.2 Las Vegas Type Attacks

We now turn to the more powerful type of oracles. Given positive η and A , we define the oracle $\text{DH}_{\eta, A}^{\text{LV}}$ as a “black box” which, given $g^x, g^y \in \mathbb{F}_p^*$, outputs the value of $\text{MSB}_{\eta, p}(g^{xy})$, with probability at least $\log^{-A} p$, (taken over random pairs $(x, y) \in \mathbb{Z}_q^2$ and possible internal coin-flips), and outputs an error audit message, \perp , otherwise.

That is, $\text{DH}_{\eta, A}^{\text{LV}}$ is a Las Vegas type oracle which outputs some useful (correct) information non-negligibly often and never returns a wrong answer (but rather gives no answer at all). Again, the case of $A = 0$ quantitatively corresponds to the “error-free” oracle which has been considered in [6, 16].

Theorem 2. *For any $\varepsilon > 0$ the following statement holds. Let*

$$\eta = \left\lceil \left(\frac{\tau \log p \log \log \log p}{\log \log p} \right)^{1/2} \right\rceil,$$

where $\tau > 0$ is an arbitrary absolute constant. For any element $g \in \mathbb{F}_p^*$ of multiplicative order $q \geq p^\varepsilon$, where q is prime, there exists a probabilistic polynomial time algorithm which for any pair $(a, b) \in \mathbb{Z}_q^2$, given the values of $g^a, g^b \in \mathbb{F}_p$, makes the expected number of $O((\log p)^{\max\{1, A\}} \log \log p)$ calls to the oracle $\text{DH}_{\eta, A}^{\text{LV}}$ and computes g^{ab} correctly with probability $1 + O(\log^{-A} p)$.

Proof. The proof is similar to that of Theorem 1, though for simplicity, we use a slightly rougher estimate. Let $\gamma = \log^{-A} p$ and let γ_y be as in the notation of the proof of Theorem 1.

To find $v \in \mathbb{Z}_q$ with at least $\gamma_{v+b} > \gamma/4$ choose a random $v \in \mathbb{Z}_q$. We check whether $b = \pm v \pmod{q}$, in which case we are done. Otherwise we choose $N = \lceil 20\gamma^{-1} \log \gamma^{-1} \rceil$ independent, random elements $u_1, \dots, u_N \in \mathbb{Z}_q$ and query the oracle $\text{DH}_{\eta, A}^{\text{LV}}$ with g^{a+u_i} and g^{b+v} . If the oracle returns $K \geq \gamma N/2$ queries then $\gamma_{v+b} > \gamma/4$ with probability $1 + O(\gamma)$. Indeed, by the Chernoff bound, see for example Section 9.3 of [24], we get that if $\gamma_{v+b} < \gamma/4$ then even after

$$M = \left\lceil \frac{8}{\gamma_{v+b}} \log(2/\gamma) \right\rceil \geq N$$

the oracle returns at most $2M\gamma_{v+b} \leq K$ queries with probability at least $1 - \gamma$.

By the Markov inequality (5), we see that after the expected number of $2\gamma^{-1}$ random choices of v we find v with $\gamma_{v+b} \geq \gamma/4$ with probability $1 + O(\gamma)$.

For this v we re-use the first $d = \lfloor 5 \log p / \eta \rfloor$ replies of the oracle $\text{DH}_{\eta, A}^{\text{LV}}$ which have been used for testing whether $b+v$ is “good” (if $K \geq d$) or get $d - K$ additional replies (if $K < d$, which will not happen for interesting γ). Thus we get d values

$$s_i = \text{MSB}_{\eta, p} \left(g^{(a+u_i)(b+v)} \right) = \text{MSB}_{\eta, p}(\alpha t_i)$$

where $\alpha = g^{a(b+v)}$ and $t_i = g^{u_i(b+v)}$, $i = 1, \dots, d$. Applying Lemma 5 we see that the algorithm of Lemma 3 finds α with probability at least $1 - 1/p$. Finally we compute $g^{ab} = \alpha g^{-av}$. \square

A recent paper by Hast [19] studies an oracle model which falls somewhere in between Monte Carlo and Las Vegas oracles. Specifically, [19] considers oracles which, for some $\varepsilon, \delta \in [0, 1]$, output \perp with probability $1 - \delta$, and where non- \perp answers are correct with advantage ε . Our Las Vegas oracles thus correspond to ones with non-negligible δ and the extreme case of $\varepsilon = 1$. For the specific case of Goldreich-Levin [13] based pseudo-random bit generator, Hast [19], shows that for a given (non-negligible) success-rate, the existence of oracles with small δ would indeed be more serious than existence of traditional oracles, (for which $\delta = 1$). Perhaps not surprisingly, Theorems 1 and 2 demonstrate this for the case of the Diffie-Hellman scheme, by comparing the complexity of the respective reductions.

5 Summary

We have extended existing hardness results on the Diffie-Hellman scheme to tolerate higher error-rates in the predictions, by a trade-off on the number of bits predicted. We also studied an alternative (and much stronger) computational prediction-model whose realization, albeit less likely than more classical models, would have more severe impact on the security of Diffie-Hellman. The idea to consider Las Vegas type attacks in these setting appears to be new and definitely deserves further studying.

We remark that the analysis using Las Vegas type predictors can be applied to the study of RSA as well, for instance, when analyzing the use of RSA to send reasonably short bit strings with random padding. (In particular, results concerning Δ -homogeneous distribution generalize to composite moduli). Qualitatively, such results could of course also have been derived from the bit security results in [20]. Nevertheless, the significantly tighter reductions possible from a Las Vegas oracle show (as one would expect) that the existence of such an oracle would indeed also be *quantitatively* more severe for the security of RSA.

Of course, the most intriguing open problem remains: show that even single, individual bits of the Diffie-Hellman scheme are hard to approximate.

References

1. M. Ajtai, R. Kumar and D. Sivakumar, 'A sieve algorithm for the shortest lattice vector problem', *Proc. 33rd ACM Symp. on Theory of Comput.*, ACM, 2001, 601–610.
2. M. Bellare, A. Desai, D. Pointcheval and P. Rogaway, 'Relations among notions of security for public-key encryption schemes', *Advances in Cryptology - Proceedings of CRYPTO'98. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **1462**, 1998, 26–46.
3. M. Bellare and P. Rogaway, 'Random oracles are practical: a paradigm for designing efficient protocols', *Proc. 1st ACM Computer and Communication Security' 93*, ACM Press, 1993, 62–73.
4. M. Blum and S. Micali, 'How to generate cryptographically strong sequences of pseudo-random bits', *SIAM J. Comp.*, **13**, 1984, 850–864.
5. D. Boneh and I. E. Shparlinski, 'On the unpredictability of bits of the elliptic curve Diffie–Hellman scheme', *Advances in Cryptology - CRYPTO 2001. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **2139**, 2001, 201–212.
6. D. Boneh and R. Venkatesan, 'Hardness of computing the most significant bits of secret keys in Diffie–Hellman and related schemes', *Advances in Cryptology - CRYPTO '96. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **1109**, 1996, 129–142.
7. D. Boneh and R. Venkatesan, 'Rounding in lattices and its cryptographic applications', *Proc. 8th Annual ACM-SIAM Symp. on Discr. Algorithms*, ACM, 1997, 675–681.
8. J. Bourgain and S. V. Konyagin, 'Estimates for the number of sums and products and for exponential sums over subgroups in fields of prime order', *Comptes Rendus Mathematique*, **337** (2003), 75–80.
9. R. Canetti, O. Goldreich and S. Halevi. 'The random oracle model, revisited', *Proc. 30th ACM Symp. on Theory of Comp.*, ACM, 1998, 209–218.
10. R. Cramer and V. Shoup 'Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption', *Advances in Cryptology - EURO-CRYPT 2002. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **2332**, 2002, 45–64.
11. E. El Mahassni, P. Q. Nguyen and I. E. Shparlinski, 'The insecurity of Nyberg–Rueppel and other DSA-like signature schemes with partially known nonces', *Cryptography and Lattices : International Conference, CaLC 2001. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **2146**, 2001, 97–109.
12. R. Fischlin and C. P. Schnorr, 'Stronger security proofs for RSA and Rabin bits', *J. Cryptology*, **13**, 2000, 221–244.
13. O. Goldreich and L. A. Levin, 'A hard core predicate for any one way function', *Proc. 21st ACM Symp. on Theory of Comput.*, ACM, 1989, 25–32.
14. M. I. González Vasco and M. Näslund, 'A survey of hard core functions', *Proc. Workshop on Cryptography and Computational Number Theory*, Singapore 1999, Birkhäuser, 2001, 227–256.
15. M. I. González Vasco, M. Näslund and I. E. Shparlinski, 'The hidden number problem in extension fields and its applications', *Proc. of LATIN 2002: Theoretical Informatics : 5th Latin American Symposium. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **2286**, 2002, 105–117.
16. M. I. González Vasco and I. E. Shparlinski, 'On the security of Diffie–Hellman bits', *Proc. Workshop on Cryptography and Computational Number Theory*, Singapore 1999, Birkhäuser, 2001, 257–268.

17. M. I. González Vasco and I. E. Shparlinski, 'Security of the most significant bits of the Shamir message passing scheme', *Math. Comp.*, **71**, 2002, 333–342.
18. M. Grötschel, L. Lovász and A. Schrijver, *Geometric algorithms and combinatorial optimization*, Springer-Verlag, Berlin, 1993.
19. G. Hast, 'Nearly one-sided tests and the Goldreich-Levin predicate', *Advances in Cryptology - EUROCRYPT 2003. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **2656**, 2003, 195–210.
20. J. Håstad and M. Näslund, 'The security of individual RSA and discrete log bits', *J. of the ACM*, (to appear).
21. D. R. Heath-Brown and S. V. Konyagin, 'New bounds for Gauss sums derived from k th powers, and for Heilbronn's exponential sum', *Quart. J. Math.*, **51**, 2000, 221–235.
22. N. A. Howgrave-Graham, P. Q. Nguyen and I. E. Shparlinski, 'Hidden number problem with hidden multipliers, timed-release crypto and noisy exponentiation', *Math. Comp.*, **72**, 2003, 1473–1485.
23. R. Kannan, 'Algorithmic geometry of numbers', *Annual Review of Comp. Sci.*, **2**, 1987, 231–267.
24. M. J. Kearns and U. V. Vazirani, *An introduction to computational learning theory*, MIT Press, Cambridge MA, 1994.
25. S. V. Konyagin and I. Shparlinski, *Character sums with exponential functions and their applications*, Cambridge Univ. Press, Cambridge, 1999.
26. W.-C. W. Li, M. Näslund and I. E. Shparlinski, 'The hidden number problem with the trace and bit security of XTR and LUC', *Advances in Cryptology - CRYPTO 2002. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **2442**, 2002, 433–448.
27. M. Näslund, I. E. Shparlinski and W. Whyte, 'On the bit security of NTRU', *Proc of Public Key Cryptography - PKC 2003. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **2567**, 2003, 62–70.
28. P. Q. Nguyen and I. E. Shparlinski, 'The insecurity of the digital signature algorithm with partially known nonces', *J. Cryptology*, **15**, 2002, 151–176.
29. P. Q. Nguyen and I. E. Shparlinski, 'The insecurity of the elliptic curve digital signature algorithm with partially known nonces', *Designs, Codes and Cryptography*, **30** (2003), 201–217.
30. P. Q. Nguyen and J. Stern, 'Lattice reduction in cryptology: An update', *Proc. of ANTS 2000. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **1838**, 2000, 85–112.
31. P. Q. Nguyen and J. Stern, 'The two faces of lattices in cryptology', *Proc. of CalC 2001. Lect. Notes in Comp. Sci.*, Springer-Verlag, Berlin, **2146** (2001), 146–180.
32. C. P. Schnorr, 'Security of almost all discrete log bits', *Electronic Colloq. on Comp. Compl.*, Univ. of Trier, **TR98-033**, 1998, 1–13.
33. V. Shoup, 'Lower bounds for discrete logarithms and related problems', *Preprint*, available from <http://www.shoup.net>.
34. I. E. Shparlinski, 'Security of most significant bits of g^{x^2} ', *Inform. Proc. Letters*, **83**, 2002, 109–113.
35. I. E. Shparlinski, *Cryptographic applications of analytic number theory*, Birkhauser, 2003.