Advances in Cryptology – Proceedings of CRYPTO '98
(August 23–27, 1998, Santa Barbara, California, USA)
H. Krawczyk (Ed.), vol. 1462 of Lecture Notes in Computer Science, pages 223–242

Cryptanalysis of the Ajtai-Dwork Cryptosystem

Phong Nguyen and Jacques Stern

École Normale Supérieure
Laboratoire d'Informatique
45, rue d'Ulm
F – 75230 Paris Cedex 05
{Phong.Nguyen,Jacques.Stern}@ens.fr
http://www.dmi.ens.fr/~{pnguyen,stern}/

Abstract. Recently, Ajtai discovered a fascinating connection between the worst-case complexity and the average-case complexity of some well-known lattice problems. Later, Ajtai and Dwork proposed a cryptosystem inspired by Ajtai's work, provably secure if a particular lattice problem is difficult in the worst-case. We present a heuristic attack (to recover the private key) against this celebrated cryptosystem. Experiments with this attack suggest that in order to be secure, implementations of the Ajtai-Dwork cryptosystem would require very large keys, making it impractical in a real-life environment. We also adopt a theoretical point of view: we show that there is a converse to the Ajtai-Dwork security result, by reducing the question of distinguishing encryptions of one from encryptions of zero to approximating some lattice problems. In particular, this settles the open question regarding the NP-hardness of the Ajtai-Dwork cryptosystem: from a recent result of Goldreich and Goldwasser, our result shows that breaking the Ajtai-Dwork cryptosystem is not NP-hard, assuming the polynomial-time hierarchy does not collapse.

1 Introduction

Lattices are discrete subgroups of some n-dimensional space and have been the subject of intense research, going back to Gauss, Dirichlet, Hermite and Minkowski, among others. More recently, lattices have been investigated from an algorithmic point of view and two basic problems have emerged: the shortest vector problem (SVP) and the closest vector problem (CVP). SVP refers to the question of computing the lattice vector with minimum non-zero euclidean length while CVP addresses the non-homogeneous analog of finding a lattice element minimizing the distance to a given vector. It has been known for some time that CVP is NP-complete [12] and Ajtai has recently proved that SVP is NP-hard for polynomial random reductions [3].

The celebrated LLL algorithm [18] provides a partial answer to SVP since it runs in polynomial time and approximates the shortest vector within a factor of $2^{n/2}$ where n denotes the dimension of the lattice. This has been improved to the bound $(1+\varepsilon)^n$ by Schnorr [21]. Babai [6] gave an algorithm that approximates the closest vector by a factor of $(3/\sqrt{2})^n$. The existence of polynomial bounds is completely open: CVP is presumably hard to approximate within a factor $2^{(\log n)^{0.99}}$ as shown in [5] but a result of Goldreich and Goldwasser [14] suggests that unless the polynomial-time hierarchy collapses, this inapproximability result cannot be extended to $\sqrt{n}$.

Recently, in a beautiful paper, Ajtai [2] found the first connection between the worst-case and the average-case complexity of SVP. He established a reduction from the problem of finding the shortest non zero element u of a lattice provided that it is "unique" (*i.e.* that it is polynomially shorter than any other element of the lattice which is not linearly related) to the problem of approximating SVP for randomly chosen instances of a specific class of lattices. This reduction was improved in [8]. Later, Ajtai and Dwork [4] proposed a cryptosystem inspired by Ajtai's work and proved that it was provably secure under the assumption that the "unique" shortest vector problem considered above is difficult in the worst-case.

Again, from a theoretical point of view, the achievement in the Ajtai-Dwork paper is a masterpiece. However, its practical significance is unclear. This is partly due to the fact, exemplified by RSA, that the success of a cryptosystem is not only dependent on the computational hardness of the problem on which it is based, but also on the performances that it displays in terms of speed, key size, expansion rate, *etc.* It is also related to the fact that, so far, use of lattices in cryptography has been directed at successfully breaking schemes [1, 22, 7, 17, 10, 24, 16, 9]: experiments have shown that lattice reduction algorithms behave surprisingly well and can provide much better approximations to SVP or CVP than expected.

At this point, it was natural to ask whether or not the security level offered by the Ajtai-Dwork cryptosystem is exactly measured by the hardness of approximating lattice problems. In other terms, is there a converse to the Ajtai-Dwork security result ? The present paper shows that this is actually the case by reducing the question of distinguishing encryptions of one from encryptions of zero to approximating CVP or SVP (recall that AD encrypts bits). More precisely, we prove that if one can approximate CVP within a factor $cn^{1.33}$, then one can distinguish encryptions with a constant advantage d, where c and d are related constants. This is especially interesting in view of the result of Goldreich and Goldwasser quoted above since it seems to rule out any form of NP-hardness for AD, which was an open question. We prove a similar result for SVP, with a more restrictive factor. This shows that AD is essentially equivalent to approximating the shortest vector within a polynomial ratio and allows to reverse the basic paradigm of AD: for dimensions where lattice reduction algorithms behave well in practice, AD is insecure.

This opened the way to a practical assessment of the security of AD for real-size parameters. We answer this question by presenting a heuristic attack suitable for implementation. First experiments showed that this attack was able to recover the private key in a short time for small parameters. Current experiments suggest that the attack is feasible even for real-size parameters: in order to be secure, implementations of the Ajtai-Dwork cryptosystem would require very large keys.

The remainder of the paper is organized as follows. In section 2, the Ajtai-Dwork cryptosystem is described. Section 3 presents our heuristic attack (to recover the private key) and practical experiments. Sections 4 and 5 deal with a converse to the Ajtai-Dwork security theorem. Section 4 uses a CVP approximation oracle, while section 5 uses a SVP approximation oracle. The reduction obtained in section 4 shows that breaking the Ajtai-Dwork cryptosystem is not NP-hard if the polynomial-time hierarchy does not collapse. Due to lack of space, section 4 and 5 do not include full proofs. These can be found in [20]. The appendix includes the missing proofs.

2 The Ajtai-Dwork Cryptosystem

In this section we recall the construction of Ajtai and Dwork [4], with the notations and the presentation of [15]. For any ε between 0 and $\frac{1}{2}$, we denote by $\mathbf{Z} \pm \varepsilon$ the set of real numbers for which the distance to the nearest integer is at most ε. We denote the inner product of two vectors in the Euclidean space $\mathbf{R}^n$ by $\langle x, y \rangle$. Given a set of n linearly independent vectors $w_1, \ldots, w_n$, the *parallelepiped spanned by the* w_i's is the set $P(w_1, \ldots, w_n)$ of all linear combinations of the w_i's with coefficients in $[0, 1[$. Its *width* is the minimum over i of the Euclidean distance between w_i and the hyperplane spanned by the other w_j's. Reducing a vector v modulo a parallelepiped $P(w_1, \ldots, w_n)$ means obtaining a vector $v' \in P$ such that $v' - v$ belongs to the lattice spanned by the w_i's, which we denote by $v' = v \pmod{P}$. To simplify the exposition, we present the scheme in terms of real numbers, but we always mean numbers with some fixed finite precision. Given a security parameter n (which is also the precision of the binary expansion for real numbers), we let $m = n^3$ and $\rho_n = 2^{n \log n}$. We denote by B_n the big n-dimensional cube of side-length ρ_n. We also denote by S_n the small n-dimensional ball of radius n^{-8}.

Given n, the private key is a uniformly chosen vector u in the n-dimensional unit ball. For such a private key, we denote by $\mathcal{H}_u$ the distribution on points in B_n induced by the following construction:

1. Pick a point a uniformly at random from $\{x \in B_n \ : \ \langle x, u \rangle \in \mathbf{Z}\}$.
2. Select $\delta_1, \ldots, \delta_n$ uniformly at random from S_n.
3. Output the point $v = a + \sum_i \delta_i$.

The public key is obtained by picking the points $w_1, \ldots, w_n, v_1, \ldots, v_m$ independently at random from the distribution $\mathcal{H}_u$, subject to the constraint that the width of the parallelepiped $w = P(w_1, \ldots, w_n)$ is at least $n^{-2}\rho_n$ (which is likely to be satisfied, see [4]).

Encryption is bit-by-bit. To encrypt a '0', uniformly select $b_1, \ldots, b_m$ in $\{0, 1\}$, and reduce the vector $\sum_{i=1}^{m} b_i v_i$ modulo the parallelepiped w. The vector obtained is the ciphertext. The ciphertext of '1' is just a randomly chosen vector in the parallelepiped w. To decrypt a ciphertext x with the private key u, compute

$\tau = \langle x, u\rangle$. If $\tau \in \mathbf{Z} \pm n^{-1}$, then x is decrypted as '0', and otherwise as '1'. Thus, an encryption of '0' will always be decrypted as '0', and an encryption of '1' has a probability of $2n^{-1}$ to be decrypted as '0'. These decryption errors can be removed (see [15]). The main result of [4] states that a probabilistic algorithm distinguishing encryptions of a '0' from encryptions of a '1' with some polynomial advantage can be used to find the shortest nonzero vector in any n-dimensional lattice where the shortest vector v is unique, in the sense that any other vector whose length is at most $n^8||v||$ is parallel to v.

3 A Practical Attack

We describe in this section a heuristic attack to recover the private key. We first present the ideas underlying our attack, the attack itself and then the experiments. Let $(u, w_1, \ldots, w_n, v_1, \ldots, v_m)$ be a set of keys. For any real $\beta > 0$, denote by Λ_β the m-dimensional lattice (in $\mathbf{R}^{n+m}$) spanned by the columns of the following matrix:

$$\begin{pmatrix} \beta v_1 & \beta v_2 & \ldots & \beta v_m \\ 1 & 0 & \ldots & 0 \\ 0 & 1 & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \ldots & 0 & 1 \end{pmatrix}$$

Each $\langle v_i, u\rangle$ belongs to $\mathbf{Z} \pm n^{-7}$: let V_i be the closest integer to $\langle v_i, u\rangle$. The following result shows that short vectors in Λ_β give information on the V_i's:

Theorem 1. *Let $x = {}^t\big(\beta(\lambda_1 v_1 + \cdots + \lambda_m v_m), \lambda_1, \ldots, \lambda_m\big)$ be a point of Λ_β, the λ_i's being integers. If $n^7\big\|\sum_{i=1}^m \lambda_i v_i\big\| + \sum_{i=1}^m |\lambda_i| < n^7$, then $\sum_{i=1}^m \lambda_i V_i = 0$. In particular, this equality is satisfied if $\beta^2 \geq \frac{1}{2n^7-1} n^{14}$ and $\|x\| < \frac{1}{\sqrt{2n^7-1}} n^7$.*

Proof. By definition of the v_i's, $\big|\langle\sum_{i=1}^m \lambda_i v_i, u\rangle - \sum_{i=1}^m \lambda_i V_i\big| \leq n^{-7} \sum_{i=1}^m |\lambda_i|$. If $|\langle\sum_{i=1}^m \lambda_i v_i, u\rangle| < 1 - n^{-7} \sum_{i=1}^m |\lambda_i|$, then the integer $\sum_{i=1}^m \lambda_i V_i$ is zero since it is strictly less than 1 in absolute value. As $\|u\| \leq 1$, a stronger condition is $\|\sum_{i=1}^m \lambda_i v_i\| < 1 - n^{-7} \sum_{i=1}^m |\lambda_i|$ by the Cauchy-Schwarz inequality, and this proves the first statement. Squared, it becomes:

$$\Big\|\sum_{i=1}^m \lambda_i v_i\Big\|^2 < 1 + n^{-14} \sum_{i=1}^m \lambda_i^2 - 2n^{-7} \sum_{i=1}^m |\lambda_i|.$$

But $|\lambda_i| \leq \lambda_i^2$ since the λ_i's are integers. This gives a new stronger condition:

$$n^7 \Big\|\sum_{i=1}^m \lambda_i v_i\Big\|^2 + (2 - n^{-7}) \sum_{i=1}^m \lambda_i^2 < n^7,$$

which is satisfied as soon as $\beta^2 \geq \dfrac{n^{14}}{2n^7 - 1}$ and $\|x\| < \dfrac{n^7}{\sqrt{2n^7 - 1}}$. □

The following combinatorial results suggest that Λ_β contains many sufficiently short vectors.

Theorem 2. *For all $\varepsilon > 0$, there exists N such that the following holds for all $n \geq N$. Let $\{i_1, i_2, \ldots, i_{m'}\}$ be a subset of $\{1, 2, \ldots, m\}$. If $m' \geq (1+\varepsilon)n^2 \log_2 n$, then there exist $\lambda_1, \lambda_2, \ldots, \lambda_{m'}$ (not all zero) in $\{-1, 0, 1\}$ such that*

$$\left\| \lambda_1 v_{i_1} + \lambda_2 v_{i_2} + \cdots + \lambda_{m'} v_{i_{m'}} \right\| \leq n^{3.5 - n/\log_2 n}.$$

Proof. Let $\alpha = n/\log_2 n$ and $\bar{v}_i = \lfloor n^\alpha v_i \rfloor$. Each vector $\bar{v}_i$ has integral entries in the set $\{-n^\alpha \rho_n, \ldots, n^\alpha \rho_n\}$. Consider all combinations of $\bar{v}_{i_1}, \ldots, \bar{v}_{i_{m'}}$ with coefficients in $\{0, 1\}$. There are $2^{m'}$ such combinations. But there are at most $(2m' n^\alpha \rho_n + 1)^n$ distinct values for such combinations. By the pigeon-hole principle, it follows that if $2^{m'} > (2m' n^\alpha \rho_n + 1)^n$, then there exist $\lambda_1, \lambda_2, \ldots, \lambda_{m'}$ (not all zero) in $\{-1, 0, 1\}$, such that $\lambda_1 \bar{v}_{i_1} + \lambda_2 \bar{v}_{i_2} + \cdots + \lambda_{m'} \bar{v}_{i_{m'}} = 0$. Hence:

$$\sum_{k=1}^{m} \lambda_k v_{i_k} = \frac{\sum_{k=1}^{m} \lambda_k \left(n^\alpha v_{i_k} - \lfloor n^\alpha v_{i_k} \rfloor \right)}{n^\alpha},$$

whose norm is less than $n^{-\alpha} \sum_{k=1}^{m} \sqrt{n} = n^{3.5 - n/\log_2 n}$. Furthermore,

$$\begin{aligned}
\log_2 (2m' n^\alpha \rho_n + 1)^n &\leq n \log_2(2m') + \alpha n \log_2 n + n \log_2 \rho_n + n \log_2(1+1) \\
&\leq 2n + 3n \log_2 n + n^2 + n^2 \log_2 n
\end{aligned}$$

We conclude since $2n + 3n\log_2 n + n^2 + n^2 \log_2 n + < (1+\varepsilon) n^2 \log_2 n$ for sufficiently large n. $\square$

Corollary 3. *For all $\varepsilon > 0$, there exists N such that for all $n \geq N$ and all $\beta > 0$, there exist at least $n^3 - (1+\varepsilon)n^2 \log_2 n$ linearly independent lattice points in Λ_β, with norm less than $\sqrt{n^3 + \beta^2 n^{7 - 2n/\log_2 n}}$.*

We now use the notion of an orthogonal lattice introduced in [19]: if L is a lattice in $\mathbf{Z}^n$, the orthogonal lattice $L^\perp$ is defined as the set of points in $\mathbf{Z}^n$ that are orthogonal to all the lattice points. Consider the vector V in $\mathbf{Z}^m$ whose coordinates are the V_i's. Theorem 1 shows that sufficiently short vectors of Λ_β correspond to vectors in $V^\perp$, which is a $m-1$ dimensional lattice in $\mathbf{Z}^m$. Corollary 3 shows that many such short vectors exist. We conjecture that there exist at least $m-1$ sufficiently short and linearly independent vectors in Λ_β.

If one knows $m-1$ linearly independent vectors in $V^\perp$, then one can determine the one-dimensional lattice $(V^\perp)^\perp$: one can find a vector $V' \in \mathbf{Z}^m$ generating $(V^\perp)^\perp$. There exists $\varepsilon \in \mathbf{Z}$ such that $V = \varepsilon V'$. If all the V_i's are coprime (which happens with overwhelming probability), then $\varepsilon = \pm 1$. Since one can exchange $-u$ for u, one can assume that $\varepsilon = 1$. And if one knows the V_i's, one can obtain an approximation of the private key u by solving a linear system. This is because each $\langle v_i, u \rangle \approx V_i$ gives rise to a linear equation whose unknowns are the coordinates of u. If A is the matrix representing $(v_1, \ldots, v_n)$ with respect to the canonical basis, then the multiplication of A^{-1} by the vector formed by

$V_1, \ldots, V_n$ is an approximation of u (one does not need to know all the V_i's, n of them are enough). The approximation is good because $\langle v_i, u\rangle$ is close to V_i (difference less than n^{-7} in absolute value) and the coefficients of A^{-1} are very small (we omit the details but one can justify that they are roughly around ρ_n^{-1} in absolute value, because the coefficients of A can almost be considered as independent and uniformly distributed over $]-\rho_n, \rho_n[$).

The attack is the following:

1. Obtain vectors in $V^\perp$ by finding short linear combinations of the v_i's using lattice reduction algorithms.
2. When enough vectors in $V^\perp$ are found, compute $(V^\perp)^\perp$ and $\pm V$.
3. Solve the linear system $A.u' = V'$ to obtain an approximation u' to the private key u, where A is the matrix representing $v_1, \ldots, v_n$ and V' is the vector formed by $V_1, \ldots, V_n$.

For step 1, we do not reduce a complete lattice Λ_β (whose dimension m might be too large): we only keep m' random columns and reduce them. There are heuristic arguments explaining why one can still expect to find short linear combinations with only m' vectors v_i's instead of m. Due to lack of space, we omit the details here: it is related to the fact that, given a low-dimensional lattice, a reduced basis for the corresponding orthogonal lattice is much smaller than for the lattice itself (see [19]). To find enough vectors, one repeats the random selection of columns. Since short vectors are found in an apparently random fashion, one can expect to find as many vectors as wanted.

For step 2, one has to compute a basis of the orthogonal lattice of a given lattice. To do so, one can use the polynomial time algorithm given in [19], which uses lattice reduction algorithms. But there is a more practical method here: since the orthogonal lattice is one-dimensional only, one can compute it by a basic cross product, that is determinant computations. Actually one does not need to compute the complete cross product: n determinants suffice instead of m, because in step 3, only n coordinates of V are used. Note that each determinant is a $n \log n$-bit integer.

We used the NTL library [23] to conduct our experiments. Timings are given for a 170 Mhz Ultra-SPARC-I. We used the floating point variant (double and quadratic precision) of the LLL algorithm as our lattice reduction algorithm: no stronger algorithms were needed. Only steps 1 and 2 are time-consuming. For $n = 8$, we reduced the complete 512-dimensional lattice Λ_β: step 1 took 3 hours, step 2 took less than half an hour and step 3 was immediate. The approximation u' was matching u with the n-bit precision. For $n = 32$, we chose $m' = 300$: each reduction of a partial 300-dimensional lattice gave 60 vectors in $V^\perp$ in less than 4 hours. Actually, this running time might be decreased, as a complete LLL reduction is unnecessary. Hence, one could expect to find enough vectors in less than 100 days on a single machine, and the computations can easily be parallelized (each random combination of 300 columns can be reduced independently) to reduce the running time. Step 2 requires the computation of 32 determinants of huge but sparse matrices. The dimension of these matrices is

32, 767, and there are less than 300 non-zero entries per line, making the computation feasible. Assuming that step 2 determined the V_i's, step 3 immediately gave a perfect approximation of u.

Hence a successful attack even for $n = 32$ looks feasible. Note that at least $n^5 \log n$ bits are required to store the public key (the v_i's): for $n = 32$, this amounts to 20 megabytes, and the ciphertext for each bit is 6144 bits long. This shows that the Ajtai-Dwork cryptosystem is hardly practical even with marginal security.

4 Deciphering with a CVP-oracle

We define an (n, k)-CVP-oracle to be any algorithm which, given a point $x \in \mathbf{R}^n$ and a n-dimensional lattice L, outputs a lattice point $\alpha \in L$ such that for every $\beta \in L$: $\mathrm{dist}(x, \alpha) \leq k\mathrm{dist}(x, \beta)$, where *dist* denotes the Euclidean distance. Each oracle call made by a Turing machine contributes by a single unit to the overall complexity of the machine.

Using such an oracle, we will see how one can distinguish in probabilistic polynomial time ciphertexts of '0' from ciphertexts of '1', thanks to some properties of the keys. To any choice of the keys, we associate a particular lattice. Given a ciphertext, one can build a vector such that: if the ciphertext is a ciphertext of '0', this vector is likely to be close to the lattice ; and if the ciphertext is a ciphertext of '1', this vector is unlikely to be close enough. To check whether this vector is close enough, one calls an oracle.

4.1 Vulnerable keys

Theorem 4. *For sufficiently large n, for any ε_1 and ε_2 in $]0, 1[$, any set of keys $(u, w_1, \ldots, w_n, v_1, \ldots, v_m)$ picked at random as described in Ajtai-Dwork's protocol satisfies the following with probability at least $(1 - \varepsilon_1)(1 - \varepsilon_2)$:*

$$\sum_{j=1}^{n} dist(\mathbf{Z}, \langle u, w_j \rangle)^2 \leq \frac{2\pi}{n^{16}\varepsilon_1} \tag{1}$$

$$E\left[\sum_{j=1}^{n} \left\langle \sum_{i=1}^{m} b_i v_i, w_j^{\perp} \right\rangle^2\right] \leq \frac{n^4 \rho_n^2}{2\varepsilon_2} \tag{2}$$

where $w_j^{\perp}$ denotes the unit vector orthogonal to the hyperplane spanned by the other w_j's, and the expectation is with respect to a uniform random choice of $(b_1, \ldots, b_m)$ in $\{0, 1\}^m$.

We show how to prove this result, which will be used afterwards. Let u be a non-zero private key: $\|u\| \leq 1$. We start with a technical lemma:

Lemma 5. *Let δ be a randomly chosen point from S_n. Then $E[\langle u, \delta \rangle] = 0$ and $Var[\langle u, \delta \rangle] = \frac{4\|u\|^2 W_n^2}{(n+2)n^{16}}$, where $W_n = \int_0^{\pi/2} \sin^n \theta d\theta$ is the n-th Wallis integral.*

Proof (Sketch). The expectation $E[\langle u, \delta\rangle]$ is clearly zero. To compute the variance, we can assume that $u = (\|u\|, 0, 0, \ldots, 0)$ since S_n is invariant by rotation. We obtain:

$$\mathrm{Var}[\langle u, \delta\rangle] = \|u\|^2 \int_{-n^{-8}}^{n^{-8}} x^2 \frac{V_{n-1}(\sqrt{n^{-16} - x^2})}{V_n(n^{-8})} dx,$$

where $V_n(r)$ denotes the volume of the n-dimensional ball of radius r. The result follows after a few simplifications using Wallis integrals. $\square$

This leads to a more general result:

Lemma 6. *Let v be a randomly chosen point from the distribution $\mathcal{H}_u$. Then:*

$$E\left[\mathit{dist}(\mathbf{Z}, \langle u, v\rangle)^2\right] \leq \frac{2\pi}{(n+2)n^{16}}.$$

Proof (Sketch). Write $v = a + \sum_i \delta_i$ where the δ_i's are independently chosen with uniform distribution over S_n. Apply the previous lemma with δ_i as δ. Conclude as $W_n^2 \leq 2\pi/n$ and $\|u\| \leq 1$. $\square$

Denote by X the random variable $\sum_{j=1}^n \mathrm{dist}(\mathbf{Z}, \langle u, w_j\rangle)^2$, where the w_j's are chosen according to Ajtai-Dwork's rules. From the previous lemma:

$$E[X] = \sum_{j=1}^n E\left[\mathrm{dist}(\mathbf{Z}, \langle u, w_j\rangle)^2\right] \leq n\frac{2\pi}{(n+2)n^{16}} \leq \frac{2\pi}{n^{16}}.$$

By Markov's inequality, it follows that (1) is satisfied with probability at least $1 - \varepsilon_1$ over the choice of $w_1, \ldots, w_n$.

Now, we assume that the the w_j's are fixed and satisfy (1). We will prove that for sufficiently large n, when $(v_1, \ldots, v_m)$ and $(b_1, \ldots, b_m)$ are independently picked at random as described in Ajtai-Dwork's protocol,

$$E\left[\sum_{j=1}^n \left\langle \sum_{i=1}^m b_i v_i, w_j^\perp \right\rangle^2\right] \leq \frac{n^4 \rho_n^2}{2}. \tag{3}$$

Thus, by Markov's inequality, (2) is satisfied with probability at least $1 - \varepsilon_2$ over the choice of $v_1, \ldots, v_m$, which completes the proof of Theorem 4.

To prove (3), it suffices to prove that for sufficiently large n, for all choice of $(b_1, \ldots, b_m)$, (3) is satisfied with respect to a random choice of $(v_1, \ldots, v_m)$. The core of this result is the following basic lemma:

Lemma 7. *Let t on the n-dimensional unit sphere. Let s be a randomly chosen point (with uniform distribution) from the hypercube B_n. Then $E[\langle s, t\rangle] = 0$ and $E[\langle s, t\rangle^2] = \rho_n^2/3$.*

Proof (Sketch). Decompose s and t with respect to the canonical basis to express the dot product $\langle s, t\rangle$. The result follows from a short computation, using the fact that the coordinates of s are independent random variables uniformly distributed over $]-\rho_n, +\rho_n[$. $\square$

Now, we fix $b_1, \ldots, b_m$ in $\{0,1\}$ and denote by X the random variable of (3), for which we want to bound the expectation.

Assume first that the v_i's are independent random variables uniformly distributed over the hypercube B_n. Then, applying Lemma 7 several times:

$$E[X] = \sum_{j=1}^{n} \sum_{i=1}^{m} b_i^2 E[\langle v_i, w_j^{\perp} \rangle^2] \leq nm \frac{\rho_n^2}{3} \leq n^4 \frac{\rho_n^2}{3}.$$

To conclude, we show how to take care of the actual distribution of the v_i's. Let a denote a point chosen at random from $\{x \in B_n \ : \ \langle x, u \rangle \in \mathbf{Z}\}$. Let λ be randomly chosen in $[0,1[$. Then, the sum $a + \lambda u$ is uniformly distributed over an n-dimensional volume C_n, which differs from B_n by points y such that the segment $[y, y+u]$ crosses the border of B_n. Such points are within distance 1 of this border. It follows that one can bound the volume of the difference of B_n and C_n by $2n\rho_n^{n-1}$. Replacing the uniformly distributed variable v_i by $a_i + \lambda_i u$ chosen according to the above distribution, one sees that $E[X]$ is modified by at most $2n\rho_n^{n-1}/\rho_n^n \times n(m\rho_n\sqrt{n})^2 = 2n^9\rho_n$ since each $\langle v_i, w_j^{\perp} \rangle$ is less than $\rho_n\sqrt{n}$. Noting that the actual v_i is obtained from some instance of a_i by adding a small perturbation vector δ_i, and that $2n^9\rho_n = o(n^4\rho_n^2/3)$ as n grows, we obtain for sufficiently large n,

$$E[X] \leq n^4 \frac{\rho_n^2}{3}(1 + 1/2) \leq \frac{n^4\rho_n^2}{2}.$$

4.2 Deciphering

For any real β, let L_β be the $n+m$-dimensional lattice (in $\mathbf{R}^{2n+m}$) spanned by the columns of the following matrix:

$$\begin{pmatrix} \beta w_1 & \ldots & \beta w_n & \beta v_1 & \ldots & \beta v_m \\ 1 & 0 & & \ldots & & 0 \\ 0 & \ddots & & & & \\ & & 1 & \ddots & & \vdots \\ \vdots & & \ddots & n^2\sqrt{n} & & \\ & & & & \ddots & 0 \\ 0 & & \ldots & & 0 & n^2\sqrt{n} \end{pmatrix}$$

The following proposition shows that a ciphertext of '0' is, in some sense, close to this lattice.

Proposition 8. *Let $\varepsilon > 0$ and $(u, w_1, \ldots, w_n, v_1, \ldots, v_m)$ satisfying (2). A ciphertext x of '0' satisfies with probability at least $1 - \varepsilon$: for all $\beta > 0$,*

$$dist\left(\begin{pmatrix} \beta x \\ 0 \end{pmatrix}, L_\beta\right) \leq \sqrt{1 + \frac{1}{2\varepsilon_2\varepsilon}}n^4.$$

Proof. Any ciphertext x of '0' is of the form $x = \sum_{i=1}^m b_i v_i + \sum_{j=1}^n \alpha_j w_j$ where $b_i \in \{0,1\}$ and $\alpha_j \in \mathbf{Z}$. We prove that the vector $X = {}^t(\beta x, 0)$ is close enough to the lattice point $Y = {}^t(\beta x, \alpha_1, \ldots, \alpha_n, b_1, \ldots, b_m)$. We have $\alpha_j = \lfloor \theta_j \rfloor$ where the θ_j's are defined by: $\sum_{i=1}^m b_i v_i = \sum_{j=1}^n \theta_j w_j$. Since the width of the parallelepiped $P(w_1, \ldots, w_n)$ is at least $n^{-2}\rho_n$, we have:

$$\sum_{j=1}^n \alpha_j^2 \leq \sum_{j=1}^n \theta_j^2 \leq \frac{n^4}{\rho_n^2} \sum_{j=1}^n \left\langle \sum_{i=1}^m b_i v_i, w_j^\perp \right\rangle^2 .$$

Applying Markov's inequality to (2), we obtain with probability at least $1 - \varepsilon$ over the choice of $b_1, \ldots, b_m$:

$$\sum_{j=1}^n \alpha_j^2 \leq \frac{n^4}{\rho_n^2} \times \frac{n^4 \rho_n^2}{2\varepsilon_2 \varepsilon} = \frac{n^8}{2\varepsilon_2 \varepsilon}.$$

Therefore: $\operatorname{dist}(X, L_\beta) \leq \operatorname{dist}(X, Y) \leq \sqrt{\frac{n^8}{2\varepsilon_2\varepsilon} + n^3 n^5} \leq \sqrt{1 + \frac{1}{2\varepsilon_2\varepsilon}}\, n^4.$ ☐

Somehow, there is a converse to the previous proposition:

Proposition 9. *Let $\varepsilon > 0$ and $(u, w_1, \ldots, w_n, v_1, \ldots, v_m)$ satisfying (1). Let y be a point in the parallelepiped $w = P(w_1, \ldots, w_n)$.*

$$\textit{If dist}\left(\begin{pmatrix} \beta y \\ 0 \end{pmatrix}, L_\beta\right) \leq \varepsilon \sqrt{\frac{\varepsilon_1}{2\pi}} n^8 \textit{ then } \langle u, y \rangle \in \mathbf{Z} \pm \varepsilon \left(1 + \sqrt{\frac{\varepsilon_1}{2\pi}} \left(1 + \frac{n^8}{\beta}\right)\right).$$

Proof. The vector βy is of the form $\beta \left(\sum_{i=1}^m b_i v_i + \sum_{j=1}^n \alpha_j w_j\right) + e$, where $||e||^2$ and $\sum_{i=1}^m b_i^2 n^5 + \sum_{i=1}^n \alpha_j^2$ are both less than $\varepsilon^2 \varepsilon_1 n^{16}/(2\pi)$. Thus,

$$\operatorname{dist}(\mathbf{Z}, \langle u, y \rangle) \leq \sum_{i=1}^m |b_i| \operatorname{dist}(\mathbf{Z}, \langle u, v_i \rangle) + \sum_{j=1}^n |\alpha_j| \operatorname{dist}(\mathbf{Z}, \langle u, w_j \rangle) + \frac{\varepsilon}{\beta} \sqrt{\frac{\varepsilon_1}{2\pi}} n^8.$$

By the Cauchy-Schwarz inequality and the fact that each $\langle v_i, u \rangle \in \mathbf{Z} \pm n^{-7}$, the first term is bounded by $\sqrt{\varepsilon^2 \varepsilon_1 n^{11}/(2\pi)} \times \sqrt{mn^{-14}} = \varepsilon \sqrt{\varepsilon_1/(2\pi)}$. Also, the second term is less than:

$$\sqrt{\sum_{j=1}^n \alpha_j^2} \times \sqrt{\sum_{j=1}^n \operatorname{dist}(\mathbf{Z}, \langle u, w_j \rangle)^2}.$$

We know that the first term of this product is less than $\varepsilon \sqrt{\varepsilon_1/(2\pi)} n^8$. And (1) bounds the second term. We conclude from all the inequalities obtained. ☐

If we collect these two propositions, we obtain a probabilistic reduction:

Theorem 10. *There exists N such that for all $\sigma, \sigma_1, \sigma_2 > 0$, there exists a polynomial time Turing machine taking a public key and a ciphertext x as an input and making a single call to a $(n+m, n^{4-(3\sigma+\sigma_1+\sigma_2)/2}/[\sqrt{\pi}(1+2n^{-\sigma-\sigma_2})])$-CVP-oracle which outputs a yes/no answer such that: for all $n \geq N$, if the keys are picked at random as described in Ajtai-Dwork's protocol, then with a probability of at least $(1-n^{-\sigma_1})(1-n^{-\sigma_2})$,*

- *If x is a ciphertext of '0', the answer is yes with probability at least $1-n^{-\sigma}$.*
- *If x is a ciphertext of '1', the answer is yes with probability at most $3n^{-\sigma}$.*

Proof. We let $\varepsilon_1 = n^{-\sigma_1}$ and $\varepsilon_2 = n^{-\sigma_2}$. For sufficiently large n (independently of σ_1 and σ_2), (1) and (2) are satisfied with probability at least $(1-\varepsilon_1)(1-\varepsilon_2)$ over the choice of the public key by Theorem 4. We let $\varepsilon = n^{-\sigma}$ and $\beta = 4n^8\sqrt{\varepsilon_1/(2\pi)}$. Calling once the CVP-oracle above, we obtain a lattice point $\alpha \in L_\beta$ such that, for all $\gamma \in L_\beta$:

$$\mathrm{dist}\left(\alpha, \begin{pmatrix} \beta x \\ 0 \end{pmatrix}\right) \leq \frac{\varepsilon\sqrt{\varepsilon_1/(2\pi)}}{\sqrt{1+1/(2\varepsilon_2\varepsilon)}} n^4 \mathrm{dist}\left(\gamma, \begin{pmatrix} \beta x \\ 0 \end{pmatrix}\right).$$

The machine outputs 'yes' if and only if:

$$\mathrm{dist}\left(\alpha, \begin{pmatrix} \beta x \\ 0 \end{pmatrix}\right) \leq \varepsilon\sqrt{\frac{\varepsilon_1}{2\pi}} n^8.$$

If x is a ciphertext of '0', Proposition 8 then ensures that the answer is 'yes' with probability at least $1-\varepsilon$. Now, if this inequality is satisfied, Proposition 9 implies that: $\langle u, x\rangle \in \mathbf{Z} \pm \varepsilon(1+\frac{1}{4}+\frac{1}{4}) = \mathbf{Z} \pm \frac{3}{2}\varepsilon$. But this happens with probability at most 3ε if x is a ciphertext of '1'. □

5 Deciphering with a SVP-oracle

We now show how to use SVP-oracles. Given a n-dimensional lattice L, an (n,k)-SVP-oracle outputs a point $\alpha \in L$ such that for every $\beta \in L$: $\|\alpha\| \leq k\|\beta\|$. The main result of this section is the following:

Theorem 11. *Let $\theta, \gamma > 0$ such that $\frac{5\gamma}{2} + 2\theta < 2$. For all $\sigma_1, \sigma_2 > 0$, there exists $N > 0$, $\sigma \in \ ]0\,;3+3/5[$ and a polynomial time oracle Turing machine calling a $(n^{2+\gamma}, n^\theta)$-SVP-oracle such that: for all $n \geq N$, if the keys are picked at random as described in Ajtai-Dwork's protocol, then with a probability of at least $(1-n^{-\sigma_1})(1-n^{-\sigma_2})$, the machine distinguishes encryptions of '0' from encryptions of '1' with polynomial advantage $n^{-\sigma}$.*

Note: recall that the advantage ε of a distinguishing algorithm $\mathcal{A}$ is such that

$$P[\mathcal{A} \text{ answers correctly}] \geq \frac{1}{2} + \varepsilon.$$

We will need a technical improvement over the computations of section 4 which reads as the following generalization of Theorem 4, proved in the appendix. The key to the improvement is to replace Markov's inequality by moments inequalities, using the multinomial formula.

Theorem 12. *Let k be a positive integer. There exists M_1 and M_2 such that for sufficiently large n: for any choice of ε_1 and ε_2 in $]0,1[$, any set of keys $(u, w_1, \ldots, w_n, v_1, \ldots, v_m)$ picked at random as described in Ajtai-Dwork's protocol satisfies the following with probability at least $(1-\varepsilon_1)(1-\varepsilon_2)$:*

$$\sum_{j=1}^{n} dist(\mathbf{Z}, \langle u, w_j \rangle)^2 \leq \frac{M_1}{n^{16}\varepsilon_1^{1/k}} \tag{4}$$

$$E\left[\left(\sum_{j=1}^{n}\left\langle \sum_{i=1}^{m} b_i v_i, w_j^{\perp}\right\rangle^2\right)^k\right] \leq \frac{n^{4k}\rho_n^{2k}M_2}{\varepsilon_2} \tag{5}$$

This leads to the following results:

Lemma 13. *For all k, there exists M_3 such that: if $(u, w_1, \ldots, w_n, v_1, \ldots, v_m)$ satisfies (5), then a random ciphertext y of '0' is, with probability at least $1-\varepsilon_3$, of the form $y = \sum_{i=1}^{m} b_i v_i + \sum_{j=1}^{n} \alpha_j w_j$, where $b_i \in \{0,1\}$, $\alpha_j \in \mathbf{Z}$ and*

$$\sum_{j=1}^{n} \alpha_j^2 \leq M_3 n^8 \frac{1}{(\varepsilon_2\varepsilon_3)^{1/k}} \tag{6}$$

Proof (Sketch). Apply Markov's inequality to the random variable of (5), then extract k-th roots. Conclude with $M_3 = M_2^{1/k}$, by bounding the sum of the α_j^2 as in the proof of Proposition 8. □

Ciphertexts of '0' satisfying (6) are called *good ciphertexts*. Note that it is possible to produce good ciphertexts, given the public key, by a polynomial time algorithm.

Lemma 14. *For all k, there exists M_4 such that: if $(u, w_1, \ldots, w_n, v_1, \ldots, v_m)$ satisfies (4), then any good ciphertext y of '0' satisfies*

$$dist(\mathbf{Z}, \langle u, y \rangle) \leq M_4 \frac{1}{n^4(\varepsilon_1\varepsilon_2\varepsilon_3)^{1/2k}}.$$

Proof (Sketch). Decompose y with the b_i's and the α_j's. Conclude by Cauchy-Schwarz thanks to (6) and (4), with $M_4 = 1 + \sqrt{M_1 M_3}$. □

We now fix some constants. Since $2\theta + \frac{5\gamma}{2} < 2$, there exist strictly positive $\gamma_1, \gamma_2, \sigma_3, k, \lambda$ such that

$$2\theta + \frac{3\gamma}{2} + \gamma_2 + \lambda + \frac{1}{2k}(\sigma_1 + \sigma_2 + \sigma_3) < 2,$$

with:

$$4/5 > \gamma_2 > \gamma_1 > \gamma, \ \gamma_1 < \gamma + \lambda, \ \text{and } \sigma_3 > 2(2 + \gamma + \gamma_1).$$

We let $\varepsilon_1 = n^{-\sigma_1}$, $\varepsilon_2 = n^{-\sigma_2}$ and $\varepsilon_3 = n^{-\sigma_3}$. We assume that the keys satisfy (4) and (5) (which happens with probability at least $(1-\varepsilon_1)(1-\varepsilon_2)$ for sufficiently large n). We will use our oracle as follows: let $\nu = n^{2+\gamma}$ and consider a sequence $(y_1, \ldots, y_\nu)$ of elements of $P(w_1, \ldots, w_n)$. Choose a random permutation p of $\{1, \ldots, \nu\}$ and apply the $(n^{2+\gamma}, n^\theta)$-SVP-oracle to the lattice spanned by the columns of the following matrix, with $\beta = n^6 n^{1+\frac{\gamma}{2}}$:

$$\begin{pmatrix} \beta y_{p(1)} & \beta y_{p(2)} & \cdots & \beta y_{p(\nu)} \\ 1 & 0 & \cdots & 0 \\ 0 & 1 & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & 1 \end{pmatrix}$$

The output is a vector $(z, \lambda_1, \ldots, \lambda_\nu)$. Say that y_i is *hit* if:

$$0 < |\lambda_{p^{-1}(i)}| \leq n^{\frac{\gamma}{2}+\theta+\lambda}.$$

The following two propositions (proved in the appendix) show that ciphertexts of '0' and '1' behave differently.

Proposition 15. *If $y_1, \ldots, y_\nu$ are ciphertexts of '1', then y_1 is hit with probability $\Omega(n^{-\gamma_1})$.*

Proposition 16. *If y_1 is a ciphertext of '1' and $y_2, \ldots, y_\nu$ are good ciphertexts of '0', then y_1 is hit with probability $\mathcal{O}\left(n^{-\gamma_2}\right)$.*

We show how to conclude. The distributions $S_\nu = (y_1, \ldots, y_\nu : y_i$ is a ciphertext of '1') and $T_\nu = (y_1, \ldots, y_\nu : y_1$ is a ciphertext of '1' and the others are good ciphertexts of '0') are distinguished by the test "y_1 is hit" with advantage $\Omega(n^{-\gamma_1})$. Using the "hybrid technique" (see [13]), we introduce the distributions $S_i = (y_1, \ldots, y_\nu : y_1, \ldots, y_i$ are ciphertexts of '1' and $y_{i+1}, \ldots, y_\nu$ are good ciphertexts of '0'). There exists i such that S_{i-1} and S_i are distinguished by the test with advantage:

$$\Omega(n^{-\gamma_1}/\nu) = \Omega(n^{-2-\gamma_1-\gamma}).$$

One can check whether a given y is a ciphertext of '0' or '1' by querying the answer of the test for $(y_1, \ldots, y_{i-1}, y, y_{i+1}, \ldots, y_\nu)$ where $y_1, \ldots, y_{i-1}$ are random ciphertexts of '1' and $y_{i+1}, \ldots, y_\nu$ are random good ciphertexts of '0'. Since the bad ciphertexts of '0' form a set of probability less than $\varepsilon_3 = n^{-\sigma_3}$ where $\sigma_3 > 2(2 + \gamma + \gamma_1)$, the distinguisher has (for sufficiently large n) polynomial advantage $n^{-\sigma}$ if $\sigma > 2 + \gamma + \gamma_1$. But:

$$2 + \gamma + \gamma_1 < 2 + \frac{4}{5} + \frac{4}{5} = 3 + \frac{3}{5}.$$

Therefore, σ can be chosen strictly less than $3 + 3/5$, and the result follows. Note: the above construction is non-uniform. Eliminating the non-uniformity requires "sampling" the test for the various distributions S_i (see [13]).

6 Conclusion

We have shown how to reduce the question of distinguishing encryptions of one from encryptions of zero in the Ajtai-Dwork cryptosystem to approximating CVP or SVP. For the sake of simplicity, our results were proved with the choice of constants from [15]. Of course, the method extends to a more general setting as well, with the same proofs. More precisely, if we let $m = n^c$ (instead of n^3) and denote by S_n the n-dimensional ball of radius n^{-d} (instead of n^{-8}), one can show that with a $(n+m, n^{d-(c+5)/2-(3\gamma+\gamma_1+\gamma_2)/2}/[\sqrt{\pi}(1+2n^{-\gamma-\gamma_2})])$-CVP-oracle, Theorem 10 remains valid. Theorem 11 also remains valid with a constant σ in $]0\,;2+2(2d-(9+c))/5[$ if θ and γ are such that $\frac{5\gamma}{2}+2\theta < d-(9+c)/2$ and we use a $(n^{2+\gamma}, n^\theta)$-SVP-oracle. In particular, the CVP-reduction implies that breaking the Ajtai-Dwork cryptosystem is unlikely to be NP-hard.

We have also presented a heuristic attack to recover the private key, given only the public key. It has been successfully implemented in the case of small parameters, and latest expriments suggest that the attack could be applied to real-life parameters in a reasonable time. This shows that unless major improvements are found, the Ajtai-Dwork cryptosystem is only of theoretical importance.

Acknowledgements. We would like to thank the anonymous referees for their helpful comments.

References

1. L. M. Adleman. On breaking generalized knapsack public key cryptosystems. In *Proc. 15th ACM STOC*, pages 402–412, 1983.
2. M. Ajtai. Generating hard instances of lattice problems. In *Proc. 28th ACM STOC*, pages 99–108, 1996. Available at [11] as TR96-007.
3. M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions. In *Proc. 30th ACM STOC*, 1998. Available at [11] as TR97-047.
4. M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proc. 29th ACM STOC*, pages 284–293, 1997. Available at [11] as TR96-065.
5. S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, 1997.
6. L. Babai. On Lovász lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–13, 1986.
7. E. Brickell. Breaking iterated knapsacks. In *Proc. CRYPTO'84*, volume 196 of *LNCS*, pages 342–358, 1985.
8. J.-Y. Cai and A. P. Nerurkar. An improved worst-case to average-case connection for lattice problems. In *Proc. 38th IEEE FOCS*, pages 468–477, 1997.
9. D. Coppersmith. Small solutions to polynomial equations, and low exponent RSA vulnerabilities. *J. of Cryptology*, 10(4):233–260, 1997.
10. M.J. Coster, A. Joux, B.A. LaMacchia, A.M. Odlyzko, C.-P. Schnorr, and J. Stern. Improved low-density subset sum algorithms. *Computational Complexity*, 2:111–128, 1992.
11. ECCC. `http://www.eccc.uni-trier.de/eccc/`. The Electronic Colloquium on Computational Complexity.

12. P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical report, Mathematische Instituut, University of Amsterdam, 1981. Report 81-04.
13. O. Goldreich. *Foundations of Cryptography (Fragments of a Book)*. Weizmann Institute of Science, 1995. Available at [11].
14. O. Goldreich and S. Goldwasser. On the limits of non-approximability of lattice problems. In *Proc. 30th ACM STOC*, 1998. Available at [11] as TR97-031.
15. O. Goldreich, S. Goldwasser, and S. Halevi. Eliminating decryption errors in the Ajtai-Dwork cryptosystem. In *Proc. of Crypto'97*, volume 1294 of *LNCS*, pages 105–111. Springer-Verlag, 1997. Available at [11] as TR97-018.
16. A. Joux and J. Stern. Lattice reduction: a toolbox for the cryptanalyst. (to appear in J. of Cryptology).
17. J.C. Lagarias and A.M. Odlyzko. Solving low-density subset sum problems. In *Proc. 24th IEEE FOCS*, pages 1–10. IEEE, 1983.
18. A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261:515–534, 1982.
19. P. Nguyen and J. Stern. Merkle-Hellman revisited: a cryptanalysis of the Qu-Vanstone cryptosystem based on group factorizations. In *Proc. of Crypto'97*, volume 1294 of *LNCS*, pages 198–212. Springer-Verlag, 1997.
20. P. Nguyen and J. Stern. A converse to the Ajtai-Dwork security proof and its cryptographic implications. Technical Report TR98-010, ECCC, 1998. Revision available at [11].
21. C.-P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theoretical Computer Science*, 53:201–224, 1987.
22. A. Shamir. A polynomial time algorithm for breaking the basic Merkle-Hellman cryptosystem. In *Proc. 23rd IEEE FOCS*, pages 145–152, 1982.
23. V. Shoup. Number Theory C++ Library (NTL) version 2.0. Can be obtained at `http://www.cs.wisc.edu/~shoup/ntl/`.
24. J. Stern. Secret linear congruential generators are not cryptographically secure. In *Proc. 28th IEEE FOCS*, pages 421–426, 1987.

A Appendix

A.1 Proof of Theorem 12

The proof is similar to the one of Theorem 4. Let u be a private key. For (4), we need to generalize Lemma 5 and 6. Let δ be a randomly chosen point from S_n:

$$E[\langle u, \delta \rangle^{2k}] \leq \frac{4W_n}{n^{16}} \int_0^1 (1-y^2)^{(n-1)/2} y^{2k} dy.$$

This integral is equal to $I(n,k) = \int_0^{\pi/2} \sin^n \theta \cos^{2k} \theta d\theta$. We have $I(n,0) = W_n$ and an integration by parts shows that: $I(n,k) = \frac{2k-1}{n+1} I(n+2, k-1)$. This implies $I(n,k) \leq W_n (2k)!/n^k$. Hence:

$$E[\langle u, \delta \rangle^{2k}] \leq \frac{4}{n^{16}} \times \frac{(2k)!}{n^k} W_n^2 \leq \frac{2\pi(2k)!}{n^{17+k}}.$$

The expectation would be equal to zero if there was an odd power instead of $2k$. Now, let $v = a + \sum_i \delta_i$ be a randomly chosen point from the distribution $\mathcal{H}_u$. We have:

$$E\left[\operatorname{dist}(\mathbf{Z}, \langle u, v\rangle)^{2k}\right] \leq E\left[\left(\sum_{i=1}^{n} \langle u, \delta_i\rangle\right)^{2k}\right].$$

If we expand this product, we obtain a sum of m^{2k} terms. But all the terms for which some $\langle u, \delta_i\rangle$ has an odd exponent disappear. By the multinomial formula and the independence of the δ_j's, this expectation is therefore equal to:

$$\sum_{i_1+\cdots+i_n=k} \frac{(2k)!}{(2i_1)!\cdots(2i_n)!} \prod_{j=1}^{n} E\left[\langle u, \delta_j\rangle^{2i_j}\right].$$

We know that each product is less than $\displaystyle\prod_{i_j>0} \frac{2\pi(2i_j)!}{n^{17+i_j}} \leq \frac{(2\pi(2k)!)^{2k}}{n^{17k+k}}$. And:

$$\sum_{i_1+\cdots+i_n=k} \frac{(2k)!}{(2i_1)!\cdots(2i_n)!} \leq \frac{(2k)!}{k!} \sum_{i_1+\cdots+i_n=k} \frac{k!}{i_1!\cdots i_n!} = \frac{(2k)!}{k!} n^k.$$

Thus:

$$E\left[\operatorname{dist}(\mathbf{Z}, \langle u, v\rangle)^{2k}\right] \leq \frac{(2k)!}{k!} n^k \times \frac{(2\pi(2k)!)^{2k}}{n^{17k+k}} \leq \frac{1}{n^{17k}} 4^k \pi^{2k} (2k)!^{2k+1}.$$

Therefore:

$$\begin{aligned} E\left[\left(\sum_{j=1}^{n} \operatorname{dist}(\mathbf{Z}, \langle u, w_j\rangle)^2\right)^k\right] &\leq \sum_{j_1+\cdots+j_n=k} \frac{k!}{j_1!\cdots j_n!} \prod_{\ell=1}^{n} E\left[\operatorname{dist}(\mathbf{Z}, \langle u, w_\ell\rangle)^{2j_\ell}\right] \\ &\leq \sum_{j_1+\cdots+j_n=k} \frac{k!}{j_1!\cdots j_n!} \frac{1}{n^{17k}} \left(4^k \pi^{2k} (2k)!^{2k+1}\right)^k \\ &\leq \frac{1}{n^{16k}} \left(4^k \pi^{2k} (2k)!^{2k+1}\right)^k \end{aligned}$$

Thus, by the moment inequality, (4) is satisfied with probability at least $1 - \varepsilon_1$ with respect to the choice of $w_1, \ldots, w_n$, if we let $M_1 = 4^k \pi^{2k} (2k)!^{2k+1}$.

For (5), as in the proof of (2), we bound the expectation when the b_i's are fixed. A first bound is obtained when the v_i's are independent random variables uniformly distributed over the hypercube B_n. Then, we show that with the actual distribution of the v_i's, the additional error is negligible, so that the bound of (5) is satisfied for sufficiently large n, thanks to Markov's inequality.

For the first bound, we generalize Lemma 7 with the same tricks we used to generalize Lemma 5. Let $t = (t_1, \ldots, t_n)$ be a vector in the n-dimensional unit sphere. Let $s = (s_1, \ldots, s_n)$ be a randomly chosen point with uniform distribution from B_n. We have:

$$E[\langle s,t\rangle^{2k}] = E\left[\left(\sum_{j=1}^{n} s_j t_j\right)^{2k}\right].$$

If we expand this product, we obtain m^{2k} terms. But all the terms for which some s_j has an odd exponent disappear. We obtain by the multinomial formula:

$$E[\langle s,t\rangle^{2k}] = \sum_{i_1+\cdots+i_n=k} \frac{(2k)!}{(2i_1)!\cdots(2i_n)!} E\left[(s_1t_1)^{2i_1}\cdots(s_nt_n)^{2i_n}\right].$$

And since the s_j's are independent:

$$E\left[(s_1t_1)^{2i_1}\cdots(s_nt_n)^{2i_n}\right] = t_1^{2i_1}\cdots t_n^{2i_n}\rho_n^{2i_1+\cdots+2i_n}\frac{1}{2i_1+1}\cdots\frac{1}{2i_n+1}.$$

Therefore:

$$E[\langle s,t\rangle^{2k}] = \rho_n^{2k}\sum_{i_1+\cdots+i_n=k}\frac{(2k)!}{(2i_1+1)!\cdots(2i_n+1)!}t_1^{2i_1}\cdots t_n^{2i_n}.$$

And this sum is less than:

$$\frac{(2k)!}{k!}\sum_{i_1+\cdots+i_n=k}\frac{k!}{i_1!\cdots i_n!}t_1^{2i_1}\cdots t_n^{2i_n} = \frac{(2k)!}{k!}(t_1^2+\cdots+t_n^2)^k = \frac{(2k)!}{k!}.$$

Thus:

$$E[\langle s,t\rangle^{2k}] \leq \frac{(2k)!}{k!}\rho_n^{2k}.$$

And this expectation would be equal to zero if there was an odd power instead of $2k$. Therefore, if we assume that the v_i's are distributed uniformly over B_n:

$$E\left[\left\langle \sum_{i=1}^{m} b_i v_i, w_j^{\perp}\right\rangle^{2k}\right] \leq \sum_{i_1+\cdots+i_m=k}\frac{(2k)!}{(2i_1)!\cdots(2i_m)!}\prod_{\ell=1}^{m}E\left[\langle v_\ell, w_j^{\perp}\rangle^{2i_\ell}\right].$$

We know that each product is less than $\displaystyle\prod_{i_\ell>0}\frac{(2i_\ell)!}{i_\ell!}\rho_n^{2i_\ell} \leq \rho_n^{2k}(2k)!^k$. And:

$$\sum_{i_1+\cdots+i_m=k}\frac{(2k)!}{(2i_1)!\cdots(2i_m)!} \leq \frac{(2k)!}{k!}\sum_{i_1+\cdots+i_m=k}\frac{k!}{i_1!\cdots i_m!} = \frac{(2k)!}{k!}m^k.$$

It follows that:

$$E\left[\left\langle \sum_{i=1}^{m} b_i v_i, w_j^{\perp}\right\rangle^{2k}\right] \leq \rho_n^{2k}(2k)!^k\frac{(2k)!}{k!}m^k = \rho_n^{2k}\frac{(2k)!^{k+1}}{k!}m^k.$$

Therefore, if we denote by X the random variable $(\sum_{j=1}^{n}\langle\sum_{i=1}^{m} b_i v_i, w_j^{\perp}\rangle^2)^k$, the multinomial formula shows that:

$$\begin{aligned}
E[X] &\leq \sum_{j_1+\cdots+j_n=k} \frac{k!}{j_1!\cdots j_n!} \prod_{\ell=1}^{n} E\left[\left\langle \sum_{i=1}^{m} b_i v_i, w_\ell^{\perp} \right\rangle^{2j_\ell}\right] \\
&\leq \sum_{j_1+\cdots+j_n=k} \frac{k!}{j_1!\cdots j_n!} \prod_{j_\ell>0} \rho_n^{2j_\ell} \frac{(2j_\ell)!^{j_\ell+1}}{j_\ell!} m^{j_\ell} \\
&\leq \sum_{j_1+\cdots+j_n=k} \frac{k!}{j_1!\cdots j_n!} \rho_n^{2k} (2k)!^{k+1} m^k \\
&\leq \rho_n^{2k} (2k)!^{k+1} m^k \times \frac{(2k)!}{k!} n^k \\
&\leq n^{4k} \rho_n^{2k} \frac{(2k)!^{k+2}}{k!}
\end{aligned}$$

With the actual distribution of the v_i's, there is an additional term which is negligible, so that the wanted bound is satisfied for sufficiently large n, with for instance: $M_4 = (2k+1)!^{k+2}/k!$.

A.2 Proof of Proposition 15

We first need a combinatorial lemma:

Lemma 17. *For sufficiently large n, for all elements $y_1, \ldots, y_\nu$ in the parallelepiped $P(w_1, \ldots, w_n)$, there exist coefficients λ_i (not all zero) in $\{-1, 0, +1\}$ such that:*

$$\left\| \sum_{i=1}^{\nu} \lambda_i y_i \right\| \leq \frac{1}{n^6}.$$

Proof (Sketch). Same reasoning as in the proof of Theorem 2. □

Lemma 18. *Let $\lambda_1, \ldots, \lambda_\nu$ be integers not all zero. If $y_1, \ldots, y_\nu$ are chosen at random in the parallelepiped $P(w_1, \ldots, w_n)$ then:*

$$Pr\left[\left\| \sum_{i=1}^{\nu} \lambda_i y_i \right\| \leq \frac{1}{2n^2}\right] \leq \frac{1}{\rho_n^n}.$$

Proof. Assume that the inequality on the norm is satisfied. Write $\sum_{i=1}^{\nu} \lambda_i y_i$ as $\sum_{j=1}^{n} \alpha_j w_j$. We have: $|\alpha_j| \leq \| \sum_{i=1}^{\nu} \lambda_i y_i \| \times n^2/\rho_n \leq 1/(2\rho_n)$. The probability is therefore bounded by the probability that each α_j is between $-\frac{1}{2\rho_n}$ and $\frac{1}{2\rho_n}$.

Each y_i is of the form $\sum_{\ell=1}^{n} \mu_{i,\ell} w_\ell$ where the $\mu_{i,\ell}$'s are independently chosen in $[0,1[$ with uniform distribution. It follows that: $\alpha_j = \sum_{i=1}^{\nu} \lambda_i \mu_{i,j}$. If λ_i is non-zero, then $\lambda_i \mu_{i,j}$ modulo 1 is uniformly distributed over $[0,1[$. Since the λ_i's are not all zero, α_j modulo 1 is therefore uniformly distributed over $[0,1[$. Furthermore, the α_j's are independent, and the result follows. □

This probabilistic lemma is the core of the following result:

Lemma 19. *Let $\tau = \gamma_1 - \gamma$. If $y_1, \ldots, y_\nu$ are chosen at random in $P(w_1, \ldots, w_n)$ then the probability that there exist $\lambda_1, \ldots, \lambda_\nu$ not all zero such that*

$$\left\| \sum_{i=1}^{\nu} \lambda_i y_i \right\| \leq \sqrt{2} \frac{1}{n^{6-\theta}} \tag{7}$$

$$\|(\lambda_1, \ldots, \lambda_\nu)\| \leq \sqrt{2} n^{1+\gamma/2+\theta} \tag{8}$$

$$|\{i : \lambda_i \neq 0\}| \leq n^{2-\tau} \tag{9}$$

is exponentially small (with respect to n).

Proof. The number of non-zero $(\lambda_1, \ldots, \lambda_\nu)$ satisfying (8) and (9) is at most

$$\binom{n^{2+\gamma}}{n^{2-\tau}} (2n^{1+\gamma/2+\theta})^{n^{2-\tau}} \leq (n^{2+\gamma})^{n^{2-\tau}} (2n^{1+\gamma/2+\gamma})^{n^{2-\tau}}.$$

Since $\theta < 3$, by Lemma 18, each vector has probability less than ρ_n^{-n} to satisfy (7). This yields an overall probability less than $(n^{2+\gamma})^{n^{2-\tau}} (2n^{1+\gamma/2+\theta})^{n^{2-\tau}} \rho_n^{-n}$. Taking logarithms we get:

$$n^{2-\tau} \left[(2+\gamma) \log_2 n + \left(1 + \frac{\gamma}{2} + \theta\right) \log_2 n + 1 \right] - n^2 \log_2 n.$$

Since $2 - \tau < 2$, the leading term is $-n^2 \log_2 n$ and the result follows. □

Now, consider the output $(z, \lambda_1, \ldots, \lambda_\nu)$ of the oracle. By Lemma 17 and by definition of the oracle, $\|z\|^2$ and $\sum_{i=1}^{\nu} \lambda_i^2$ are both less than:

$$n^{2\theta} \left(\beta^2 \frac{1}{n^{12}} + \nu \right) \leq n^{2\theta} \left(n^{2+\gamma} + n^{2+\gamma} \right) = 2n^{2+\gamma+2\theta}.$$

Therefore:

$$\left\| \lambda_1 v_{p(1)} + \cdots + \lambda_\nu v_{p(\nu)} \right\| \leq \sqrt{2} n^{6-\theta} \text{ and } \|(\lambda_1, \ldots, \lambda_\nu)\| \leq \sqrt{2} n^{1+\gamma/2+\theta}.$$

This means that (7) and (8) are satisfied if we use the $y_{p(i)}$'s instead of the y_i's. Since the λ_i's are not all zero and $y_1, \ldots, y_\nu$ are ciphertexts of '1', Lemma 19 implies that with overwhelming probability, (9) is not satisfied: at least $n^{2-\tau}$ coefficients are non zero. By symmetry, the probability that y_i is hit does not depend on i. Furthermore, (8) implies that the number x of $(\lambda_1, \ldots, \lambda_\nu)$'s such that $|\lambda_i| \geq n^{\gamma/2+\theta+\lambda}$ is such that:

$$x n^{\gamma+2\theta+2\lambda} \leq \|(\lambda_1, \ldots, \lambda_\nu)\|^2 \leq 2n^{2+\gamma+2\theta}.$$

Hence:

$$x \leq 2n^{2-2\lambda}.$$

Since $\lambda > \tau$ (because $\gamma_1 < \gamma + \lambda$), this number is negligible with respect to $n^{2-\tau}$. Now, the probability that λ_i is hit is:

$$\Omega\left(\frac{n^{2-\tau}}{n^{2+\gamma}}\right) = \Omega\left(\frac{1}{n^{\gamma+\tau}}\right) = \Omega\left(\frac{1}{n^{\gamma_1}}\right).$$

A.3 Proof of Proposition 16

As in the proof of Proposition 15, consider the output $(z, \lambda_1, \ldots, \lambda_\nu)$ of the oracle. $\|z\|$ and $\|(\lambda_1, \ldots, \lambda_\nu)\|$ are still less than $\sqrt{2}n^{1+\theta+\gamma/2}$. And we have:

$$\lambda_{p^{-1}(1)} y_1 = \frac{1}{\beta} z - \sum_{i=2}^{\nu} \lambda_{p^{-1}(i)} y_i.$$

Since $y_2, \ldots, y_\nu$ are good ciphertexts of '0', Lemma 14 implies that for all $i \geq 2$:

$$\operatorname{dist}(\mathbf{Z}, \langle u, y_i \rangle) \leq M_4 \frac{1}{n^4 (\varepsilon_1 \varepsilon_2 \varepsilon_3)^{1/2k}}.$$

Therefore, by the Cauchy-Schwarz inequality:

$$\begin{aligned}
\operatorname{dist}\left(\mathbf{Z}, \left\langle \sum_{i=2}^{\nu} \lambda_{p^{-1}(i)} y_i, u \right\rangle\right) &\leq \sqrt{\sum_{i=1}^{\nu} \lambda_{p^{-1}(i)}^2} \times \sqrt{\nu M_4^2 \frac{1}{n^8 (\varepsilon_1 \varepsilon_2 \varepsilon_3)^{1/k}}} \\
&\leq \sqrt{2} n^{1+\theta+\gamma/2} M_4 n^{1+\gamma/2-4} \frac{1}{(\varepsilon_1 \varepsilon_2 \varepsilon_3)^{1/2k}} \\
&\leq M_4 \frac{\sqrt{2}}{(\varepsilon_1 \varepsilon_2 \varepsilon_3)^{1/2k}} n^{\theta+\gamma-2}.
\end{aligned}$$

Furthermore:

$$\operatorname{dist}(\mathbf{Z}, \langle z/\beta, u \rangle) \leq \sqrt{2} n^{\theta-6}.$$

Therefore, for sufficiently large n:

$$\operatorname{dist}(\mathbf{Z}, \langle \lambda_{p^{-1}(1)} y_1, u \rangle) \leq M_4 \frac{\sqrt{3}}{(\varepsilon_1 \varepsilon_2 \varepsilon_3)^{1/2k}} n^{\theta+\gamma-2}.$$

If $\lambda_{p^{-1}(1)}$ is a fixed integer, since y_1 is a random vector in the parallelepiped, the latter inequality is satisfied with probability at most:

$$2M_4 \frac{\sqrt{3}}{(\varepsilon_1 \varepsilon_2 \varepsilon_3)^{1/2k}} n^{\theta+\gamma-2}.$$

But if y_1 is hit, then:

$$|\lambda_{p^{-1}(1)}| \in \left\{1, 2, \ldots, n^{\frac{\gamma}{2}+\theta+\lambda}\right\}.$$

Hence, y_1 is hit with probability at most:

$$2M_4 \frac{2\sqrt{3}}{(\varepsilon_1 \varepsilon_2 \varepsilon_3)^{1/2k}} n^{\theta+\gamma-2} 2n^{\gamma/2+\theta+\lambda}.$$

As n grows, this is:

$$\mathcal{O}\left(n^{2\theta+3\gamma/2+\lambda-2+(\sigma_1+\sigma_2+\sigma_3)/(2k)}\right) = \mathcal{O}\left(\frac{1}{n^{\gamma_2}}\right).$$

And this concludes the proof.