A Partial Key Exposure Attack on RSA Using a 2-Dimensional Lattice

Ellen Jochemsz* and Benne de Weger*

Department of Mathematics and Computer Science,
Eindhoven University of Technology, Eindhoven, The Netherlands
{e.jochemsz, b.m.m.d.weger}@tue.nl

Abstract. We describe an attack on the RSA cryptosystem when the private exponent d is chosen to be 'small', under the condition that a sufficient amount of bits of d is available to the attacker. The attack uses a 2-dimensional lattice and is therefore (in the area of the keyspace where it applies) more efficient than known attacks using Coppersmith techniques. Moreover, we show that the attacks of Wiener and Verheul/Van Tilborg, using continued fractions techniques, are special deterministic cases of our attack, which in general is heuristic.

Keywords: RSA, cryptanalysis, partial key exposure, lattice basis reduction, inhomogeneous diophantine approximation.

1 Introduction

Since the introduction of the RSA cryptosystem in 1977, people have been looking for its vulnerabilities. A summary of attacks on RSA up to 1999 was given by Boneh in [2]. Although none of these attacks totally break RSA, they provide a certain guideline for the use of RSA and show in which cases the cryptosystem is unsafe.

For instance, it is known that using a small private exponent d can be dangerous. In 1990, Wiener showed in [13] that if the size of d is less than $\frac{1}{4}$th of the size of the modulus N, it can be found by continued fractions methods. Verheul and Van Tilborg [11] generalized this result in 1997, to obtain an attack based on continued fractions that works if d is slightly larger than $N^{\frac{1}{4}}$. In 2000, Boneh and Durfee [3] extended Wiener's bound to $d < N^{0.292}$.

The concept of partial key exposure attacks on RSA was introduced in 1997 by Boneh, Durfee and Frankel in [4], and deals with the situation where an attacker has obtained some bits of the private exponent d. The main question is: How much information on the bits of d is needed such that an attacker can reconstruct d, thereby breaking the RSA instance?

* The work described in this paper has been supported in part by the European Commission through the IST Programme under Contract IST-2002-507932 ECRYPT. The information in this document reflects only the author's views, is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

S.K. Katsikas et al. (Eds.): ISC 2006, LNCS 4176, pp. 203–216, 2006.

The motivation for exploring partial key exposure attacks comes from side-channel attacks such as power analysis, timing attacks, etc. Using a side-channel, an attacker can expose a part of d, generally an MSB (most significant bit) part or LSB (least significant bit) part.

In all the subsequent papers about partial key exposure attacks, the assumption is made (besides knowledge of MSBs/LSBs of d) that one of the exponents e, d is chosen to be small (at least significantly smaller than the modulus N). This is a common practice, since a small exponent yields faster modular exponentiation. For instance, $e = 2^{16} + 1 = 65537$ is a popular choice, and for signing operations on constrained devices such as smartcards, it is useful to have the private (signing) exponent d to be small, though obviously larger than $N^{0.292}$.

The first partial key exposure attacks by Boneh, Durfee, and Frankel [4] required the public exponent e to be smaller than $N^{\frac{1}{2}}$. Blömer and May extended their result in [5] with attacks for $e \in [N^{0.5}, N^{0.725}]$. Ernst, Jochemsz, May, and De Weger [7] recently showed attacks for both the situations where the private exponent d or the public exponent e is chosen to be small. Both their attacks work up to full size exponents.

In the papers [3,5,7], lattice methods are used instead of continued fractions methods. Generally, one starts by describing an RSA situation in terms of an integer polynomial that has a small (unknown) root, or a polynomial that has a small (unknown) root modulo a known constant. After that, one uses the theory initiated by Coppersmith [6], to construct a lattice with polynomials with the same root, and reduce the lattice to obtain a polynomial, again having the same root, whose coefficients are small enough to find the root.

These attacks using lattice methods are asymptotic, meaning that if one comes close to the maximal value for the unknown part of d for which an attack should work, the lattices involved are very large. This implies that the lattice reduction phase, for which the LLL-algorithm [8] is used, may take a prohibitively long time.

Therefore, it may be useful to look at very small lattices instead of very large. In this paper, we explore for which sizes of d, one can mount an attack in a few seconds with a very simple method using a 2-dimensional lattice. Our result is summarized in the following theorem.

Theorem 1. *Under a reasonable heuristic assumption that we specify in Assumption 1, the following holds: Let $N = pq$ be an n-bit RSA-modulus, and p, q primes of bitsize $\frac{n}{2}$. Let $0 < \beta < \frac{1}{2}$, and let e, d satisfy $ed \equiv 1 \mod \phi(N)$ with bitsize(e) $= n$ and bitsize(d) $= \beta n$. Given a (total) amount of $(2\beta - \frac{1}{2})n$ MSBs and/or LSBs of d (see Figure 1), N can be factored very efficiently, using a 2-dimensional lattice.*

We will comment on what 'very efficiently' means in Section 5, when we compare the performance of this attack on small d to the method of Ernst et al. [7]. Moreover, we show that the results of Wiener and Verheul/Van Tilborg can be obtained by our attack on small d and are simply special (homogeneous and deterministic) cases. One could also say that our partial key exposure attack

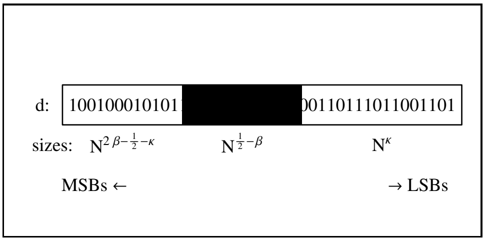

Fig. 1. Partition of d for small d

is the inhomogeneous counterpart of the results by Wiener and Verheul/Van Tilborg. We will comment on this, and on the heuristic assumption in the other cases, in Section 4.

The rest of this paper is organized as follows. In Section 2, we will state preliminaries on RSA and on lattice techniques, define some notation and introduce Assumption 1. Section 3 will contain the description of the attack for small d. In Section 4, we will comment on the cases where our attack does not depend on Assumption 1 and is therefore deterministic, and at the experimental results for the heuristic in the cases where we do need Assumption 1. In Section 5, we look at the efficiency of our method and compare our 2-dimensional attack with the existing partial key exposure attacks on small d of [7]. Finally, we will give a conclusion in Section 6.

2 Preliminaries on RSA and Lattices

In this section, we state some basic properties of RSA, the cryptosystem we are attacking and of 2-dimensional lattices, the tool we use to do so.

Let p, q, N, d, e be as usual, i.e. p and q are distinct primes, $N = pq$ is taken as modulus, and the encryption exponent e and decryption exponent d satisfy $ed \equiv 1 \pmod{\phi(N)}$. For the attack in this paper, we assume that p and q have the same bitsize, thus $p + q < 3N^{\frac{1}{2}}$. Let $k \in \mathbb{Z}$ be defined by the RSA key equation

$$ed - 1 = k\phi(N), \quad \text{where } \phi(N) = (p-1)(q-1) = N - (p + q - 1).$$

In our attack in this paper, we assume that the private exponent d is chosen to be small, for efficient modular computations. From the RSA key equation, it follows directly that $k < d$.

We define a 2-dimensional lattice L as the set of all integer linear combinations of two linearly independent vectors $\{\mathbf{b_1}, \mathbf{b_2}\}$, which are basis vectors. We usually say that L is the lattice spanned by the columns of the matrix $\Gamma = (\mathbf{b_1}\ \mathbf{b_2})$. The determinant of L is $\det(L) = |\det(\Gamma)|$, and though there are infinitely many bases possible, the determinant is always the same.

To find a small, so-called reduced basis $\{\mathbf{r}, \mathbf{s}\}$, one can use a reduction algorithm. For a 2-dimensional lattice, the Lagrange reduction algorithm (which is simply a generalization of Euclid's algorithm) finds a reduced basis, and this basis also contains the smallest nonzero vector of the lattice. We are interested in how small the reduced basis vectors are in norm.

We use the following notation for size-computations in this paper. With $u \approx N^{\lambda}$, we mean that u 'has the size of' N^{λ}, that is $|u| = C_u N^{\lambda}$ for some number C_u that does not deviate much from 1. Naturally, $\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} \approx \begin{pmatrix} N^{\lambda_1} \\ N^{\lambda_2} \end{pmatrix}$ is a short notation for $v_1 \approx N^{\lambda_1}$ and $v_2 \approx N^{\lambda_2}$.

When we reduce the matrix Γ to $\Gamma_{\text{red}} = (\mathbf{r}\ \mathbf{s})$, with $\mathbf{r}$ the smaller reduced basis vector and $\mathbf{s}$ the larger reduced basis vector, it holds that $||\mathbf{r}|| \cdot ||\mathbf{s}|| \approx \det(L)$. So, we assume $||\mathbf{r}|| \approx a^{-1} \det(L)^{\frac{1}{2}}$ and $||\mathbf{s}|| \approx a \det(L)^{\frac{1}{2}}$ for some $a \geq 1$.

Hence,

$$\Gamma_{\text{red}} = (\mathbf{r}\ \mathbf{s}) = \begin{pmatrix} r_1 & s_1 \\ r_2 & s_2 \end{pmatrix}, \text{ and } \Gamma_{\text{red}}^{-1} = \tfrac{1}{\det(\Gamma)} \begin{pmatrix} s_2 & -s_1 \\ -r_2 & r_1 \end{pmatrix} = \begin{pmatrix} \mathbf{s'^T} \\ \mathbf{r'^T} \end{pmatrix}.$$

It follows that the first row $\mathbf{s}'$ of Γ_{red} satisfies $||\mathbf{s}'|| \approx a \det(L)^{\frac{1}{2}}$. Analogously, $||\mathbf{r}'|| \approx a^{-1} \det(L)^{\frac{1}{2}}$.

If the two reduced basis vectors $\mathbf{r}$, $\mathbf{s}$ are 'nearly-equal' in length, that is when a does not deviate much from 1, then $||\mathbf{r}|| \approx ||\mathbf{s}|| \approx \det(L)^{\frac{1}{2}}$. In other words, all reduced basis vectors of L have a norm of size $\det(L)^{\frac{1}{2}}$. However, it is also possible that there is one 'extremely small' basis vector, which makes the lattice 'unbalanced'. For the attacks in this paper, we make the following assumption.

Assumption 1. *The reduced basis vectors given by the columns of Γ_{red} both have a norm of size* $\det(L)^{\frac{1}{2}}$ *. In other words, the parameter a used to describe the unbalancedness of the lattice is near to* 1.

In Section 4, we comment on how this assumption holds in practice.

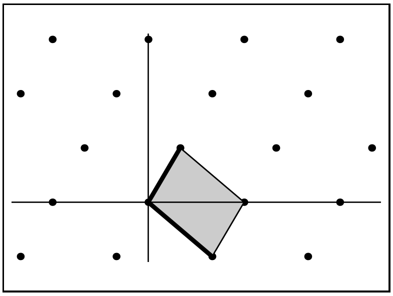

Fig. 2. $a \approx 1$

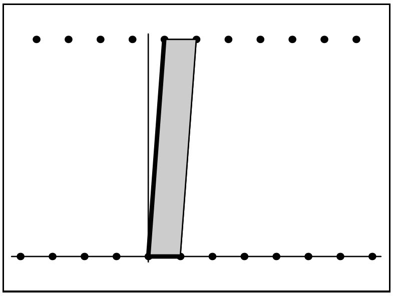

Fig. 3. $a \gg 1$

Having discussed the necessary preliminaries, we are now ready to explain the 2-dimensional partial key exposure attack on RSA for small d.

3 The Attack on Small d

3.1 Description of the Attack

Let $d = N^{\beta} < N^{\frac{1}{2}}$ and $e < \phi(N) < N$. In this section, we will prove the statement in Theorem 1, namely that we can factor N very efficiently if we know a (total) amount of $(2\beta - \frac{1}{2})n$ MSBs and/or LSBs of d.

This implies that our method will work if the 'unknown middle part' of d is of size N^{δ} with $\delta < \beta - (2\beta - \frac{1}{2}) = \frac{1}{2} - \beta$. The situation is sketched in Figure 4.

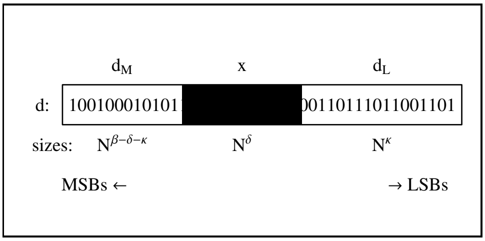

Fig. 4. Partition of d when MSBs and/or LSBs are known

Let d_L be the known LSB part of d of size N^{κ}, followed by an unknown middle part x of size N^{δ}, which itself is followed by a known MSB part d_M, of size $N^{\beta-\kappa-\delta}$. Hence, we can write

$$d = d_L + 2^{\lfloor \kappa n \rceil} x + 2^{\lfloor \kappa n \rceil + \lfloor \delta n \rceil} d_M,$$

where $\lfloor \ \rceil$ is simply rounding to the nearest integer.

When we substitute the partition of d in the RSA key equation, we obtain

$$e 2^{\lfloor \kappa n \rceil} x + e d_L + e 2^{\lfloor \kappa n \rceil + \lfloor \delta n \rceil} d_M - 1 = k(N - (p + q - 1)).$$

Therefore, we must find the solution $(x, y, z) = (x, k, p + q - 1)$ of the trivariate equation

$$e 2^{\lfloor \kappa n \rceil} x - Ny + yz + R - 1 = 0, \text{ with } R = e d_L + e 2^{\lfloor \kappa n \rceil + \lfloor \delta n \rceil} d_M.$$

The equation above implies that

$$|e 2^{\lfloor \kappa n \rceil} x - Ny + R| = |1 - yz| \leq |k(p + q - 1)| \leq |d(p + q)| \leq 3 N^{\beta + \frac{1}{2}}.$$

This is an inhomogeneous diophantine approximation problem in the unknowns x and y. To solve it, we define a lattice L spanned by the columns of Γ, with

$$\Gamma = \begin{pmatrix} C & 0 \\ e 2^{\lfloor \kappa n \rceil} & N \end{pmatrix}, \quad \text{and } \mathbf{v} = \begin{pmatrix} 0 \\ -R \end{pmatrix},$$

where C is a convenient integer of size $N^{\beta - \delta + \frac{1}{2}}$.

The lattice point $\Gamma \begin{pmatrix} x \\ -y \end{pmatrix}$ is close to $\mathbf{v}$, since

$$\Gamma \begin{pmatrix} x \\ -y \end{pmatrix} - \mathbf{v} = \begin{pmatrix} Cx \\ e2^{\lfloor \kappa n \rceil} x - Ny + R \end{pmatrix} \approx \begin{pmatrix} N^{\beta+\frac{1}{2}} \\ N^{\beta+\frac{1}{2}} \end{pmatrix}.$$

Our strategy to find x and y is therefore to start with a lattice vector $\mathbf{v}'$ close to $\mathbf{v}$, and add small multiples of the reduced basis vectors of the lattice L until we get $\Gamma \begin{pmatrix} x \\ -y \end{pmatrix}$. To do so, we apply lattice basis reduction to the columns of Γ, and obtain a reduced matrix Γ_{red}, whose columns still span L. We aim to find an integer pair (z_1, z_2) for which

$$\Gamma_{\text{red}} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \Gamma \begin{pmatrix} x \\ -y \end{pmatrix} - \Gamma_{\text{red}} \left\lfloor \Gamma_{\text{red}}^{-1} \mathbf{v} \right\rceil,$$

where $\left\lfloor \Gamma_{\text{red}}^{-1} \mathbf{v} \right\rceil = \mathbf{v}'$ is the vector we get from rounding the elements of $\Gamma_{\text{red}}^{-1} \mathbf{v}$ to nearest integers. Alternatively, one could also solve the closest vector problem to obtain a lattice vector $\mathbf{v}'$ to start with, but in practice the closest vector will almost immediately appear in this way as well.

It can be checked that

$$\Gamma_{\text{red}} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \left(\Gamma \begin{pmatrix} x \\ -y \end{pmatrix} - \mathbf{v}\right) - \left(\Gamma_{\text{red}} \left\lfloor \Gamma_{\text{red}}^{-1} v \right\rceil - \mathbf{v}\right) \approx \begin{pmatrix} N^{\beta+\frac{1}{2}} \\ N^{\beta+\frac{1}{2}} \end{pmatrix} + \Gamma_{\text{red}} \begin{pmatrix} \epsilon_1 \\ \epsilon_2 \end{pmatrix},$$

with $|\epsilon_i| < \frac{1}{2}$. Therefore

$$\begin{aligned} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} &\approx \Gamma_{\text{red}}^{-1} \begin{pmatrix} N^{\beta+\frac{1}{2}} \\ N^{\beta+\frac{1}{2}} \end{pmatrix} + \begin{pmatrix} \epsilon_1 \\ \epsilon_2 \end{pmatrix} = \begin{pmatrix} \mathbf{s'^T} \\ \mathbf{r'^T} \end{pmatrix} \begin{pmatrix} N^{\beta+\frac{1}{2}} \\ N^{\beta+\frac{1}{2}} \end{pmatrix} + \begin{pmatrix} \epsilon_1 \\ \epsilon_2 \end{pmatrix} \\ &\lessapprox \begin{pmatrix} a \det(L)^{-\frac{1}{2}} N^{\beta+\frac{1}{2}} + \epsilon_1 \\ a^{-1} \det(L)^{-\frac{1}{2}} N^{\beta+\frac{1}{2}} + \epsilon_2 \end{pmatrix} \approx \begin{pmatrix} a N^{\frac{1}{2}(\beta+\delta-\frac{1}{2})} + \epsilon_1 \\ a^{-1} N^{\frac{1}{2}(\beta+\delta-\frac{1}{2})} + \epsilon_2 \end{pmatrix}. \end{aligned}$$

Each pair (z_1, z_2) leads to a pair $(x, -y)$. If we substitute x as the unknown part of d, and y as k, we can find a ϕ that satisfies $ed - 1 = k\phi$. First we test whether ϕ, computed as $\frac{ed-1}{k}$ is integral (unfortunately we see no way how to use this condition earlier). The next test will be to solve for the integer roots p, q of the quadratic equation $X^2 - (N + 1 - \phi)X + N = 0$.

The number of pairs (z_1, z_2) to try is of size

$$(a N^{\frac{1}{2}(\beta+\delta-\frac{1}{2})}) \cdot \max\{a^{-1} N^{\frac{1}{2}(\beta+\delta-\frac{1}{2})}, 1\}.$$

Hence, the number of pairs (z_1, z_2) to try is either

- $\mathrm{O}(N^{\beta+\delta-\frac{1}{2}})$, when $a < N^{\frac{1}{2}(\beta+\delta-\frac{1}{2})}$, or
- $\mathrm{O}(a N^{\frac{1}{2}(\beta+\delta-\frac{1}{2})})$, when $a > N^{\frac{1}{2}(\beta+\delta-\frac{1}{2})}$.

Note that in the latter case, $z_2 = 0$, but we do have to check for all z_1 separately.

In the next section, we show the relation between our method and the attacks of Wiener [13] and Verheul/Van Tilborg [11], which are special cases of this attack. For these situations, we show that the attacks are deterministic instead of heuristic, simply because the lattice vector $\Gamma \begin{pmatrix} x \\ -y \end{pmatrix}$ is small enough to ensure that the search region does not depend on a.

However, if we are outside the range of Wiener's and Verheul/Van Tilborg's attacks, it is highly unusual that the lattice involved contains an exceptionally small nonzero vector, which would make the lattice unbalanced and the attack inefficient. By Assumption 1, we take a to be close to 1. Under this heuristic, the number of pairs (z_1, z_2) to try is $\mathrm{O}(N^{\beta+\delta-\frac{1}{2}})$. In Section 4.2, we will show that this assumption is reasonable in practice.

Under our heuristic assumption, and provided that δ is smaller than or at most only marginally larger than $\frac{1}{2} - \beta$, then we can efficiently try all pairs (z_1, z_2) and find the factorization of N.

One might note that by knowing MSBs of d, one can also obtain an MSB part of k. However, splitting k into a known and an unknown part results in more combinations of variables, which we can only represent in a 3-dimensional lattice instead of a 2-dimensional one. The 3-dimensional lattice attack will give a worse analysis then the method described in this section. This is an example of a common phenomenon in lattice based cryptanalysis, namely that sometimes one can get better results by leaving out information that one knows, just by the monomials of the equation involved.

3.2 Complexity

We now study the total complexity of the above attack.

Firstly it requires one lattice basis reduction for a 2-dimensional lattice. This is just Lagrange reduction, which takes at most $\mathrm{O}((\log N)^3)$ bit operations.

Secondly, a number of $\mathrm{O}(N^{\beta+\delta-\frac{1}{2}})$ pairs (z_1, z_2) have to be checked for coming from a solution. For each vector this check takes $\mathrm{O}((\log N)^2)$ bit operations.

It follows that the bit complexity of our attack is $\mathrm{O}((\log N)^3)$ when $\delta \leq \frac{1}{2} - \beta$, which is polynomial. When $\delta = \frac{1}{2} - \beta + \epsilon$ the bit complexity becomes exponential, namely $\mathrm{O}(N^{\epsilon}(\log N)^2)$. This results in an increased workload by a factor N^{ϵ}. In other words, for an additional amount of r unknown bits, the complexity is equivalent to an exhaustive search over r bits. Furthermore, in the case that we let both d and the unknown part of d grow r bits, such that the known part of d stays of the same size, one can check that the extra workload will be an exhaustive search over $2r$ bits. This relates directly to a result of Verheul and Van Tilborg [11], on which we shall comment in Section 4.1.

3.3 Examples

We have done several experiments for this attack. A typical case is with 2048 bit N and $\delta = 0.156$, $\beta = 0.350$ (e.g. $\epsilon = 0.006$), meaning that d has about 717 bits, of which at most the 320 least significant bits are unknown.

Then $N^{\frac{1}{2}(\delta+\beta-\frac{1}{2})} \approx 70$. Indeed, we typically find a hit with $\|z\| \lessapprox 200$. A search area like this takes only a few seconds with Mathematica 5 on a 2GHz Pentium 4 PC. And with $\delta \leq \frac{1}{2} - \beta$ typically $\|z\| \approx 1$, and the computation time is only a fraction of a second.

Here's a baby example for $\{\delta = 0.156, \beta = 0.35\}$. Let the 128-bit public key be given by

$$\begin{aligned} N &= 269866491905568049204176579604167754067, \\ e &= 222981052634419442506270512141611354797. \end{aligned}$$

Now suppose we know some MSBs of d, hence we know an approximation

$$\tilde{d} = 24584250313023$$

of d for which $d_0 = d - \tilde{d}$ is $0.156 \cdot 128 \approx 20$ bits. We take

$$\begin{aligned} &C = 2^{\lfloor 128\cdot(0.35-0.156+0.5)\rceil} = 2^{89}, \text{ and} \\ &R = e\tilde{d} = 5481822013025924218218657989757723471271758362621331, \end{aligned}$$

and we know that we are looking for $\{d_0, k\}$ such that

$$\Gamma \cdot \begin{pmatrix} d_0 \\ -k \end{pmatrix} - \mathbf{v} = \begin{pmatrix} C & 0 \\ e & N \end{pmatrix} \cdot \begin{pmatrix} d_0 \\ -k \end{pmatrix} - \begin{pmatrix} 0 \\ -R \end{pmatrix}$$

is a small vector. Then Γ_{red} is given by

$$\begin{pmatrix} 93923748720621086836871453999104 & -645630915298759729739927100850176 \\ 223858603616044679201441362439981 & 239654325473299927083414831489037 \end{pmatrix}$$

and $\lfloor \Gamma_{\text{red}}^{-1} v \rceil = \begin{pmatrix} -21188034626414783992 \\ -3082348742879388262 \end{pmatrix}$.

We then enumerate the pairs $\{z_1, z_2\}$, for each value computing

$$\begin{pmatrix} x \\ -y \end{pmatrix} = \Gamma^{-1} \left(\Gamma_{\text{red}} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} + \Gamma_{\text{red}} \lfloor \Gamma_{\text{red}}^{-1} v \rceil \right).$$

We try $d = \tilde{d} + x$ and $k = y$, and solve $N + 1 - \left(p + \dfrac{N}{p}\right) = \dfrac{ed-1}{k}$ to get a possible factor p.

At $z = \begin{pmatrix} -2 \\ -1 \end{pmatrix}$ we have a hit, namely $x = 1016998$, $y = 20313089635876$, so we find that $d = 24584251330021$, and $k = 20313089635876$.

It follows that $\phi(N) = 269866491905568049171299025219693706736$, and then we obtain the factors

$$\begin{aligned} p &= 15833051453602685849, \\ q &= 17044502930871361483. \end{aligned}$$

4 The Deterministic and Heuristic Cases of the Attack

4.1 Wiener and Verheul/Van Tilborg

In [11,13], attacks were described for small d. Wiener showed that when $d < N^{\frac{1}{4}}$, it can be found in polynomial time. Verheul and Van Tilborg's extension of Wiener's result shows the price for d slightly larger than this. Their attacks can be seen as homogeneous diophantine approximation problems, and continued fraction techniques are used to solve them.

In this section, we will show that Wiener's and Verheul/Van Tilborg's attacks are special cases of our method. Moreover, we will show that in these cases the method is deterministic, in other words, it does not depend on the size of a (the parameter that describes the unbalancedness of the lattice).

Wiener [13] bases his attack on the fact that $\dfrac{k}{d}$ can be found as a convergent of $\dfrac{e}{N}$ if

$$\left|\frac{e}{N} - \frac{k}{d}\right| < \frac{1}{2d^2}.$$

It is commonly known (see for instance [9]) that this can also be described using a 2-dimensional lattice. When we assume no part of d is known ($d_M = d_L = 0$), it follows that $R = 0$ and

$$\Gamma = \begin{pmatrix} C & 0 \\ e & N \end{pmatrix}, \quad \mathbf{v} = \mathbf{0},$$

with C of size $N^{\beta-\delta+\frac{1}{2}} = N^{\frac{1}{2}}$, will reproduce Wiener's result, namely that the method will work if $\beta < \frac{1}{4}$. Later in this section we will show that the solution will be found by the shortest lattice vector only, making this case deterministic.

Verheul and Van Tilborg [11] have given an extension of Wiener's attack, where d is at most slightly larger than $N^{\frac{1}{4}}$ and no bits are known. To find $\frac{k}{d}$, they look not only at convergents of $\frac{e}{N}$, but also at 'linear combinations' of consecutive convergents, which, be it not the best, nevertheless are pretty good approximations. To be precise, when $\dfrac{p_{i-1}}{q_{i-1}}, \dfrac{p_i}{q_i}$ are consecutive convergents, then they also look for approximations to $\dfrac{e}{N}$ of the form $\dfrac{\lambda p_i + \mu p_{i-1}}{\lambda q_i + \mu q_{i-1}}$ for parameters $\lambda, \mu \in \mathbb{N}$. Then they have a weaker inequality to satisfy, of the form of

$$\left|\frac{e}{N} - \frac{k}{d}\right| < \frac{c}{d^2},$$

where the exact value for c depends on the search region for λ and μ. In this way they show that in order to extend Wiener's result for $d < N^{\frac{1}{4}}$ by r bits, one has to do an additional computation of the complexity of an exhaustive search over $2r$ bits.

In the language of lattices this becomes immediately clear. With Γ as above and $\mathbf{v} = \mathbf{0}$ (as we're still in the homogeneous case), the results of Section 3.2 show that for $\delta = \beta = \frac{1}{4} + \epsilon$, the complexity of the attack is $\mathrm{O}(N^{2\epsilon}(\log N)^2)$.

The example given in [11] will go as follows in our method. We start with the lattice

$$\Gamma = \begin{pmatrix} 2^{38} & 0 \\ e & N \end{pmatrix} = \begin{pmatrix} 2^{38} & 0 \\ 7\,115\,167\,804\,808\,765\,210\,427 & 31\,877\,667\,548\,624\,237\,348\,233 \end{pmatrix}$$

(note that in [11] the value of e contains a misprint).

We compute the reduced basis

$$\Gamma_{\text{red}} = \begin{pmatrix} 42\,694\,311\,384\,449\,024 & 87\,227\,281\,088\,446\,464 \\ 34\,997\,160\,860\,155\,755 & -133\,735\,834\,148\,055\,649 \end{pmatrix}.$$

The lattice point we need is $\Gamma \begin{pmatrix} 2d \\ -k \end{pmatrix} = \Gamma \begin{pmatrix} 3\,295\,186 \\ -735\,493 \end{pmatrix} = \Gamma_{\text{red}} \begin{pmatrix} 11 \\ 5 \end{pmatrix}$. Here $2d$ appears instead of d because in [11] $ed \equiv 1 \pmod{\operatorname{lcm}(p-1, q-1)}$ is taken, and in this case $\gcd(p-1, q-1)$ appears to be equal to 2.

This shows that, at least in this example, the efficiency of our method is comparable to [11], since we had to search for the numbers 11 and 5 of resp. 3.5 and 2.3 bits, together less than 7 bits (rather than 6 bits, because we have to allow negative values for one of the coordinates).

The fact that Verheul and Van Tilborg require a computation of the complexity of a $2r$ bit exhaustive search to allow r unknown bits more than $\frac{1}{4}$th of N for both d and the unknown part of d (which, in this case, are of course the same), corresponds to our complexity results of Section 3.2. However, it does not directly imply that their method can be used in a partial key exposure setting. In that sense our result, with the homogeneous case being a special case of the general case, implies the result of [11], but not the other way around. We believe that the method of Verheul and Van Tilborg can be combined with the method of Baker and Davenport [1], for solving inhomogeneous diophantine approximation problems, but we see no advantages above our uniform and clean lattice method.

Finally, we will show that the cases of Wiener and Verheul/Van Tilborg are *deterministic* situations in our method.

Recall that we look for a small pair (d, k) such that

$$\begin{pmatrix} C & 0 \\ e & N \end{pmatrix} \begin{pmatrix} d \\ -k \end{pmatrix} = \begin{pmatrix} Cd \\ ed - kN \end{pmatrix} \approx \begin{pmatrix} N^{\beta+\frac{1}{2}} \\ N^{\beta+\frac{1}{2}} \end{pmatrix}.$$

We will argue that if $d < N^{\frac{1}{4}}$ (Wiener's case), this small vector is actually the smallest nonzero lattice vector, which will be found by the Lagrange reduction.

Suppose it is not the smallest vector. Then the smallest vector cannot be linearly independent from it, for else the product of their sizes is smaller than $N^{2\beta+1} < N^{\frac{3}{2}}$, whereas the determinant of the lattice is $\det(L) = CN = N^{\frac{3}{2}}$. This is a contradiction. The other option when $\begin{pmatrix} Cd \\ ed - kN \end{pmatrix}$ is not the smallest vector, is that the smallest vector is

$$\begin{pmatrix} Cx \\ ex - yN \end{pmatrix} = \alpha \begin{pmatrix} Cd \\ ed - kN \end{pmatrix}, \text{ for some } \alpha \in [-1, 1].$$

It follows that $d = \frac{1}{\alpha}x$ and $k = \frac{1}{\alpha}y$, and since $ed - k\phi(N) = 1$, it must hold that

$$ex - y\phi(N) = \alpha.$$

Since the left hand side is an integer, $\alpha \neq 0$, and $\alpha \in [-1, 1]$, it follows that $|\alpha| = 1$. Therefore, $d = |x|$ and $k = |y|$. Hence, the shortest reduced basis vector immediately gives us d and k. Thus, the method is clearly deterministic.

In the case of Verheul/Van Tilborg's attack, $d = N^{\frac{1}{4}+\epsilon}$, so

$$\begin{pmatrix} Cd \\ ed - kN \end{pmatrix} \approx \begin{pmatrix} N^{\beta+\frac{1}{2}} \\ N^{\beta+\frac{1}{2}} \end{pmatrix} = \begin{pmatrix} N^{\frac{3}{4}+\epsilon} \\ N^{\frac{3}{4}+\epsilon} \end{pmatrix},$$

so this vector is not the smallest reduced vector. However, one can see that the smallest vector must be linearly independent of it, so we know that

$$a^{-1}\det(L)^{\frac{1}{2}} \cdot N^{\frac{3}{4}+\epsilon} \geq \det(L).$$

It follows that $a < \det(L)^{-\frac{1}{2}}N^{\frac{3}{4}+\epsilon} = \det(L)^{-\frac{1}{2}}N^{\beta+\frac{1}{2}} = N^{\frac{1}{2}(\beta+\delta-\frac{1}{2})}$ and from the computations in Section 3.1, we know that this means that the search area is $O(N^{\beta+\delta-\frac{1}{2}}) = O(N^{2\epsilon})$. So one can see that in this case, one also does not depend on Assumption 1.

4.2 Comments on the Size of a in Other Cases

When we are outside the regions where the known continued fractions methods from Wiener and Verheul/Van Tilborg apply, the attack depends on Assumption 1, namely that the elements of Γ_{red} are all of size $\det(L)^{\frac{1}{2}}$. In this section, we will comment on how this assumption holds in practice.

Let m be the maximal entry of Γ_{red}, and $m = a\det(L)^{\frac{1}{2}}$. We want to check that for the matrices involved in the attacks of this paper, a is close to 1. Therefore, we performed tests for the attacks for small d in the following setup: N is an 2048 bit modulus, $\beta \in [0.25, 0.5]$, $\epsilon \in [0, 0.1]$, and $\delta = \mathrm{Min}\{\beta, \frac{1}{2} - \beta + \epsilon\}$.

For this case, the lattices behaved as expected. In 500 experiments, the average value of a was approximately 1.9, and the maximal value of a was approxi mately 39.

5 Efficiency of the Attack

Now let us give some intuition on how our attack compares in running time to the other known results on partial key exposure attacks on small d, by Ernst et al. [7].

Figure 5 and 6 are two pictures of the attacks that are currently known and that use knowledge of MSBs or LSBs of d for relatively small d. The pictures show, for each value of β (the size parameter of d) what fraction of d we need to know in order to mount a successful attack. The area where the attack of this paper applies is dark.

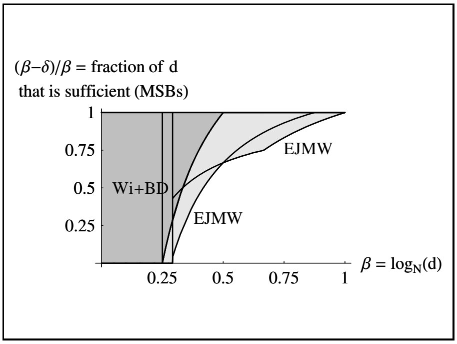

Fig. 5. Small d with known MSBs

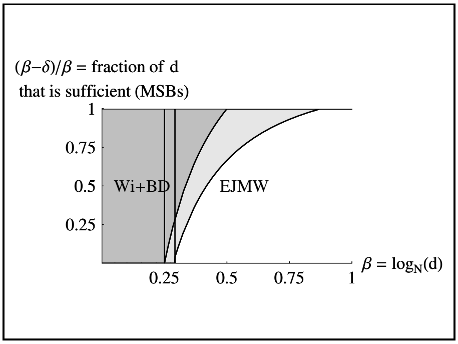

Fig. 6. Small d with known LSBs

One can see that our results do not exceed or match the optimal bounds which are already available. This is not surprising, since the attacks of [7] use large lattices, containing shifts of the RSA key equation and therefore combinations of monomials. Besides that, the attacks of [7] have asymptotic bounds. Hence, in those cases, to be able to perform close to the theoretic bound, there can be immensely large lattices involved, of which the reduction takes hours, days, or longer. Our attack belongs to those situations of partial key exposure that require only the reduction of a 2-dimensional lattice to solve, which an attacker can perform in just a few seconds. Moreover we can even exceed the theoretical bounds of the attacks with a small value ϵ.

To give some intuition of what this means in practice, we matched the 2-dimensional attack on small d and known MSBs against those of Ernst et al. [7]. The result is shown in the following table. For different values of β and δ, and different moduli N of 2048 bits, we computed the time to perform an attack for both methods.

This time includes:

- lattice reduction, resultant computations, and using the root $p+q-1$ of the resultant polynomials to find p and q, for [7],
- lattice reduction and trying all pairs (z_1, z_2) to find p, q, for the 2-dimensional method.

In the table, it shows that for $\beta = 0.30$, and $\delta = 0.205$, our attack works in approximately 2 seconds (this is an average over 50 experiments). It uses a simple Mathematica program that runs on a computer with Pentium III processor of 733 MHz.

On the other hand, for the same parameters we need about 40 minutes to solve the problem using one of the methods of [7], and the smallest lattice for which their attack works is of dimension 30 in this case. These experiments were done using Shoup's Number Theory Library [10], on a shared server with a Pentium IV Xeon processor of 2.80 GHz.

For the cases $\{\beta = 0.30, \delta = 0.210\}$, $\{\beta = 0.35, \delta = 0.150\}$, and $\{\beta = 0.35, \delta = 0.160\}$, one can see from the table that there are 'breaking points' for the methods

β	δ	Dim. lattice [7]	Time method [7]	Time 2D-method
0.30	0.050	10	35 sec.	1 sec.
0.30	0.100	10	35 sec.	1 sec.
0.30	0.150	10	35 sec.	1 sec.
0.30	0.200	30	40 min.	1 sec.
0.30	0.205	30	40 min.	2 sec.
0.30	0.210	30 / 50	40 min. / $4\frac{1}{2}$ hrs.	21 min.
0.35	0.050	10	35 sec.	1 sec.
0.35	0.100	10	35 sec.	1 sec.
0.35	0.150	14 / 30	1 min. / 40 min.	1 sec.
0.35	0.155	30	40 min.	2 sec.
0.35	0.160	30 / 50	40 min. / $4\frac{1}{2}$ hrs.	21 min.
0.40	0.050	10	35 sec.	1 sec.
0.40	0.100	14	1 min.	1 sec.
0.40	0.105	14	1 min.	2 sec.
0.40	0.110	14	1 min.	21 min.
0.45	0.050	14	1 min.	1 sec.
0.45	0.055	14	1 min.	2 sec.
0.45	0.060	14	1 min.	21 min.

Fig. 7. Experimental results: Comparison with [7]

of Ernst et al. For instance, if $\beta = 0.30$ and $\delta = 0.210$, the 30-dimensional lattice attack of [7] might suffice in some situations, whereas in others it will not lead to the solution. Therefore, for these parameters, it is possible that the attack takes either 40 minutes (if the attack using the 30-dimensional lattice works), or approximately $4\frac{1}{2}$ hours (if the 30-dimensional attack does not work and one has to use the 50-dimensional lattice attack).

6 Conclusion

We have shown how to perform a partial key exposure attack on RSA using a 2-dimensional lattice. The attack applies when the private exponent d is chosen to be small, which occurs in practice. In most cases, the attack is heuristic, but the underlying assumption is a reasonable one and supported by experiments. Although the attack does not achieve the theoretic bounds of known partial key exposure attacks using Coppersmith's method, it is much faster in the area where it applies. Moreover, the attack shows what you can achieve with the simplest lattices possible, and also provides a link with the known attacks based on continued fractions techniques, as they appear as special deterministic cases of our attack.

References

1. A. Baker, H. Davenport, "The equations $3x^2 - 2 = y^2$ and $8x^2 - 7 = z^2$", *Quarterly Journal of Mathematics (Oxford) (2)* **20** [1969], pp. 129–137.
2. Dan Boneh, "Twenty years of Attacks on the RSA Cryptosystem", *Notices of the American Mathematical Society* **46** [1999], pp. 203–213.

3. Dan Boneh, Glenn Durfee, " Cryptanalysis of RSA with Private Key d less than $N^{0.292}$", *IEEE Transactions on Information Theory* **46** [2000], pp. 1339–1349.
4. Dan Boneh, Glenn Durfee, Yair Frankel, "An Attack on RSA given a Small Fraction of the Private Key Bits", *Proceedings of ASIACRYPT 1998, LNCS* **1514** [1998], pp. 25–34.
5. Johannes Blömer, Alexander May, "New Partial Key Exposure Attacks on RSA", *Proceedings of CRYPTO 2003, LNCS* **2729** [2003], pp. 27–43.
6. Don Coppersmith, "Small Solutions to Polynomial Equations, and Low Exponent RSA Vulnerabilities", *Journal of Cryptology* **10** [1997], pp. 233–260.
7. Matthias Ernst, Ellen Jochemsz, Alexander May, Benne de Weger, "Partial Key Exposure Attacks on RSA up to Full Size Exponents", *Proceedings of EUROCRYPT 2005, LNCS* **3494** [2005], pp. 371–386.
8. Arjen Lenstra, Hendrik Lenstra, Jr., László Lovász, "Factoring Polynomials with Rational Coefficients", *Mathematische Annalen* **261** [1982], pp. 515–534.
9. Alexander May, "New RSA Vulnerabilities Using Lattice Reduction Methods", *PhD Thesis, University of Paderborn* [2003]
10. Victor Shoup, "Online Number Theory Library", *http://www.shoup.net/ntl*
11. Erik Verheul, Henk van Tilborg, "Cryptanalysis of 'less short' RSA Secret Exponents", *Applicable Algebra in Engineering, Communication and Computing* **8** [1997], pp. 425–435.
12. Benne de Weger, "Algorithms for Diophantine Equations", *CWI Tract 65, Centre for Mathematics and Computer Science, Amsterdam* [1989].
13. Michael Wiener, "Cryptanalysis of Short RSA Secret Exponents", *IEEE Transactions on Information Theory* **36** [1990], pp. 553–558.