

Attacking CRT-RSA with Small Exponents via Unravelled Linearization

Zhaopeng Ding¹, Zhaopeng Dai², Yanshuo Zhang¹, Ziyang Yan¹, and Ying Chen¹

¹ Department of Cryptographic Science and Technology, Beijing Electronic Science and Technology Institute, Beijing, China

16678784491@163.com, zhang_yanshuo@163.com, 985368553@qq.com, ychen@besti.edu.cn

² School of Mathematics and Statistics, Qingdao University, Qingdao, China
dzpeng@amss.ac.cn

Abstract. We develop a new Coppersmith lattice attack on balanced CRT-RSA with small CRT exponents d_p, d_q using unravelled linearization. The new attack matches the practical effectiveness of the Takayasu–Lu–Peng attack (JoC 2019) while using lattices of roughly half the dimension. In contrast to their attack, for which there is a noticeable gap between the theoretical predictions and experimental results, our theoretical estimates closely match the experimental results for all tested parameter sets. Our analysis further shows that the asymptotic bound remains $d_p, d_q < N^{0.122}$, so the new attack does not achieve the asymptotic improvement anticipated by Takayasu, Lu, and Peng. Nevertheless, performing LLL reduction on lower-dimensional lattices substantially reduces the running time, allowing us to attack much larger d_p, d_q in practice. Furthermore, through elegant variable substitutions, we extend our attack to two most significant bit (MSB) leakage models, obtaining the first attack exploiting MSB leakage of $p + q$ and an improved attack exploiting MSB leakage of d_p, d_q . In both models, the asymptotic bounds improve as the amount of leakage increases and recover the Takayasu–Lu–Peng bound in the absence of leakage.

Keywords: CRT-RSA · Coppersmith’s method · unravelled linearization · MSB leakage

1 Introduction

Since its introduction, RSA has become one of the most important and widely used public-key cryptosystems. However, RSA can be attacked in polynomial time when some secret parameters are sufficiently small. Wiener [19] first showed that RSA can be broken in polynomial time when the private exponent satisfies $d \leq N^{1/4}$. Subsequently, using Coppersmith’s method [6, 7], Boneh and Durfee [5] improved this bound first to $d \leq N^{0.284}$ and then to $d \leq N^{0.292}$. Although many later works have studied this problem, $d \leq N^{0.292}$ remains the best known bound for attacks on RSA with a small private exponent.

CRT-RSA. Compared with the case of a small private exponent, attacks on RSA with small CRT exponents seem to be more difficult. May [12] first showed at CRYPTO 2002 that the Boneh–Durfee attack [5] can be extended to this setting, but his result only works when p and q have unbalanced bit lengths. This result was later improved in [3]. For the balanced case, Jochemsz and May [11] first showed that RSA can be attacked in polynomial time using Coppersmith’s method when $d_p, d_q \leq N^{0.073}$. Subsequently, Herrmann and May [9] used unravelled linearization to reveal the lower-dimensional sublattice hidden in the Jochemsz–May construction. Their method greatly simplified the lattice construction, but did not improve the asymptotic bound.

The bound remained unchanged until EUROCRYPT 2017, when Takayasu, Lu, and Peng [17] improved it to $d_p, d_q < N^{0.091}$. In their later journal version [18], they further improved the bound to $d_p, d_q < N^{0.122}$. We call this the TLP attack.

However, the TLP attack works much better in experiments than its theoretical analysis suggests, especially for small lattice dimensions. For example, for a 31-dimensional lattice, the theoretical analysis gives no positive bound on δ in $d_p, d_q < N^\delta$, while experiments succeed for CRT exponents of size $N^{0.03}$. This difference suggests that the analysis does not fully use the structure of the lattice. Based on this observation, Takayasu, Lu, and Peng [18] suggested that a further analysis of the lattice structure, aimed at bridging the gap between theory and experiments, might lead to a better theoretical bound.

Partial key exposure attacks. Besides attacks on small CRT exponents, CRT-RSA can also be attacked when part of the secret key is leaked. Such attacks use the leaked secret bits together with public information to recover the secret key. At CRYPTO 2003, Blömer and May [4] studied the leakage of the most significant bits (MSBs) or least significant bits (LSBs) of one CRT exponent, and gave provable attacks for small public exponents. Subsequently, Takayasu and Kunihiko [16] improved these bounds at ACNS 2015 and extended the attacks to full-size public exponents by using some leaked MSBs or LSBs of both CRT exponents. For small CRT exponents, May, Nowakowski, and Sarkar [13] proposed a heuristic partial key exposure attack at ASIACRYPT 2021. Building on the TLP construction, their attack exploits partial LSBs of both d_p and d_q to extend the attackable range beyond the no-leakage bound of approximately $N^{0.122}$. At EUROCRYPT 2022, the same authors [14] introduced an attack based on approximate divisor multiples. For $e \asymp N^{1/12}$, knowledge of approximately one third of the bits of each CRT exponent, either MSBs or LSBs, suffices to factor N in polynomial time. Their MSB attack is provable, whereas the LSB variant relies on a standard Coppersmith-type heuristic.

Our Contributions. We present a new attack on RSA with small CRT exponents using unravelled linearization [8,9,1]. Compared with the TLP attack, our method reduces the lattice dimension by about half and provides a theoretical explanation for the gap between its experimental results and theoretical analysis. Further analysis shows that these lower-dimensional lattices give the same asymptotic bound of about $N^{0.122}$ as the TLP attack, and do not improve the

asymptotic bound as expected by Takayasu, Lu, and Peng [18]. Nevertheless, our method performs much better than the TLP attack in finite dimensions: for CRT exponents of the same size, it uses lower-dimensional lattices and greatly reduces the running time. We implement the algorithm and verify its effectiveness through experiments. For example, for a 1000-bit RSA modulus N , our method recovers 64-bit d_p and d_q in about 1/30 of the time required by the TLP attack. Figure 1 compares the theoretical bounds with the experimental results.




Fig. 1. Finite-dimensional bounds of our approach and the TLP attack.




Fig. 2. Required MSB leakage of $p + q$.




Fig. 3. Required MSB leakage of d_p, d_q .

However, the above construction appears to fully exploit the known algebraic relations of CRT-RSA, and our attempts with several alternative approaches have not produced an asymptotic bound beyond approximately 0.122. We therefore turn to partial information leakage and investigate whether it can improve this bound. Our unravelled-linearization framework extends naturally to two MSB leakage models: leakage of the MSBs of $p + q$ and of d_p, d_q . For $e \asymp N$, both models yield asymptotic bounds that improve with the amount of leakage and reduce to the TLP bound in the no-leakage case. The corresponding results are shown in Figures 2 and 3. Moreover, to the best of our knowledge, this is the first result for MSB leakage of $p + q$. For MSB leakage of d_p, d_q , our bound improves on that of Takayasu and Kunihiro [16] for δ slightly above 0.122, requiring fewer leaked MSBs.

Organization. The rest of this paper is organized as follows. Section 2 briefly reviews Coppersmith’s method. Section 3 presents our attack on CRT-RSA with small d_p, d_q using unravelled linearization and analyzes its theoretical bounds. Section 4 extends the attack to two MSB leakage settings. Section 5 reports our experimental results and compares our method with the TLP attack.

2 Coppersmith’s method

This section briefly reviews Coppersmith’s method for finding small modular roots of polynomials. For a more detailed introduction, we refer the reader to the relevant literature [6,7,10,18,15]. Let $f_1, \dots, f_n \in \mathbb{Z}[x_1, \dots, x_k]$ be n polynomials sharing a common root $\mathbf{r} = (r_1, \dots, r_k)$ modulo M , where the root satisfies the bounds $|r_j| < X_j$. For a suitable positive integer m , we construct a family of so-called shift polynomials.

$$g_i(x_1, \dots, x_k) = x_1^{j_1} \cdots x_k^{j_k} f_1^{i_1} \cdots f_n^{i_n} M^{m-(i_1+\dots+i_n)}$$

where each such shift satisfies $g_i(r_1, \dots, r_k) \equiv 0 \pmod{M^m}$.

Using the coefficient vectors of the shift polynomials, we construct a lattice L . If the so-called feasibility condition

$$(\det L)^{1/\dim L} < M^m \tag{1}$$

is satisfied, then one can obtain, in polynomial time, k polynomials

$$h_1(x_1, \dots, x_k), \dots, h_k(x_1, \dots, x_k),$$

such that $\mathbf{r}$ is not only a root of these polynomials modulo M^m , but also a common root over the integers. However, the polynomials $h_1, \dots, h_k$ do not necessarily reveal the root $\mathbf{r}$ directly. In practice, they typically generate an ideal whose variety is zero-dimensional.

$$\mathfrak{a} = \langle h_1, \dots, h_k \rangle \subseteq \mathbb{Q}[x_1, \dots, x_k].$$

The root $\mathbf{r}$ can then be obtained in polynomial time by computing a Gröbner basis of $\mathbf{a}$. However, there is no provable guarantee that the variety of $\mathbf{a}$ is always zero-dimensional. Consequently, many Coppersmith-type results rely on the following heuristic assumption.

Assumption 1. *The polynomials output by Coppersmith’s method in the multi-variate case generate an ideal defining a zero-dimensional variety.*

Since Assumption 1 may fail in some cases, experimental verification is important. For example, Bernstein et al. [2] observed that the short polynomials obtained after lattice reduction can all be multiples of a single linear polynomial, leading to a positive-dimensional solution set.

3 Attacks on CRT-RSA with Small d_p, d_q

Recall the defining equations of CRT-RSA

$$ed_p = k_p(p - 1) + 1, \quad (2)$$

$$ed_q = k_q(q - 1) + 1. \quad (3)$$

Here $N = pq$ is the RSA modulus, e is the public exponent, d_p, d_q are the corresponding CRT exponents, and $k_p, k_q \in \mathbb{N}$. Write $e = N^\alpha$, and for some $\alpha, \delta \in \mathbb{R}$, upper-bound $d_p, d_q \leq N^\delta$. Then k_p and k_q can be estimated as

$$k_p = \frac{ed_p - 1}{p - 1} < \frac{ed_p}{p - 1} = \Theta\left(\frac{ed_p}{N^{1/2}}\right) \leq \Theta\left(N^{\alpha+\delta-1/2}\right),$$

$$k_q = \frac{ed_q - 1}{q - 1} < \frac{ed_q}{q - 1} = \Theta\left(\frac{ed_q}{N^{1/2}}\right) \leq \Theta\left(N^{\alpha+\delta-1/2}\right),$$

In the usual RSA setting, we have $p, q = \Theta(N^{1/2})$. In 2019, Takayasu, Lu, and Peng [18] proposed a Coppersmith lattice attack that exploits the algebraic structure of CRT-RSA. For the case $\alpha = 1$, their asymptotic analysis gives the bound $\delta < 0.122$. However, their attack performs significantly better in experiments than predicted by the theoretical analysis, suggesting that the lattice construction may admit a tighter analysis and that its finite-dimensional behavior is not fully captured by the asymptotic bound.

In this section, we use unravalled linearization to characterize the sublattice structure underlying the TLP attack and explain the experimental behavior reported in [18]. Compared with the lattice construction of [18], our construction has substantially lower dimension, while its theoretical bounds closely match the experimental results presented in Section 5.

We now describe the attack in detail.

3.1 Construction of Modular Equations and Linearization

We first derive three fundamental modular equations from (2) and (3). These equations form the basic building blocks for the subsequent construction of the shift polynomials. First, reducing these equations modulo e gives

$$k_p p \equiv k_p - 1 \pmod{e}, \quad k_q q \equiv k_q - 1 \pmod{e}.$$

Multiplying these congruences and using $pq = N$, we obtain

$$f = (N - 1)k_p k_q + k_p + k_q - 1 \equiv 0 \pmod{e}$$

Next, multiplying the original integer equations (2) and (3) yields

$$(1 + k_p(p - 1))(1 + k_q(q - 1)) = e^2 d_p d_q.$$

Expanding the left-hand side and using $pq = N$, we obtain

$$\begin{aligned} g &= (N + 1)k_p k_q + k_p p + k_q q \\ &\quad - (k_p + k_q - 1) - (p + q)k_p k_q \equiv 0 \pmod{e^2} \end{aligned}$$

Finally, adding (2) and (3) gives

$$h = k_p p + k_q q - (k_p + k_q - 1) + 1 = e(d_p + d_q) \equiv 0 \pmod{e}$$

We introduce the following five linearization variables:

$$\begin{aligned} A &= k_p k_q, \\ B &= k_p + k_q - 1, \\ T &= k_p p + k_q q, \\ L &= p + q, \\ C &= k_p p + k_q q - (k_p + k_q - 1) - (p + q)k_p k_q. \end{aligned}$$

Their absolute values are bounded by X_A, X_B, X_T, X_L, X_C , respectively, where

$$\begin{aligned} X_A &= \Theta(N^{2\alpha+2\delta-1}), & X_B &= \Theta(N^{\alpha+\delta-1/2}), & X_C &= \Theta(N^{2\alpha+2\delta-1/2}), \\ X_T &= \Theta(N^{\alpha+\delta}), & X_L &= \Theta(N^{1/2}). \end{aligned}$$

We can then express f , g , and h in terms of A, B, C, T, L as follows:

$$\begin{cases} f = B + (N - 1)A \equiv 0 \pmod{e}, \\ g = C + (N + 1)A \equiv 0 \pmod{e^2}, \\ h = T - B + 1 \equiv 0 \pmod{e}. \end{cases}$$

3.2 Construction of Shift Polynomials

We first introduce the monomial set used to construct the lattice. Let m be a positive even integer, set $n = m/2$, and let $\tau \in [1/2, 1]$ be a parameter. We define

$$\mathcal{M}_{\text{core}} = \left\{ A^a B^b C^c T^t \left| \begin{array}{l} t \in \{0, 1\}, \quad a, b, c \in \mathbb{Z}_{\geq 0}, \\ 1 \leq a + b + c \leq n - t \end{array} \right. \right\},$$

and

$$\mathcal{M}_{\text{tail}} = \left\{ B^b C^c T^t L^\ell \left| \begin{array}{l} t \in \{0, 1\}, \quad b, c, \ell \in \mathbb{Z}_{\geq 0}, \\ 1 \leq b + c \leq n - t, \\ 1 \leq \ell \leq b + c - 1, \\ c + \ell + t \leq \tau(b + 2c + t) \end{array} \right. \right\}.$$

The corresponding polytopes are shown in Figure 4. The red surface represents the inclined plane controlled by the parameter τ .




Fig. 4. Polytopes for $\mathcal{M}_{\text{core}}$ (left) and $\mathcal{M}_{\text{tail}}$ (right) for fixed t .

Therefore, we choose the following monomial set to index the columns of the lattice.

$$\mathcal{M} = \mathcal{M}_{\text{core}} \cup \mathcal{M}_{\text{tail}}.$$

For each monomial $M = A^a B^b C^c T^t L^\ell \in \mathcal{M}$, we uniformly define the shift polynomial

$$F_M = e^{m-t-b-2c} h^t A^a L^\ell f^b g^c.$$

Since $e \mid h(r)$, $e \mid f(r)$, and $e^2 \mid g(r)$, each shift satisfies $F_M(r) \equiv 0 \pmod{e^m}$.

3.3 Unravelled Linearization

To obtain a square lattice with a triangular basis matrix, we further apply unravelled linearization to the shift polynomials using the algebraic relations among A, B, C, T, L . In particular, these variables satisfy two types of relations:

$$\begin{aligned} AL + B + C - T &= 0, \\ BLT + (B + C)L - N(B + 1)^2 + 4NA - T^2 &= 0. \end{aligned}$$

These relations can be readily derived by a Gröbner basis computation under the lexicographic order ($p \succ q \succ k_p \succ k_q \succ L \succ T \succ C \succ B \succ A$). They reveal a useful structural property. Suppose we choose a monomial order such that f , g , and h have leading monomials B , C , and T , respectively. Then A is not the leading monomial of any of these polynomials. However, the first relation allows us to rewrite AL in terms of T , B , and C , while the second allows us to reduce higher powers of T . These substitutions preserve the values of the shift polynomials at the target root and lead to a lower-dimensional lattice.

Accordingly, we define the following transformation operator.

Definition 1. Let F be a polynomial in the variables A, B, C, T, L . We define $\text{trans}(F)$ by repeatedly applying the following reduction rules to the monomials of F . If a monomial contains both A and L but not T , apply

$$AL \mapsto T - B - C. \quad (\text{R1})$$

If a monomial contains A , L , and one factor of T , apply

$$ALT \mapsto BLT + (B + C)L - N(B + 1)^2 + 4NA - (B + C)T. \quad (\text{R2})$$

We repeat these reductions until no monomial contains both A and L .

Lemma 1. For every polynomial F , the above reduction process terminates regardless of the order in which the rules are applied.

Proof. By (R1) and (R2), each reduction strictly decreases the total degree in A and L of every resulting monomial. Since this degree is a nonnegative integer, only finitely many reductions can occur along any branch. Therefore, the above reduction process terminates. $\square$

Lemma 2. Under the reduction procedure in Definition 1, let $r, b, c \in \mathbb{Z}_{\geq 0}$ and $\ell \in \mathbb{Z}_{\geq 1}$. If $r \geq 1$, every monomial appearing in $\text{trans}(A^r B^b C^c L^\ell)$ has L -degree strictly less than ℓ . Moreover,

$$\text{trans}(A^r B^b C^c T L^\ell) = B^{b+r} C^c T L^\ell + R,$$

where every monomial M appearing in R satisfies

$$\deg_L M < \ell \quad \text{or} \quad (\deg_L M = \ell \text{ and } \deg_T M = 0).$$

Proof. Neither reduction rule increases the L -degree. If the L -degree is unchanged, the T -degree cannot increase.

For the first claim, the first reduction must use (R1). This lowers the L -degree from ℓ to $\ell - 1$ in every resulting monomial, and later reductions cannot increase it.

For the second claim, the case $r = 0$ is immediate. Suppose $r \geq 1$. In (R2), BLT is the only term that preserves both the L -degree and the T -degree, and its coefficient is one. Thus the only way to preserve both degrees is to follow

this term at each step, replacing one factor of A by B . After r steps, this gives $B^{b+r}C^cTL^\ell$ with coefficient one.

All other terms produced along this sequence have either L -degree less than ℓ , or L -degree ℓ and T -degree zero. Further reductions preserve this property, so every monomial in the remainder R satisfies the stated conditions. $\square$

Lemma 3. *For any monomial $M \in \mathcal{M}$, we have*

$$\text{supp}(\text{trans}(F_M)) \subseteq \mathcal{M}.$$

Proof. We omit the scalar factor $e^{m-t-b-2c}$, since it does not affect the support.

First, let $M = A^aB^bC^cT^t \in \mathcal{M}_{\text{core}}$. Every monomial $A^iB^jC^kT^s$ in the expansion of $h^tA^af^bg^c$ satisfies

$$s \in \{0, 1\}, \quad i + j + k \geq a + b + c \geq 1, \quad i + j + k + s \leq a + b + c + t \leq n.$$

Hence every such monomial belongs to $\mathcal{M}_{\text{core}}$. Since none contains L , no reduction is needed.

Now let $M = B^bC^cT^tL^\ell \in \mathcal{M}_{\text{tail}}$. For a monomial $U = A^iB^jC^kT^sL^v$, define

$$\begin{aligned} q(U) &= i + j + k + s, \\ \zeta(U) &= i + j + k - v, \\ \Omega_\tau(U) &= \tau(i + j + 2k + s) - (k + s + v). \end{aligned}$$

These quantities are additive under multiplication of monomials. Every monomial U in the expansion of $h^tL^\ell f^bg^c$ has T -degree at most one and satisfies

$$\begin{aligned} q(U) &\leq b + c + t \leq n, \\ \zeta(U) &\geq b + c - \ell \geq 1, \\ \Omega_\tau(U) &\geq \tau b + (2\tau - 1)c + (\tau - 1)t - \ell \\ &= \tau(b + 2c + t) - (c + t + \ell) \geq 0. \end{aligned}$$

For the bound on $\Omega_\tau(U)$, each monomial chosen from a factor f , g , or h contributes at least τ , $2\tau - 1$, or $\tau - 1$, respectively, while L^ℓ contributes $-\ell$.

We next show that the reductions preserve these three bounds. The following table gives the changes in q , ζ , and Ω_τ when the left-hand monomial of each rule

is replaced by a monomial on the right-hand side. We ignore coefficients.

Rule	Right-hand monomial	Δq	$\Delta \zeta$	$\Delta \Omega_\tau$
(R1)	T	0	0	0
	B	0	1	1
	C	0	1	τ
(R2)	BLT	0	0	0
	BL	-1	0	$1 - \tau$
	CL	-1	0	0
	B^2	0	2	2
	A, B	-1	1	$2 - \tau$
	1	-2	0	$2 - 2\tau$
	BT	0	1	1
	CT	0	1	τ

Since $\tau \in [1/2, 1]$, q never increases, while ζ and Ω_τ never decrease. By additivity, this also holds when either rule is multiplied by any monomial.

Thus every monomial $V = A^i B^j C^k T^s L^v$ in the final polynomial satisfies

$$q(V) \leq n, \quad \zeta(V) \geq 1, \quad \Omega_\tau(V) \geq 0.$$

Both rules preserve the condition that the T -degree is at most one, so $s \in \{0, 1\}$. At termination, no monomial contains both A and L , so $iv = 0$.

If $v = 0$, the first two bounds give

$$1 \leq i + j + k \leq n - s,$$

so $V \in \mathcal{M}_{\text{core}}$. If $v \geq 1$, then $i = 0$, and the three bounds give

$$1 \leq v \leq j + k - 1, \quad j + k \leq n - s, \quad k + v + s \leq \tau(j + 2k + s).$$

Hence $V \in \mathcal{M}_{\text{tail}}$. This proves the claimed support containment. $\square$

Lemma 4. *Let $M = A^a B^b C^c T^\ell L^\ell$ be a selected monomial. Under the lexicographic order $L \succ T \succ C \succ B \succ A$, we have*

$$\text{LT}(\text{trans}(F_M)) = \begin{cases} e^{m-t-b-2c} M, & \ell = 0, \\ e^{m-b-2c} M, & \ell > 0, t = 0, \\ e^{m-1-b-2c} N^b M, & \ell > 0, t = 1. \end{cases}$$

Proof. We first omit the nonzero scalar factor $e^{m-t-b-2c}$ and set $Q_M = h^t A^a L^\ell f^b g^c$.

Case 1: $\ell = 0$. No reduction is needed. Since $\text{LT}(h) = T$, $\text{LT}(f) = B$, and $\text{LT}(g) = C$, we have $\text{LT}(\text{trans}(Q_M)) = M$.

Case 2: $\ell > 0$ and $t = 0$. Here $a = 0$, so $Q_M = L^\ell f^b g^c$. The only term in its expansion that does not contain A is $L^\ell B^b C^c = M$, with coefficient one. By

Lemma 2, all other terms reduce to monomials of L -degree less than ℓ . Thus the leading term is M .

Case 3: $\ell > 0$ and $t = 1$. Again $a = 0$. Using $h = T - B + 1$, we write

$$Q_M = TL^\ell f^b g^c + (1 - B)L^\ell f^b g^c.$$

By Lemma 2, every monomial obtained by reducing the second summand has either L -degree less than ℓ , or L -degree ℓ and T -degree zero.

For the first summand, the same lemma shows that the terms with L -degree ℓ and T -degree one are obtained by replacing A with B in $f^b g^c$ and multiplying by TL^ℓ . Since $f|_{A=B} = NB$ and $g|_{A=B} = C + (N + 1)B$, we obtain

$$\text{trans}(Q_M) = N^b TL^\ell B^b (C + (N + 1)B)^c + R,$$

where every monomial in R has either L -degree less than ℓ , or L -degree ℓ and T -degree zero. The remainder therefore does not affect the leading term. Since $C \succ B$, we have $\text{LT}(\text{trans}(Q_M)) = N^b M$.

Restoring the scalar factor $e^{m-t-b-2c}$ gives the stated leading term in each case. $\square$

3.4 Improving the TLP Attack

We first present our lattice construction and prove that it achieves the same asymptotic bound as the TLP attack.

Theorem 1. *Fix a constant*

$$0 \leq \delta < \frac{1}{2} - \frac{1}{\sqrt{7}} \approx 0.122.$$

Let $N = pq$ be a sufficiently large RSA modulus, where p and q are distinct primes of the same bit-size. Let e be a public exponent satisfying $1 < e < \phi(N)$, $\gcd(e, \phi(N)) = 1$, and $e \asymp N$. If the corresponding CRT exponents satisfy $d_p, d_q \leq N^\delta$, then, under Assumption 1, N can be factored from (N, e) in time polynomial in $\log N$.

Proof. We first compute $\gcd(N, e)$. If it is greater than one, it gives a nontrivial factor of N , since $e < N$. We may therefore assume that $\gcd(N, e) = 1$.

Write $\alpha = \log_N e$. Let m be a positive even integer and $\tau \in [1/2, 1]$, both to be chosen below, and set $\omega = |\mathcal{M}|$. For each $M = A^a B^b C^c T^t L^\ell \in \mathcal{M}$, Lemmas 3 and 4 give

$$P_M := \text{trans}(F_M) = e^{\nu_M} (c_M M + R_M), \quad \nu_M = m - t - b - 2c \geq 0,$$

where $c_M = N^b$ if $\ell > 0$ and $t = 1$, and $c_M = 1$ otherwise. Every monomial in R_M belongs to $\mathcal{M}$ and is smaller than M under the lexicographic order $L \succ T \succ C \succ B \succ A$.

We next normalize the leading coefficients. Since $\gcd(c_M, e) = 1$, choose an integer $0 \leq u_M < e^m$ such that $u_M c_M \equiv 1 \pmod{e^m}$, and set $\hat{P}_M = e^{\nu_M}(M + u_M R_M)$. Then

$$\hat{P}_M - u_M P_M = e^{\nu_M}(1 - u_M c_M)M \equiv 0 \pmod{e^m}.$$

Thus $\hat{P}_M$ still vanishes at the target root modulo e^m , and its leading term is $e^{\nu_M}M$.

Scale the variables in the normalized polynomials by the integer bounds X_A, X_B, X_C, X_T, X_L , and let Λ be the lattice generated by the resulting coefficient vectors. Order the rows and columns increasingly by the above monomial order. The basis matrix is then lower triangular, with diagonal entry

$$d_M = e^{m-t-b-2c} X_A^a X_B^b X_C^c X_T^t X_L^\ell$$

corresponding to M . Thus $\det(\Lambda) = \prod_{M \in \mathcal{M}} d_M$, and the determinant condition (1) requires $\prod_{M \in \mathcal{M}} d_M < e^{m\omega}$.

Substituting the variable bounds and comparing the leading m^4 terms in the exponents of N , we obtain the following sufficient asymptotic condition:

$$\begin{aligned} & \frac{\alpha(1+5\tau)}{96} + \frac{2\alpha+2\delta-1}{192} + \frac{\tau(\alpha+\delta-\frac{1}{2})}{48} \\ & + \frac{(5\tau-1)(2\alpha+2\delta-\frac{1}{2})}{192} + \frac{7\tau^2-5\tau+1}{384} < \frac{\alpha\tau}{8}. \end{aligned}$$

The contribution from X_T is of order m^3 , so it does not affect the leading term. The inequality simplifies to $4\alpha + 28\tau\delta + 7\tau^2 - 14\tau < 0$, or equivalently,

$$\delta < \frac{1}{2} - \frac{\tau}{4} - \frac{\alpha}{7\tau}.$$

Since $e \asymp N$, we have $\alpha \rightarrow 1$ as $N \rightarrow \infty$. The limiting right-hand side is maximized at $\tau = 2/\sqrt{7} \in [1/2, 1]$, where it equals $\frac{1}{2} - \frac{1}{\sqrt{7}}$. We choose this value of τ .

Since δ is strictly below this threshold, we can choose a sufficiently large even m , depending only on δ , so that the determinant condition holds for all sufficiently large N . For this fixed m , the lattice dimension and polynomial degrees are independent of N . Under Assumption 1, lattice reduction and root recovery then yield a factorization of N in time polynomial in $\log N$. $\square$

We next give a quantitative comparison. For a fixed even m and $\alpha = 1$, keeping all degree sums in the determinant condition gives the finite-dimensional bound

$$\delta < \delta_m(\tau) := \frac{\sum_{M \in \mathcal{M}} (b+c-\ell-2a)}{2 \sum_{M \in \mathcal{M}} (2a+b+2c+t)}, \quad (4)$$

where $M = A^a B^b C^c T^t L^\ell$ and $\mathcal{M}$ is defined using the chosen value of τ .

The set $\mathcal{M}$, and hence $\delta_m(\tau)$, can change only at finitely many breakpoints of the form $\tau = (c + \ell + t)/(b + 2c + t)$, where

$$t \in \{0, 1\}, \quad b, c \in \mathbb{Z}_{\geq 0}, \quad 1 \leq b + c \leq m/2 - t, \quad 1 \leq \ell \leq b + c - 1.$$

Between these breakpoints, $\delta_m(\tau)$ is constant. Since the inequality defining $\mathcal{M}_{\text{tail}}$ is non-strict, it is enough to evaluate (4) at $\tau = 1/2$ and all breakpoints in $[1/2, 1]$. This gives the exact optimal bound within our shift family for each fixed m . Table 1 lists the optimal bounds for the first few even values of m , together with the corresponding values of τ and lattice dimensions.

Table 1. Finite-dimensional optimization results for $\alpha = 1$.

m	Our lattice			TLP lattice [18]		
	τ	dimension	δ	τ	dimension	δ
2	1/2	3	0	1	12	$-1/16 \approx -0.0625$
4	3/4	15	$1/28 \approx 0.0357$	3/4	31	0
6	5/6	42	$9/166 \approx 0.0542$	5/6	84	$6/161 \approx 0.0373$
8	6/7	89	$59/904 \approx 0.0653$	7/8	179	$25/451 \approx 0.0554$
10	5/6	155	$14/191 \approx 0.0733$	6/7	316	$43/646 \approx 0.0666$
12	5/6	252	$7/88 \approx 0.0795$	5/6	501	$273/3656 \approx 0.0747$
14	9/11	371	$533/6316 \approx 0.0844$	5/6	758	$259/3209 \approx 0.0807$

Our construction gives the same asymptotic bound as the TLP attack [18], but uses lattices of about half the dimension. For large m , TLP requires LLL reduction on a lattice of dimension approximately $\frac{\tau}{4}m^3$, whereas our lattice has dimension approximately $\frac{\tau}{8}m^3$. Figure 1 compares the bounds on d_p, d_q achieved by the two methods for finite values of m . Besides allowing us to attack larger CRT exponents in practice, our construction has another important advantage: its experimental results are fully explained by our theoretical analysis. The same analysis also explains the experimental results reported in [18], closing the gap between the original analysis and the observed performance.

4 Attacks with MSB Leakage

Our results suggest that the asymptotic bound of approximately 0.122 may be optimal for Coppersmith attacks based on the currently known algebraic relations of CRT-RSA. The polynomials f , g , and h capture essentially all the nontrivial algebraic information we have found, and our attempts to use these relations in different ways have not improved this bound. A better asymptotic bound may therefore require new algebraic relations, as illustrated by the improvement in [18]

over [11]. We therefore turn to partial information leakage. The construction in Section 3 provides a simple way to improve the asymptotic bound when some of the most significant bits (MSBs) of certain CRT-RSA quantities are leaked. We consider the following two MSB leakage models.

4.1 MSB Leakage of $p + q$

We use the CRT-RSA setting of Section 3, with $e \asymp N$ and $d_p, d_q \leq N^\delta$. Suppose that a known integer S satisfies $|p + q - S| = O(N^\rho)$ for some $0 \leq \rho \leq 1/2$. For MSB leakage, we obtain S by appending zeros to the known prefix of $p + q$. If the u least significant bits are unknown, then $0 \leq p + q - S < 2^u$, so this model corresponds to $u = \rho \log_2 N + O(1)$.

We use this information by setting $x = L - S$ and $C_S = C + SA = T - B - Ax$, while leaving A, B, T unchanged. The basic modular equations then become

$$\begin{cases} f = B + (N - 1)A \equiv 0 \pmod{e}, \\ g_S = C_S + (N + 1 - S)A \equiv 0 \pmod{e^2}, \\ h = T - B + 1 \equiv 0 \pmod{e}. \end{cases}$$

Using $C_S = T - B - Ax$, we choose integer root bounds of the following orders:

$$\begin{aligned} X_A &= \Theta(N^{1+2\delta}), & X_B &= \Theta(N^{1/2+\delta}), & X_{C_S} &= \Theta(N^{1+2\delta+\rho}), \\ X_T &= \Theta(N^{1+\delta}), & X_x &= \Theta(N^\rho). \end{aligned}$$

Theorem 2. *Under Assumption 1 for the combined polynomial system specified below, N can be factored from (N, e, S) in time polynomial in $\log N$ whenever $0 \leq \delta < \Delta(\rho)$, where*

$$\Delta(\rho) = \begin{cases} \frac{5}{14} - \frac{\rho}{2}, & 0 \leq \rho \leq \frac{2}{7}, \\ \frac{1}{2} - \sqrt{\frac{2\rho}{7}}, & \frac{2}{7} \leq \rho \leq \frac{1}{2}. \end{cases}$$

Proof. Fix a positive even integer m and a parameter $\tau \in [1/2, 1]$. Let $\mathcal{M}_S$ be the monomial set obtained from Section 3.2 by replacing C, L with C_S, x , respectively. For each $M = A^a B^b C_S^c T^t x^\ell \in \mathcal{M}_S$, define

$$F_{M,S} = e^{m-t-b-2c} h^t A^a x^\ell f^b g_S^c.$$

Each shift vanishes at the target root modulo e^m .

The centered variables satisfy the reduction rules $Ax \mapsto T - B - C_S$ and

$$\begin{aligned} AxT &\mapsto BxT + (B + C_S)x - N(B + 1)^2 + (4N - S^2)A \\ &\quad + ((S - 1)B - C_S - S)T + 2S(B + C_S). \end{aligned}$$

These identities follow by substituting $L = S + x$ and $C = C_S - SA$ into the algebraic relations in Section 3.3.

Each reduction strictly decreases $\deg_A + \deg_x$ in every branch. Compared with the original second rule, the only additional monomials are C_S and T . Relative to AxT , their changes in $(q, \zeta, \mathcal{Q}_\tau)$ from Lemma 3 are $(-1, 1, 1)$ and $(-1, 0, 1-\tau)$, respectively. Thus the proofs of termination and support containment in Section 3.3 still apply for $\tau \in [1/2, 1]$.

In the second rule, BxT is the only term with the same x -degree and T -degree as AxT . Therefore, the leading-term argument in Lemma 4 still applies under the order $x \succ T \succ C_S \succ B \succ A$. Apart from the factor $e^{m-t-b-2c}$, the leading coefficient is still N^b when $\ell > 0$ and $t = 1$, and is 1 otherwise.

We may assume that $\gcd(N, e) = 1$; otherwise, the gcd immediately gives a factor of N . Applying the same modular normalization as in Section 3.4, we obtain a lattice Λ_S with a triangular basis and determinant

$$\det(\Lambda_S) = \prod_{M \in \mathcal{M}_S} e^{m-t-b-2c} X_A^a X_B^b X_{C_S}^c X_T^t X_x^\ell.$$

Compared with Section 3, only the root bounds for C and L change, and both exponents in N decrease by $1/2 - \rho$. Thus the finite-dimensional bound becomes

$$\delta < \delta_{m,\rho}(\tau) := \delta_m(\tau) + \left(\frac{1}{2} - \rho\right) \frac{\sum_{M \in \mathcal{M}_S} (c + \ell)}{\sum_{M \in \mathcal{M}_S} (2a + b + 2c + t)},$$

where $\delta_m(\tau)$ is the bound defined in Section 3.4.

The degree sums over the same core–tail monomial set give

$$\begin{aligned} \sum_{M \in \mathcal{M}_S} (c + \ell) &= \frac{7\tau^2}{192} m^4 + O(m^3), \\ \sum_{M \in \mathcal{M}_S} (2a + b + 2c + t) &= \frac{7\tau}{96} m^4 + O(m^3). \end{aligned}$$

Together with the asymptotic bound in Section 3.4, these expressions yield

$$\delta_{m,\rho}(\tau) = \frac{1}{2} - \frac{\rho\tau}{2} - \frac{1}{7\tau} + O\left(\frac{1}{m}\right).$$

Hence it suffices to have $\delta < 1/2 - \rho\tau/2 - 1/(7\tau)$. Over $\tau \in [1/2, 1]$, the right-hand side is maximized at $\tau = 1$ if $\rho \leq 2/7$, and at $\tau = \sqrt{2/(7\rho)}$ otherwise. This gives the stated expression for $\Delta(\rho)$. $\square$

When $\rho = 1/2$, we have $\Delta(1/2) = 1/2 - 1/\sqrt{7}$, which recovers the no-leakage bound of Section 3. For each fixed m and τ , the centered and original constructions have exactly the same lattice dimension. Figure 5 shows the corresponding finite-dimensional bounds.




Fig. 5. Finite-dimensional bounds under MSB leakage of $p + q$.

4.2 MSB Leakage of d_p, d_q

We again use the setting of Section 3, with $e \asymp N$ and $d_p, d_q \leq N^\delta$. Suppose that the known MSBs give integers D_p, D_q such that

$$d_p = D_p + x, \quad d_q = D_q + y, \quad 0 \leq x, y < N^\gamma,$$

where $0 \leq \gamma \leq \delta \leq 1/2$. Each unknown suffix therefore has at most $\gamma \log_2 N + O(1)$ bits. We use this information to center only $A = k_p k_q$; the errors x, y are not used as lattice variables.

The key equations give $A(N + 1 - L) = (ed_p - 1)(ed_q - 1)$. We define the public integer

$$A_0 = \left\lfloor \frac{(eD_p - 1)(eD_q - 1)}{N + 1} + \frac{1}{2} \right\rfloor$$

and set $\tilde{A} = A - A_0$. To bound $\tilde{A}$, note that

$$A - \frac{(eD_p - 1)(eD_q - 1)}{N + 1} = \frac{e^2(D_p y + D_q x + xy) - e(x + y) + AL}{N + 1}.$$

Using the bounds from Section 3, we obtain

$$|\tilde{A}| \ll N^{1+\delta+\gamma} + N^{1/2+2\delta} + 1 \ll N^{1+\delta+\gamma}.$$

We therefore choose an integer bound $X_{\tilde{A}} = \Theta(N^{1+\delta+\gamma})$ and keep the original bounds for B, C, T, L .

Theorem 3. *Let δ and γ be fixed and satisfy the conditions above. Under Assumption 1 for the combined polynomial system specified below, N can be factored from (N, e, D_p, D_q) in time polynomial in $\log N$ for all sufficiently large N , provided that*

$$\gamma < \Gamma(\delta) := 14\delta^2 - 13\delta + \frac{3}{2}. \quad (5)$$

Proof. Fix a positive even integer m and a parameter $\tau \in [1/2, 1]$. Let $\hat{P}_M$ be the normalized polynomials constructed in the proof of Theorem 1, and define

$$Q_M(\tilde{A}, B, C, T, L) = \hat{P}_M(\tilde{A} + A_0, B, C, T, L), \quad M \in \mathcal{M}.$$

We apply the translation after the original reductions and normalization, so each Q_M still vanishes at the translated target root modulo e^m .

Decreasing the A -degree of a core monomial gives another core monomial, except possibly for 1 and T . Tail monomials contain no A and are therefore unchanged. Thus the support of every translated polynomial is contained in

$$\mathcal{M}^+ = \left\{ \tilde{A}^a B^b C^c T^t L^\ell : A^a B^b C^c T^t L^\ell \in \mathcal{M} \right\} \cup \{1, T\}.$$

To complete the basis, we add e^m and $e^{m-1}(T - B + 1)$. Both vanish at the target root modulo e^m , since $T - B + 1 = e(d_p + d_q)$. Under the order $L \succ T \succ C \succ B \succ \tilde{A}$, the translation preserves each original leading term, with A replaced by $\tilde{A}$, while the two added polynomials have leading terms e^m and $e^{m-1}T$. Thus the scaled coefficient matrix is lower triangular and has dimension $\omega_+ = |\mathcal{M}| + 2$.

The determinant of the resulting lattice Λ_{MSB} is

$$\det(\Lambda_{\text{MSB}}) = e^{2m-1} X_T \prod_{A^a B^b C^c T^t L^\ell \in \mathcal{M}} e^{m-t-b-2c} X_{\tilde{A}}^a X_B^b X_C^c X_T^t X_L^\ell.$$

The determinant condition (1) requires $\det(\Lambda_{\text{MSB}}) < e^{m\omega_+}$. Substituting the root bounds and comparing the leading m^4 terms in the exponents of N , as in Section 3.4, gives the sufficient condition

$$\begin{aligned} & \frac{1+5\tau}{96} + \frac{1+\delta+\gamma}{192} + \frac{\tau(\frac{1}{2}+\delta)}{48} \\ & + \frac{(5\tau-1)(\frac{3}{2}+2\delta)}{192} + \frac{7\tau^2-5\tau+1}{384} < \frac{\tau}{8}. \end{aligned}$$

The contribution from X_T is of order m^3 , while the two added rows contribute only terms of order m . Neither therefore affects the leading m^4 terms. After simplification, the inequality becomes $4 + (28\tau - 2)\delta + 2\gamma + 7\tau^2 - 14\tau < 0$, or equivalently,

$$\gamma < -\frac{7}{2}\tau^2 + 7(1-2\delta)\tau + \delta - 2.$$

The right-hand side is maximized at $\tau = 1 - 2\delta$, where it equals $\Gamma(\delta) = 14\delta^2 - 13\delta + 3/2$. The condition $\gamma < \Gamma(\delta)$, together with $0 \leq \gamma$ and $\delta \leq 1/2$, implies $\delta < 1/4$, so this choice of τ lies in $[1/2, 1]$. $\square$

Figure 6 shows the corresponding finite-dimensional bounds.




Fig. 6. Finite-dimensional bounds under MSB leakage of d_p, d_q .

Comparison with the no-leakage and previous MSB bounds. Setting $\gamma = \delta$ in (5) recovers the no-leakage condition $\delta < \delta_0 := 1/2 - 1/\sqrt{7}$. For $e \asymp N$, the double-MSB bound of Takayasu and Kunihiro [16, Theorem 7] becomes

$$\Gamma_{\text{TK}}(\delta) = \sup_{\sigma > 0} \frac{(6 - 12\delta)\sigma^2 + (4 - 16\delta)\sigma + 1 - 4\delta}{24\sigma^3 + 54\sigma^2 + 40\sigma + 10}.$$

Let $U_\delta(\sigma)$ and $V(\sigma)$ denote the numerator and denominator, respectively. For $\delta_0 < \delta \leq 1/8$, all coefficients of $(3/32)V - U_\delta$ are strictly positive, which gives

$$\Gamma_{\text{TK}}(\delta) < \frac{3}{32} \leq \Gamma(\delta).$$

Thus our bound strictly improves the previous bound throughout this interval. In particular, at $\delta = 1/8$, our condition is $\gamma < 3/32$, whereas optimizing the previous bound gives $\gamma < 0.0566321\dots$. For exponents of bit length $(1/8)\log_2 N + O(1)$, these conditions require revealing slightly more than 25% and approximately 54.69% of the MSBs of each exponent, respectively. Figure 3 gives a detailed comparison.

5 Experiments

We first run a series of experiments on the CRT-RSA lattice attack from Section 3. Our experiments have two main goals. First, we check how closely the theoretical bounds match the experimental results. Second, since the attack is heuristic, we test Assumption 1 (regarding the variables k_p, k_q, p, q) and confirm that the target roots can be recovered efficiently in the tested instances.

We also reimplemented the TLP attack and compared it with our method. Table 2 shows that our method substantially reduces the time needed for LLL reduction. For example, the speedup is about 10-fold for $m = 6$ and about 30-fold for $m = 8$. This improvement mainly comes from reducing the lattice dimension by about half. The smaller lattices also allow us to test parameter sets that were previously too computationally expensive to handle.

Moreover, the largest sizes of d_p and d_q recovered in our experiments agree almost exactly with those predicted by the theoretical bound on δ . This close agreement provides strong experimental evidence that our analysis is tight for this lattice construction.

Table 2. Experimental comparison between our lattice and the TLP lattice.

m	N	d_p, d_q	Lattice dimension		LLL time	
			Ours	TLP [18]	Ours	TLP [18]
4	1000 bit	35 bit	15	31	0.02 s	0.22 s
6	1000 bit	53 bit	42	84	7 s	70 s
8	1000 bit	64 bit	89	179	337 s	9902 s
10	1000 bit	72 bit	155	316	8404 s	–
12	1000 bit	78 bit	252	501	32.8 h	–
4	2000 bit	70 bit	15	31	0.14 s	0.75 s
6	2000 bit	107 bit	42	84	25 s	261 s
8	2000 bit	129 bit	89	179	1601 s	12.4 h

To further assess our construction, we counted the nonzero polynomials obtained after lattice reduction that vanish at the target small root for each parameter set. The results are reported in Table 3. For every tested parameter set, our method produces exactly one fewer vanishing polynomial than the lattice dimension. Under the same parameters, the TLP attack produces the same number of vanishing polynomials, despite using a larger lattice. These results strongly indicate that our construction is optimal, in the sense that no lower-dimensional Coppersmith-type dense sublattice exists. Indeed, it is impossible to obtain as many vanishing polynomials as the lattice dimension, since doing so would amount to recovering the target small root directly.

Table 3. Comparison of lattice dimensions and the numbers of vanishing polynomials.

m	N	d_p, d_q	Our lattice		TLP lattice [18]	
			Dimension	numbers	Dimension	numbers
4	1000 bit	35 bit	15	14	31	14
6	1000 bit	53 bit	42	41	84	41
8	1000 bit	64 bit	89	88	179	88
10	1000 bit	72 bit	155	154	316	–
12	1000 bit	78 bit	252	251	501	–

In addition, we further evaluate the attacks under the $p + q$ MSB leakage model and the d_p, d_q MSB leakage model introduced in Sections 4.1 and 4.2, respectively. The detailed experimental results are reported in Tables 4 and 5.

Table 4. Experimental results for $p + q$ MSB leakage.

m	N	Lattice dimension	d_p, d_q under $p + q$ MSB leakage				
			0%	5%	10%	20%	50%
4	1000 bit	15	35 bit	42 bit	49 bit	63 bit	106 bit
6	1000 bit	42	53 bit	62 bit	70 bit	87 bit	137 bit
8	1000 bit	89	64 bit	74 bit	82 bit	100 bit	154 bit

Table 5. Experimental results for d_p, d_q MSB leakage.

m	N	Lattice dimension	d_p, d_q under MSB leakage			
			0%	25%	50%	75%
4	1000 bit	17	35 bit	36 bit	37 bit	38 bit
6	1000 bit	44	53 bit	54 bit	56 bit	58 bit
8	1000 bit	91	64 bit	66 bit	68 bit	70 bit

All experiments were performed in SageMath 10.9 using its built-in LLL implementation, on a 14-inch MacBook Pro with an Apple M5 processor and 16 GB of memory.

References

1. Bauer, A., Vergnaud, D., Zapalowicz, J.C.: Inferring sequences produced by non-linear pseudorandom number generators using Coppersmith's methods. In: Public Key Cryptography – PKC 2012. Lecture Notes in Computer Science, vol. 7293, pp. 609–626. Springer (2012). https://doi.org/10.1007/978-3-642-30057-8_36
2. Bernstein, D.J., Chang, Y.A., Cheng, C.M., Chou, L.P., Heninger, N., Lange, T., van Someren, N.: Factoring rsa keys from certified smart cards: Coppersmith in the wild. In: Advances in Cryptology – ASIACRYPT 2013. Lecture Notes in Computer Science, vol. 8270, pp. 341–360. Springer (2013). https://doi.org/10.1007/978-3-642-42045-0_18
3. Bleichenbacher, D., May, A.: New attacks on RSA with small secret CRT-exponents. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.) Public Key Cryptography – PKC 2006. Lecture Notes in Computer Science, vol. 3958, pp. 1–13. Springer (2006). https://doi.org/10.1007/11745853_1
4. Blömer, J., May, A.: New partial key exposure attacks on RSA. In: Boneh, D. (ed.) Advances in Cryptology – CRYPTO 2003. Lecture Notes in Computer Science, vol. 2729, pp. 27–43. Springer, Berlin, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45146-4_2
5. Boneh, D., Durfee, G.: Cryptanalysis of RSA with private key d less than $n^{0.292}$. IEEE Transactions on Information Theory **46**(4), 1339–1349 (2000). <https://doi.org/10.1109/18.850673>
6. Coppersmith, D.: Finding a small root of a bivariate integer equation; factoring with high bits known. In: Advances in Cryptology – EUROCRYPT 1996. Lecture Notes in Computer Science, vol. 1070, pp. 178–189. Springer (1996). https://doi.org/10.1007/3-540-68339-9_16
7. Coppersmith, D.: Small solutions to polynomial equations, and low exponent RSA vulnerabilities. Journal of Cryptology **10**(4), 233–260 (1997). <https://doi.org/10.1007/s001459900030>
8. Herrmann, M., May, A.: Attacking power generators using unravelled linearization: When do we output too much? In: Advances in Cryptology – ASIACRYPT 2009. Lecture Notes in Computer Science, vol. 5912, pp. 487–504. Springer (2009). https://doi.org/10.1007/978-3-642-10366-7_29
9. Herrmann, M., May, A.: Maximizing small root bounds by linearization and applications to small secret exponent RSA. In: Nguyen, P.Q., Pointcheval, D. (eds.) Public Key Cryptography – PKC 2010. Lecture Notes in Computer Science, vol. 6056, pp. 53–69. Springer (2010). https://doi.org/10.1007/978-3-642-13013-7_4
10. Jochemsz, E., May, A.: A strategy for finding roots of multivariate polynomials with new applications in attacking RSA variants. In: Lai, X., Chen, K. (eds.) Advances in Cryptology – ASIACRYPT 2006. Lecture Notes in Computer Science, vol. 4284, pp. 267–282. Springer (2006). https://doi.org/10.1007/11935230_18
11. Jochemsz, E., May, A.: A polynomial time attack on RSA with private CRT-exponents smaller than $N^{0.073}$. In: Menezes, A. (ed.) Advances in Cryptology – CRYPTO 2007. Lecture Notes in Computer Science, vol. 4622, pp. 395–411. Springer (2007). https://doi.org/10.1007/978-3-540-74143-5_22
12. May, A.: Cryptanalysis of unbalanced RSA with small CRT-exponent. In: Yung, M. (ed.) Advances in Cryptology – CRYPTO 2002. Lecture Notes in Computer Science, vol. 2442, pp. 242–256. Springer (2002). https://doi.org/10.1007/3-540-45708-9_16
13. May, A., Nowakowski, J., Sarkar, S.: Partial key exposure attack on short secret exponent CRT-RSA. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology

- ASIACRYPT 2021, Part I. Lecture Notes in Computer Science, vol. 13090, pp. 99–129. Springer (2021). https://doi.org/10.1007/978-3-030-92062-3_4
- 14. May, A., Nowakowski, J., Sarkar, S.: Approximate divisor multiples – factoring with only a third of the secret CRT-exponents. In: Dunkelman, O., Dziembowski, S. (eds.) *Advances in Cryptology – EUROCRYPT 2022, Part III*. Lecture Notes in Computer Science, vol. 13277, pp. 147–167. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07082-2_6
- 15. Ryan, K.: Solving multivariate Coppersmith problems with known moduli. In: Fehr, S., Fouque, P.A. (eds.) *Advances in Cryptology – EUROCRYPT 2025*. Lecture Notes in Computer Science, vol. 15606, pp. 355–384. Springer (2025). https://doi.org/10.1007/978-3-031-91095-1_13
- 16. Takayasu, A., Kunihiro, N.: Partial key exposure attacks on CRT-RSA: Better cryptanalysis to full size encryption exponents. In: Malkin, T., Kolesnikov, V., Lewko, A.B., Polychronakis, M. (eds.) *Applied Cryptography and Network Security – ACNS 2015*. Lecture Notes in Computer Science, vol. 9092, pp. 518–537. Springer (2015). https://doi.org/10.1007/978-3-319-28166-7_25
- 17. Takayasu, A., Lu, Y., Peng, L.: Small CRT-exponent RSA revisited. In: Coron, J.S., Nielsen, J.B. (eds.) *Advances in Cryptology – EUROCRYPT 2017*. Lecture Notes in Computer Science, vol. 10211, pp. 130–159. Springer (2017). https://doi.org/10.1007/978-3-319-56614-6_5
- 18. Takayasu, A., Lu, Y., Peng, L.: Small CRT-exponent RSA revisited. *Journal of Cryptology* **32**(4), 1337–1382 (2019). <https://doi.org/10.1007/s00145-018-9282-3>
- 19. Wiener, M.J.: Cryptanalysis of short RSA secret exponents. *IEEE Transactions on Information Theory* **36**(3), 553–558 (1990). <https://doi.org/10.1109/18.54902>