Rounding and Chaining LLL: Finding Faster Small Roots of Univariate Polynomial Congruences

Jingguo Bi[1], Jean-Sébastien Coron[2], Jean-Charles Faugère[3,4,5], Phong Q. Nguyen[6,1], Guénaël Renault[4,3,5], and Rina Zeitoun[7,4,3,5]

[1] Tsinghua University, Institute for Advanced Study, Beijing 100084, China
jingguobi@mail.tsinghua.edu.cn
[2] University of Luxembourg
jean-sebastien.coron@uni.lu
[3] INRIA, POLSYS, Centre Paris-Rocquencourt, F-78153, Le Chesnay, France
[4] Sorbonne Universités, UPMC Univ Paris 06, Équipe POLSYS, LIP6 UPMC, F-75005, Paris, France
[5] CNRS, UMR 7606, LIP6 UPMC, F-75005, Paris, France
jean-charles.faugere@inria.fr
guenael.renault@lip6.fr
[6] INRIA, France
http://www.di.ens.fr/~pnguyen
[7] Oberthur Technologies, 420 rue d'Estienne d'Orves, CS 40008, 92705 Colombes, France
r.zeitoun@oberthur.com

Abstract. In a seminal work at EUROCRYPT '96, Coppersmith showed how to find all small roots of a univariate polynomial congruence in polynomial time: this has found many applications in public-key cryptanalysis and in a few security proofs. However, the running time of the algorithm is a high-degree polynomial, which limits experiments: the bottleneck is an LLL reduction of a high-dimensional matrix with extra-large coefficients. We present in this paper the first significant speedups over Coppersmith's algorithm. The first speedup is based on a special property of the matrices used by Coppersmith's algorithm, which allows us to provably speed up the LLL reduction by rounding, and which can also be used to improve the complexity analysis of Coppersmith's original algorithm. The exact speedup depends on the LLL algorithm used: for instance, the speedup is asymptotically quadratic in the bit-size of the small-root bound if one uses the Nguyen-Stehlé L^2 algorithm. The second speedup is heuristic and applies whenever one wants to enlarge the root size of Coppersmith's algorithm by exhaustive search. Instead of performing several LLL reductions independently, we exploit hidden relationships between these matrices so that the LLL reductions can be somewhat chained to decrease the global running time. When both speedups are combined, the new algorithm is in practice hundreds of times faster for typical parameters.

Keywords: Coppersmith's Algorithm, Small Roots of Polynomial Equations, LLL, Complexity, Speedup, RSA.

1 Introduction

At EUROCRYPT '96, Coppersmith [Cop96b,Cop96a,Cop97] showed how to find efficiently all small roots of polynomial equations (modulo an integer, or over the integers). The simplest (and perhaps most popular) result is the following: Given an integer N of unknown factorization and a monic polynomial $f(x) \in \mathbb{Z}[x]$ of degree δ, Coppersmith's lattice-based algorithm finds all integers $x_0 \in \mathbb{Z}$ such that $f(x_0) \equiv 0 \pmod{N}$ and $|x_0| \leq N^{1/\delta}$ in time polynomial in $\log N$ and δ. This has many applications in public-key cryptanalysis (*e.g.* attacking special cases of RSA and factoring with a hint), but also in a few security proofs (such as in RSA-OAEP [Sho02]). Accordingly, Coppersmith's seminal work has been followed up by dozens of articles (see May's survey [May10] for references), which introduced new variants, generalizations, simplifications and applications.

All these small-root algorithms are based on the same idea of finding new polynomial equations using lattice basis reduction: it reduces the problem of finding small roots to finding LLL-short vectors in a lattice. This can theoretically be done in polynomial time using the LLL algorithm [LLL82], but is by no means trivial in practice: the asymptotic running time is a high-degree polynomial, because the lattice is huge. More precisely, May's recent survey [May10] gives for Coppersmith's lattice-based algorithm the complexity upper bound $O(\delta^5 \log^9 N)$ using the Nguyen-Stehlé L^2 algorithm [NS09] as the reduction algorithm. A careful look gives a slightly better upper bound: asymptotically, one may take a matrix of dimension $O(\log N)$, and bit-size $O((\log^2 N)/\delta)$, resulting in a complexity upper bound $O((\log^9 N)/\delta^2)$ using L^2. In typical applications, δ is small ≤ 9 but $\log N$ is the bit-size of an RSA modulus, *i.e.* at least 1024 bits, which makes the theoretical running time daunting: $\log^9 N$ is already at least 2^{90}. For more powerful variants of Coppersmith's algorithm, the running time is even worse, because the lattice dimension and/or the bit-size increase: for instance, Coron [Cor07] gives the upper bound $O(\log^{11} W)$ for finding small roots over bivariate equations over the integers (W plays a role similar to N in the univariate congruence case), using L^2.

The bottleneck of all Coppersmith-type small-root algorithms is the LLL reduction. Despite considerable attention, no significant improvement on the running time has been found, except that LLL algorithms have improved since [Cop97], with the appearance of L^2 [NS09] and $\tilde{L}^1$ [NSV11]. And this issue is reflected in experiments (see [CNS99]): in practice, one settles for sub-optimal parameters, which means that one can only find small roots up to a bound lower than the asymptotic bound. To illustrate this point, the celebrated Boneh-Durfee attack [BD00] on RSA with short secret exponent has the theoretical bound $d \leq N^{1-1/\sqrt{2}} \approx N^{0.292}$, but the largest d in the Boneh-Durfee experiments is only $d \approx N^{0.280}$ with a 1000-bit N, and much less for larger N, *e.g.* $d \approx N^{0.265}$ for 4000-bit N.

OUR RESULTS. We present two speedups over Coppersmith's algorithm for finding small roots of univariate polynomial congruences, which can be combined in practice.

The first speedup is provable and depends on the LLL algorithm used: if one uses L^2 [NS09], the total bit-complexity is upper bounded by $O(\log^7 N)$, which gives a speedup $\Theta((\log^2 N)/\delta^2)$ quadratic in the bit-size of the small-root bound $N^{1/\delta}$; and if one uses $\tilde{L}^1$, the total complexity is upper bounded by $O(\log^{6+\varepsilon} N)$ for any $\varepsilon > 0$ using fast integer arithmetic, which gives a speedup $O((\log N)/\delta)$ linear in the bit-size of the small-root bound $N^{1/\delta}$. This speedup comes from combining LLL reduction with rounding: instead of LLL-reducing directly a matrix with huge entries, we suitably round the coefficients before LLL reduction to make them much smaller, and show that the LLL output allows to derive sufficiently short vectors in the original lattice. In practice, this means that for any instantiation of Coppersmith's algorithm achieving a small-root bound X, we can drastically reduce the size of the coefficients of the matrix to be LLL-reduced and achieve essentially the same small-root bound: asymptotically, the bit-size is reduced by a factor $(\log N)/\delta$, which implies that the speedup is quadratic when using the popular L^2 algorithm, or quasi-linear using the more theoretical $\tilde{L}^1$ algorithm. This rounding strategy is very natural, but it is folklore that it fails in the worst case: when an arbitrary non-singular matrix is rounded, it may even become singular, and the situation is worse for LLL reduction. However, we show that a well-chosen rounding strategy surprisingly works for the special matrices used by Coppersmith's algorithm: this is because the matrices to be reduced are triangular matrices whose diagonal entries are reasonably balanced, which can be exploited. Interestingly, this peculiar property can also be used to improve the complexity upper bound of Coppersmith's original algorithm, without changing the algorithm: if one uses $\tilde{L}^1$ [NSV11], one can obtain the same complexity upper bound as in our rounding-based algorithm, up to constants.

Our second speedup is heuristic and applies whenever one wants to enlarge the root size X of Coppersmith's algorithm by exhaustive search: it is well-known that any root size X can be extended to mX by applying m times the algorithm on "shifted" polynomials. This enlargement

is necessary when one wants to go beyond Coppersmith's bound $N^{1/\delta}$, but it is also useful to optimize the running time below $N^{1/\delta}$: beyond a certain root size below $N^{1/\delta}$, it is folklore that it is faster to use exhaustive search than Coppersmith's algorithm with larger parameters. In this setting, one applies Coppersmith's algorithm with the same modulus N but different polynomials which are all "shifts" of the initial polynomial $f(x)$: $f_t(x) = f(X \cdot t + x)$ for varying t, where $0 \leqslant t < N^{1/\delta}/X$. We show that this creates hidden relationships between the matrices to be LLL reduced, which can be exploited in practice: instead of performing LLL reductions independently of say, matrices B_1 and B_2, we chain the LLL reductions. More precisely, after LLL reducing B_1 into a reduced basis C_1, we reduce a matrix of the form $C_1 \cdot P$ for some well-chosen matrix P, instead of the matrix B_2. And this process can be iterated to drastically reduce the global running time.

When both speedups are combined, the new algorithm is in practice hundreds of times faster for typical parameters. Finally, our work helps to clarify the asymptotic complexity of Coppersmith's algorithm for univariate polynomial congruences. Despite the importance of the algorithm, it seems that the dependence on the polynomial degree δ was not well-understood: as previously mentioned, May's survey [May10] gave an upper bound including a factor δ^5, and Coppersmith's journal article [Cop97] gave an upper bound growing exponentially in δ. Our final complexity upper bound is independent of δ: it only depends on the bit-size of the modulus N.

Surprisingly, our improvements only apply for now to Coppersmith's algorithm for finding all small roots of univariate polynomial equations, and not to more sophisticated variants such as the gcd generalization used for factoring with a hint. This seems to be the first significant difference between Coppersmith's algorithm and its gcd generalization. It is an interesting open problem to obtain significant speedup for other small-root algorithms.

RELATED WORK. Our first speedup is based on rounding. Rounding has been used in lattice reduction before: for instance, Buchmann [Buc94] used rounding to rigorously estimate when a computation with real lattices can be alternatively performed using integer bases; and the $\tilde{L}^1$ [NSV11] algorithm is also based on rounding. However, it seems that none of the previous work identified the special structure of Coppersmith matrices which we exploit. Our second speedup is based on chaining. Chaining has also been used in lattice reduction before, *e.g.* in the MIMO context [NJD11], but our technique and analysis seem to be a bit different. Thus, both rounding and chaining are folklore strategies, but our work seems to be their first application to Coppersmith's algorithm.

ROADMAP. In Sect. 2, we recall background on lattices and Coppersmith's small-root algorithm. In Sect. 3, we present and analyze our first speedup of Coppersmith's algorithm: rounding LLL. In Sect. 4, we present and analyze our second speedup of Coppersmith's algorithm: chaining LLL. In Sect. 5, we present experimental results with both speedups. Finally, we discuss the case of other small-root algorithms in Sect. 6.

2 Background and Notation

We use row representation for matrices: vectors are row vectors denoted by bold lowercase letters, matrices are denoted by uppercase letters, and their coefficients are denoted by lowercase letters. All logarithms are in base 2. Let $\|\|$ and $\langle , \rangle$ be the Euclidean norm and inner product of $\mathbb{R}^n$. The Euclidean norm is naturally extended to polynomials as follows: if $f(x) = \sum_{i=0}^{n} f_i x^i \in \mathbb{R}[x]$, then $\|f\| = (\sum_{0 \leq i \leq n} f_i^2)^{1/2}$. We use the following matrix norms: if $M = (m_{i,j})$ is an $n \times m$ matrix, then $\|M\|_2 = \max_{\|\mathbf{x}\| \neq 0} \frac{\|\mathbf{x}M\|}{\|\mathbf{x}\|}$, and $\|M\|_\infty = \max_{1 \leq j \leq m} \sum_{i=1}^{n} |m_{i,j}|$. Then: $\|M\|_2 \leq \sqrt{n}\|M\|_\infty$. If $x \in \mathbb{R}$, we denote by $\lceil x \rfloor$ a closest integer to x.

2.1 Lattices

LATTICES. A lattice L is a discrete subgroup of $\mathbb{R}^m$: there exist $n(\leq m)$ linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^m$ s.t. L is the set $\mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of all integral linear combinations of the $\mathbf{b}_i$'s, *i.e.*

$$\mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_n) = \left\{ \sum_{i=1}^{n} x_i \mathbf{b}_i : x_i \in \mathbb{Z} \right\}.$$

Then the matrix $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ is called a *basis* of L and n is the *rank* (or *dimension*) of L. Here, we mostly consider full-rank lattices, *i.e.* $n = m$. The (co-)volume of L is $\mathrm{vol}(L) = \sqrt{\det(BB^t)}$ for any basis B of L, where B^t denotes B's transpose. If B is square, then $\mathrm{vol}(L) = |\det B|$, and if B is further triangular, then $\mathrm{vol}(L)$ is simply the product of the diagonal entries of B in absolute value.

GRAM-SCHMIDT ORTHOGONALIZATION. Let $\mathbf{b}_1, \cdots, \mathbf{b}_n \in \mathbb{R}^m$ be linearly independent vectors. The Gram-Schmidt orthogonalization (GSO) is the family $(\mathbf{b}_1^\star, \ldots, \mathbf{b}_n^\star)$ defined recursively as: $\mathbf{b}_1^\star = \mathbf{b}_1$ and for $i \geq 2$, $\mathbf{b}_i^\star$ is the component of the vector $\mathbf{b}_i$ which is orthogonal to the linear span of $\mathbf{b}_1, \cdots, \mathbf{b}_{i-1}$. Then $\mathbf{b}_i^\star = \mathbf{b}_i - \sum_{j=1}^{i-1} \mu_{i,j} \mathbf{b}_j^\star$, where $\mu_{i,j} = \langle \mathbf{b}_i, \mathbf{b}_j^\star \rangle / \|\mathbf{b}_j^\star\|^2$ for $1 \leq j < i \leq n$.

SIZE-REDUCTION. A basis $B = (\mathbf{b}_1, \cdots, \mathbf{b}_n)$ is *size-reduced* if its GSO satisfies $|\mu_{i,j}| \leq 1/2$, for all $1 \leq j < i \leq n$. There is a classical (elementary) algorithm which size-reduces a basis $(\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of an integer lattice $L \subseteq \mathbb{Z}^m$, in polynomial time, without ever modifying the Gram-Schmidt vectors $\mathbf{b}_i^\star$: this algorithm is included in the original LLL algorithm [LLL82] (e.g. it is the sub-algorithm `RED` in the description of LLL in [Coh93]). In the special case that the input basis is (square) lower-triangular, the running-time of this size-reduction algorithm is $O(n^3 b^2)$ without fast integer arithmetic, and $n^3 \tilde{O}(b)$ using fast-integer arithmetic, where $b = \max_{1 \leq i \leq n} \log \|\mathbf{b}_i\|$.

LLL AND SHORT LATTICE VECTORS. Coppersmith's small-root method requires the ability to efficiently find reasonably short vectors in a lattice, namely a non-zero vector $\mathbf{v} \in L$ s.t. $\|\mathbf{v}\| \leq c^n \mathrm{vol}(L)^{1/n}$ where c is some constant and n is the lattice rank. This can be achieved by the celebrated LLL algorithm [LLL82]: given a basis $(\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of an integer lattice $L \subseteq \mathbb{Z}^m$, LLL outputs a non-zero $\mathbf{v} \in L$ s.t. $\|\mathbf{v}\| \leq 2^{\frac{n-1}{4}} \mathrm{vol}(L)^{1/n}$ in time $O(n^5 m b^3)$ (resp. $n^3 m b \tilde{O}(n) \tilde{O}(b)$) without (resp. with) fast integer arithmetic, where $b = \max_{1 \leq i \leq n} \log \|\mathbf{b}_i\|$: strictly speaking, this vector is actually the first vector of the basis output by the algorithm. Nguyen and Stehlé [NS09] introduced the L^2 algorithm, a faster variant of LLL which can output similarly short vectors in time $O(n^4 m (n + b) b)$ (resp. $n^2 m (n + b) b \tilde{O}(n)$) without (resp. with) fast integer arithmetic. The recent $\tilde{L}^1$ algorithm by Novocin *et al.* [NSV11] can output similarly short vectors for a full-rank lattice in time $O(n^{5+\varepsilon} b + n^{\omega+1+\varepsilon} b^{1+\varepsilon})$ for any $\varepsilon > 0$ using fast integer arithmetic, where $\omega \leq 2.376$ is the matrix multiplication complexity constant. However, this algorithm is considered to be mostly of theoretical interest for now: $\tilde{L}^1$ is currently not implemented anywhere, as opposed to L^2. When assessing the complexity of LLL reduction, it is therefore meaningful to mention two complexities: one (closer to the real world) using L^2 without fast integer arithmetic, and another using $\tilde{L}^1$ using fast integer arithmetic and fast linear algebra.

The complexity upper bound of LLL reduction can sometimes be decreased by some polynomial factor. In particular, when the Gram-Schmidt norms of the input basis are balanced, the LLL algorithm requires fewer loop iterations than in the worst case. More precisely, [DV94, Th. 1.1] showed that the classical upper bound $O(n^2 b)$ on the number of iterations can be replaced by $O\left(n^2 \log \frac{max\|\mathbf{b}_i^\star\|}{min\|\mathbf{b}_i^\star\|}\right)$.

2.2 Coppersmith's method for finding small roots

At EUROCRYPT '96, Coppersmith [Cop96b,Cop96a,Cop97] showed how to find efficiently all small roots of polynomial equations (modulo an integer, or multivariate over the integers),

which is surveyed in [May10,Ngu09]. We now review the simplest result, following the classical Howgrave-Graham approach [HG97]: In Sect. 6, we will discuss the main variants of this result.

Theorem 1 (Coppersmith [Cop96b,Cop97]). *There is an algorithm (Alg. 1) which, given as input an integer N of unknown factorization and a monic polynomial $f(x) \in \mathbb{Z}[x]$ of degree δ and coefficients in $\{0, \ldots, N-1\}$, outputs all integers $x_0 \in \mathbb{Z}$ such that $f(x_0) \equiv 0 \, (mod \, N)$ and $|x_0| \leq N^{1/\delta}$ in time polynomial in $\log N$ and δ.*

It is straightforward that when the polynomial is of degree one or when this degree is sufficiently large in comparison to N (it is sufficient that $\delta + 1 \geq (\log N)/2$ in our analysis) then a direct approach or a brute force search of the solutions provides a comparable complexity result. In a similar way, if $\delta > \log(N)$, then we have $N^{1/\delta} < 1$, which means that the only possible root is zero. Thus, in the sequel we implicitly consider polynomials not in these cases, that is to say verifying $2 < \delta + 1 < (\log N)/2$. It is also important to note that Coppersmith's algorithm (Alg. 1) does not directly achieve the bound $N^{1/\delta}$: indeed, it finds efficiently all roots up to some bound $X \, (< N^{1/\delta})$ depending on an integer parameter $h \geq 2$, chosen asymptotically to be $h = O((\log N)/\delta)$. When h is sufficiently large, then X becomes sufficiently close to $N^{1/\delta}$ so that one can find all roots up to $N^{1/\delta}$. However, it is well-known that the bound $X = N^{1/\delta}$ should not be reached by taking such a large h. Instead, it is faster to use a smaller h, and perform exhaustive search on the most significant bits of the solutions, as depicted in Alg. 1 (see Section 4 for more details).

Algorithm 1 Coppersmith's Method

Input: Two integers $N \geq 1$ and $h \geq 2$, a univariate degree-δ monic polynomial $f(x) \in \mathbb{Z}[x]$ with coefficients in $\{0, \ldots, N-1\}$ and $2 < \delta + 1 < (\log N)/2$.
Output: All $x_0 \in \mathbb{Z}$ s.t. $|x_0| \leq N^{1/\delta}$ and $f(x_0) \equiv 0 \mod N$.
1: Let $n = h\delta$, X the bound given in (3), and $t = 0$.
2: **while** $Xt < N^{1/\delta}$ **do**
3:      $f_t(x) = f(Xt + x) \in \mathbb{Z}[x]$.
4:      Build the $n \times n$ lower-triangular matrix B whose rows are the $g_{i,j}(xX)$'s defined by (1).
5:      Run the L^2 algorithm [NS09] on the matrix B.
6:      The first vector of the reduced basis corresponds to a polynomial of the form $v(xX)$ for some $v(x) \in \mathbb{Z}[x]$.
7:      Compute all the roots x_0' of the polynomial $v(x) \in \mathbb{Z}[x]$ over $\mathbb{Z}$.
8:      Output $x_0 = Xt + x_0'$ for each root x_0' which satisfies $f_t(x_0') \equiv 0 \, (\text{mod } N)$ and $|x_0'| \leq X$.
9:      $t \leftarrow t + 1$.
10: **end while**

We now explain Coppersmith's algorithm (Alg. 1). The core idea consists in reducing the problem to solving univariate polynomial equations over the integers, by transforming modular roots into integral roots. More precisely, it constructs a polynomial $g(x) \in \mathbb{Z}[x]$ such that: if $x_0 \in \mathbb{Z}$ is such that $f(x_0) \equiv 0 \, (\text{mod } N)$ and $|x_0| \leq X$, then $g(x_0) = 0$ and can be solved easily over $\mathbb{Z}$. To do so, it uses the following elementary criterion:

Lemma 1 (Howgrave-Graham [HG97]). *Let $g(x) \in \mathbb{Z}[x]$ be a polynomial with at most n non-zero coefficients. Let M be an integer ≥ 1. Assume that $\|g(xX)\| < \frac{M}{\sqrt{n}}$ for some $X \in \mathbb{R}$. If $x_0 \in \mathbb{Z}$ is such that $g(x_0) \equiv 0 \, (mod \, M)$ and $|x_0| \leq X$, then $g(x_0) = 0$.*

Lemma 1 will be used with $M = N^{h-1}$ and $g(x)$ found by lattice reduction. Let $h \geq 2$ be an integer and define the following family of $n = h\delta$ polynomials:

$$g_{i,j}(x) = (x)^j N^{h-1-i} f^i(x) \quad 0 \leq i < h, 0 \leq j < \delta \tag{1}$$

These n polynomials satisfy: if $f(x_0) \equiv 0 \, (\text{mod } N)$ for some $x_0 \in \mathbb{Z}$, then $g_{i,j}(x_0) \equiv 0 \, (\text{mod } N^{h-1})$. In order to apply Lemma 1 for a bound $X \geq 1$ to be determined later, Coppersmith's algorithm

constructs the n-dimensional lattice L spanned by the rows of the $n \times n$ matrix B formed by the n coefficient vectors of $g_{i,j}(xX)$, where the polynomials are ordered by increasing degree (*e.g.* in the order $(i,j) = (0,0), (0,1), \cdots, (0, \delta - 1), (1, 0), \cdots (h - 1, \delta - 1)$) and the coefficients are ordered by increasing monomial degree: the first coefficient is thus the constant term of the polynomial. The matrix B is lower triangular, and its n diagonal entries are:

$$\left(N^{h-1}, N^{h-1}X, \ldots, N^{h-1}X^{\delta-1}, \ldots, N^0 X^{\delta h - \delta}, \ldots, N^0 X^{\delta h - 2}, N^0 X^{\delta h - 1} \right), \qquad (2)$$

because $f(x)$ is monic. In other words, the exponent of X increases by one at each row, while the exponent of N decreases by one every δ rows. It follows that $\mathrm{vol}(L) = \det(B) = N^{\frac{1}{2}n(h-1)}X^{\frac{1}{2}n(n-1)}$. Alg. 1 applies the LLL algorithm to the matrix B, which provides a non-zero polynomial $v(x) \in \mathbb{Z}[x]$ such that $\|v(xX)\| \leq 2^{\frac{n-1}{4}}\mathrm{vol}(L)^{\frac{1}{n}} = 2^{\frac{n-1}{4}}N^{\frac{h-1}{2}}X^{\frac{n-1}{2}}$. It follows that the polynomial $v(x)$ satisfies Lemma 1 with $M = N^{h-1}$ and $g(x) = v(x)$ if

$$X \leq \frac{1}{\sqrt{2}}N^{\frac{h-1}{n-1}}(n+1)^{-\frac{1}{n-1}}. \qquad (3)$$

The dimension of B is $n = h\delta$, and the entries of the matrix B have bit-size $O(h \log N)$, therefore the running time of L^2 in Step 5 without fast integer arithmetic is $O(\delta^6 h^7 \log N + \delta^5 h^7 \log^2 N)$, which is $O(\delta^5 h^7 \log^2 N)$ because $\delta + 1 < (\log N)/2$, and the running time of $\tilde{L}^1$ in Step 5 is $O(h^{6+\varepsilon}\delta^{5+\varepsilon}\log N + h^{\omega+2+2\varepsilon}\ \delta^{\omega+1+\varepsilon}\log^{1+\varepsilon} N)$ for any $\varepsilon > 0$ using fast integer arithmetic, where $\omega \leq 2.376$ is the matrix multiplication complexity constant. We obtain the following concrete version of Th. 1:

Corollary 2 *Alg. 1 of Th. 1 with $h = \lfloor \log N/\delta \rfloor$ and $X = \lfloor 2^{-1/2}N^{\frac{h-1}{n-1}}(n+1)^{-\frac{1}{n-1}} \rfloor$ runs in time $O((\log^9 N)/\delta^2)$ without fast integer arithmetic using L^2 in Step 5, or $O((\log^{7+\varepsilon} N)/\delta)$ for any $\varepsilon > 0$ using fast integer arithmetic and $\tilde{L}^1$ in Step 5.*

Proof. As a first step, we emphasize that the bottleneck of Coppersmith's algorithm is the reduction step (Step 5). Indeed, let ℓ be the maximal bit-size of the coefficients of $v(x) \in \mathbb{Z}[x]$ in Step 7: we know that $\ell \leq h \log N$, and the degree of $v(x)$ is $\leq n$. Then Step 7 can be performed in time $O(n^3(\ell + \log n)) = O(h \log^4 N) = O((\log^5 N)/\delta)$ using Schönhage's root isolation algorithm [Sch82, Sec. 5.2]. Hence, the cost of Step 7 is less than the one of Step 5. As a second step, we show that Alg. 1 performs a constant number of loop iterations. Indeed, consider the bound $X = \lfloor 2^{-1/2}N^{\frac{h-1}{n-1}}(n+1)^{-\frac{1}{n-1}} \rfloor$ achieved by one loop of Alg. 1. By definition, Alg. 1 performs at most $O(N^{1/\delta}/X)$ loop iterations. Recall that we have $2 < \delta + 1 < (\log N)/2$, since in the other cases a direct approach provides a much better complexity. We have $O(X) = O(N^{\frac{h-1}{n-1}})$ and:

$$N^{1/\delta}/N^{\frac{h-1}{n-1}} = N^{1/\delta - \frac{h-1}{h\delta - 1}} = N^{\frac{\delta-1}{\delta(h\delta-1)}} \leq N^{1/(h\delta-1)}.$$

If $h = \lfloor \log N/\delta \rfloor$, then we have $h\delta - 1 \geq \log N - \delta - 1 \geq (\log N)/2$. Therefore, we conclude that $N^{1/\delta}/N^{\frac{h-1}{n-1}} = O(1)$.

Furthermore, we already proved the correctness and the running time for L^2 and $\tilde{L}^1$ algorithms in Step 5. Thus, from the analysis preceding Cor. 2, the asymptotic complexity of Coppersmith's algorithm is the one of one call to LLL (L^2 or $\tilde{L}^1$), with $h = \lfloor \log N/\delta \rfloor$. $\qquad \square$

We will later see that the complexity upper bounds of Cor. 2 with L^2 and $\tilde{L}^1$ can actually be decreased. Indeed, we will uncover a special property of Coppersmith's matrix (see Lemma 2), which implies that $O\left(\frac{max\|\mathbf{b}_i^\star\|}{min\|\mathbf{b}_i^\star\|}\right) = O(N)$, so that the number of loop iterations $O\left(n^2 \log \frac{max\|\mathbf{b}_i^\star\|}{min\|\mathbf{b}_i^\star\|}\right)$ on the input basis used by Coppersmith's algorithm is $O(n^2 \log N)$ instead of the all-purpose bound

$O(n^2 h \log N)$ [DV94]. By taking this observation into account, the upper bounds $O((\log^8 N)/\delta)$ and $O(\log^{6+\varepsilon} N)$ are respectively achieved for the L^2 and $\tilde{L}^1$ algorithms. In the sequel, we present another method improving Cor. 2, based on the same special property of Coppersmith's matrix, and which can be easily implemented.

3 Speeding up Coppersmith's Algorithm by Rounding

Our first main result is the following speedup over Coppersmith's algorithm:

Theorem 3. *There is an algorithm (namely, Alg. 2) which, given as input an integer N of unknown factorization and a monic polynomial $f(x) \in \mathbb{Z}[x]$ of degree δ and coefficients in $\{0, \ldots, N-1\}$, outputs all integers $x_0 \in \mathbb{Z}$ such that $f(x_0) \equiv 0 \pmod{N}$ and $|x_0| \leq N^{1/\delta}$ in time $O(\log^7 N)$ without fast integer arithmetic using the L^2 algorithm [NS09], or $O(\log^{6+\varepsilon} N)$ for any $\varepsilon > 0$ using fast integer arithmetic and the $\tilde{L}^1$ algorithm [NSV11] in Step 7.*

As explained in the beginning of Section 2, we recall that only the non trivial case of polynomial of degree δ with $2 < \delta + 1 < (\log N)/2$ will be implicitly considered in the sequel.

3.1 Rounding for Coppersmith's Algorithm

The bottleneck of Coppersmith's algorithm is the LLL reduction of the matrix B, whose dimension is $n = h\delta$, and whose entries have bit-size $O(h \log N)$. Asymptotically, we have $h = O(\log N/\delta)$ so the dimension is $O(\log N)$ and the bit-size is $O((\log^2 N)/\delta)$. We will modify Coppersmith's algorithm in such a way that we only need to LLL-reduce a matrix of the same dimension but with much smaller entries, namely bit-length $O(\log N)$.

To explain the intuition behind our method, let us first take a closer look at the matrix B and uncover some of its special properties:

Lemma 2. *Let $X \leq N^{1/\delta}$. The maximal diagonal coefficient of matrix B defined in Step. 4 of Alg. 1 is $N^{h-1} X^{\delta-1} < N^h$, the minimal diagonal coefficient is $X^{h\delta-\delta} \leq N^{h-1}$, and $\frac{N^{h-1} X^{\delta-1}}{X^{h\delta-\delta}} \geq N^{1-1/\delta}$ if $h \geq 2$. Furthermore, if $X \geq \Omega(N^{\frac{h-1}{n-1}})$, $h \geq 2$ and $h\delta = O(\log N)$ then we have:*

$$X^{h\delta-\delta} \geq N^{h-O(1)}. \tag{4}$$

Proof. The $n = h\delta$ diagonal coefficients of B are naturally split into h blocks of δ coefficients: the i-th block is formed by the leading coefficients of the polynomials $g_{i,j}(xX)$ for $0 \leq j < \delta$. Since the leading coefficient of $g_{i,j}(xX)$ is $X^j N^{h-1-i} X^{\delta i}$, it follows that the maximal and minimal coefficients in the i-th block are located respectively at the end and at the beginning: their values are respectively $X^{\delta(i+1)-1} N^{h-1-i} = N^{h-1}(X^\delta/N)^i X^{\delta-1}$ and $N^{h-1-i} X^{\delta i} = N^{h-1}(X^\delta/N)^i$. If $X \leq N^{1/\delta}$, we obtain that the maximal diagonal coefficient is $N^{h-1} X^{\delta-1}$ reached in the 0-th block, and the minimal diagonal coefficient is $X^{(h-1)\delta}$ reached in the $(h-1)$-th block. And the ratio $\frac{N^{h-1} X^{\delta-1}}{X^{h\delta-\delta}}$ is exactly $N^{h-1}/X^{h\delta-2\delta+1}$ which is clearly $\geq N^{1-1/\delta}$ if $h \geq 2$.
Now, let $X_0 = N^{\frac{h-1}{n-1}}$ so that $X = \Omega(X_0)$. We have $N^{1/\delta}/N^{\frac{h-1}{n-1}} \leq N^{1/(h\delta-1)}$ by the proof of Cor. 2. Therefore we have $X_0 \geq N^{1/\delta-1/(h\delta-1)} = N^{(h\delta-1-\delta)/(\delta(h\delta-1))}$. Hence $X_0^\delta \geq N^{(h\delta-1-\delta)/(h\delta-1)} = N^{1-\delta/(h\delta-1)}$ and thus $X_0^{h\delta-\delta}$ is such that $X_0^{h\delta-\delta} \geq N^{(h-1)-\delta(h-1)/(h\delta-1)} > N^{h-2}$. Since $X = \Omega(X_0)$ and $h\delta = O(\log N)$, we obtain (4). $\square$

This implies that the diagonal coefficients of B are somewhat balanced: the matrix B is not far from being reduced. In fact, the first row of B has norm N^{h-1} which is extremely close to the bound $N^{h-1}/\sqrt{n}$ required by Lemma 1: intuitively, this means that it should not be too difficult to find a lattice vector shorter than $N^{h-1}/\sqrt{n}$.

To take advantage of the structure of B, we first size-reduce B to make sure that the subdiagonal coefficients are smaller than the diagonal coefficients. Then we round the entries of B so that the smallest diagonal coefficient becomes $\lfloor c \rfloor$ where $c > 1$ is a parameter. More precisely, we create a new $n \times n$ triangular matrix $\tilde{B} = (\tilde{b}_{i,j})$ defined by:

$$\tilde{B} = \left\lfloor cB/X^{h\delta - \delta} \right\rceil \tag{5}$$

By Lemma 2, we have:

$$b_{i,i} \geq X^{h\delta - \delta} \quad \text{and} \quad \tilde{b}_{i,i} \geq \lfloor c \rfloor . \tag{6}$$

We LLL-reduce the rounded matrix $\tilde{B}$ instead of B: let $\tilde{\mathbf{v}} = \mathbf{x}\tilde{B}$ be the first vector of the reduced basis obtained. If we applied to B the unimodular transformation that LLL-reduces $\tilde{B}$, we may not even obtain an LLL-reduced basis in general. However, because of the special structure of B, it turns out that $\mathbf{v} = \mathbf{x}B$ is still a short non-zero vector of L, as shown below:

Lemma 3. *Let $B = (b_{i,j})$ be an $n \times n$ lower-triangular matrix over $\mathbb{Z}$ with strictly positive diagonal. Let $c > 1$. If $\tilde{B} = \lfloor cB/\min_{i=1}^{n} b_{i,i} \rceil$ and $\mathbf{x}\tilde{B}$ is the first vector of an LLL-reduced basis of $\tilde{B}$, then:*

$$0 < \|\mathbf{x}B\| < \left(n\|\tilde{B}^{-1}\|_2 + 1 \right) 2^{\frac{n-1}{4}} \det(B)^{\frac{1}{n}} .$$

Proof. Let $\alpha = \min_{i=1}^{n} b_{i,i}/c$, so that $\tilde{B} = \lfloor B/\alpha \rceil$. Define the matrix $\bar{B} = \alpha\tilde{B}$ whose entries are $\bar{b}_{i,j} = \alpha\tilde{b}_{i,j}$. Then $0 \leq b_{i,j} - \bar{b}_{i,j} < \alpha$, therefore $\|B - \bar{B}\|_2 < n\alpha$. We have:

$$\|\mathbf{x}B\| \leq \|\mathbf{x}(B - \bar{B})\| + \|\mathbf{x}\bar{B}\| \leq \|\mathbf{x}\| \times \|B - \bar{B}\|_2 + \alpha\|\mathbf{x}\tilde{B}\| < n\|\mathbf{x}\|\alpha + \alpha\|\mathbf{x}\tilde{B}\|.$$

Let $\tilde{\mathbf{v}} = \mathbf{x}\tilde{B}$. Then $\|\mathbf{x}\| \leq \|\tilde{\mathbf{v}}\|\|\tilde{B}^{-1}\|_2$, and we obtain $\|\mathbf{x}B\| < \left(n\|\tilde{B}^{-1}\|_2 + 1 \right) \alpha\|\tilde{\mathbf{v}}\|$. The matrix $\tilde{B}$ is lower-triangular with all diagonal coefficients strictly positive because $c > 1$. Since $\tilde{\mathbf{v}} = \mathbf{x}\tilde{B}$ is the first vector of an LLL-reduced basis of $\tilde{B}$, and $\tilde{B}$ is non-singular, $\mathbf{x}B \neq 0$ and we have:

$$\alpha\|\tilde{\mathbf{v}}\| \leq \alpha 2^{\frac{n-1}{4}} \det(\tilde{B})^{\frac{1}{n}} = 2^{\frac{n-1}{4}} \det(\bar{B})^{\frac{1}{n}} \leq 2^{\frac{n-1}{4}} \det(B)^{\frac{1}{n}},$$

where we used the fact that matrices $\tilde{B}$, $\bar{B}$ and B are lower-triangular. The result follows by combining both inequalities. $\qquad\square$

If $\mathbf{x}B$ is sufficiently short, then it corresponds to a polynomial of the form $v(xX)$ for some $v(x) \in \mathbb{Z}[x]$ satisfying Lemma 1, and the rest proceeds as in Alg. 1. The whole rounding algorithm is given in Alg. 2, which will be shown to admit a lower complexity upper-bound than Alg. 1 to compute all roots up to $N^{1/\delta}$.

Algorithm 2 Coppersmith's Method with Rounding

Input: Two integers $N \geq 1$ and $h \geq 2$, a univariate degree-δ monic polynomial $f(x) \in \mathbb{Z}[x]$ with coefficients in $\{0, \ldots, N-1\}$ and $2 < \delta + 1 < (\log N)/2$.
Output: All $x_0 \in \mathbb{Z}$ s.t. $|x_0| \leq N^{1/\delta}$ and $f(x_0) \equiv 0 \mod N$.
1: Let $n = h\delta$, X the bound given in Th. 4, $c = (3/2)^n$ and $t = 0$.
2: **while** $Xt < N^{1/\delta}$ **do**
3:     $f_t(x) = f(Xt + x) \in \mathbb{Z}[x]$.
4:     Build the $n \times n$ matrix B whose rows are the $g_{i,j}(xX)$'s defined by (1).
5:     Size-reduce B without modifying its diagonal coefficients.
6:     Compute the matrix $\tilde{B} = \lfloor cB/X^{h\delta - \delta} \rceil$ obtained by rounding B.
7:     Run the L^2 algorithm [NS09] on the matrix $\tilde{B}$.
8:     Let $\tilde{\mathbf{v}} = \mathbf{x}\tilde{B}$ be the first vector of the reduced basis obtained.
9:     The vector $\mathbf{v} = \mathbf{x}B$ corresponds to a polynomial of the form $v(xX)$ for some $v(x) \in \mathbb{Z}[x]$.
10:     Compute all the roots x_0' of the polynomial $v(x) \in \mathbb{Z}[x]$ over $\mathbb{Z}$.
11:     Output $x_0 = x_0' + Xt$ for each root x_0' which satisfies $f_t(x_0') \equiv 0 \pmod{N}$ and $|x_0'| \leq X$.
12:     $t \leftarrow t + 1$.
13: **end while**

We now justify the bound X given in Alg. 2. In order for Lemma 3 to be useful, we need to exhibit an upper bound for $\|\tilde{B}^{-1}\|_2$, as given in the following lemma; we provide the proof in Appendix B.

Lemma 4. *Let $B = (b_{i,j})$ be an $n \times n$ size-reduced lower-triangular matrix over $\mathbb{Z}$ with strictly positive diagonal. Let $c > 1$. If $\tilde{B} = \lfloor cB/\min_{i=1}^{n} b_{i,i}\rfloor$, then:*

$$\|\tilde{B}^{-1}\|_2 \leq \sqrt{n} \left(\frac{3c-2}{2c-2}\right)^{n-1} / \lfloor c\rfloor.$$

By combining Lemmas 3 and 4, we obtain the following small-root bound X for Alg. 2:

Theorem 4. *Given as input two integers $N \geq 1$ and $h \geq 2$, a rational $c > 1$, and a univariate degree-δ monic polynomial $f(x) \in \mathbb{Z}[x]$ with coefficients in $\{0, \ldots, N-1\}$, one loop of Alg. 2, corresponding to $t < N^{1/\delta}/X$, outputs all $x_0 = Xt + x_0' \in \mathbb{Z}$ s.t. $|x_0'| \leq X$ and $f(x_0) = 0 \mod N$, where $n = h\delta$ and*

$$X = \left\lfloor \frac{N^{\frac{h-1}{n-1}} \kappa_1^{-2/(n-1)}}{\sqrt{2}\, n^{1/(n-1)}} \right\rfloor \quad with \quad \kappa_1 = n^{3/2} \left(\frac{3c-2}{2c-2}\right)^{n-1} \lfloor c\rfloor^{-1} + 1. \tag{7}$$

Proof. Combining Lemma 4 with Lemma 3 where $\det(B)^{1/n} = N^{\frac{h-1}{2}} X^{\frac{n-1}{2}}$, we get $0 < \|\mathbf{x}B\| < \kappa_1 2^{\frac{n-1}{4}} N^{\frac{h-1}{2}} X^{\frac{n-1}{2}}$. Furthermore, Lemma 1 is satisfied with $M = N^{h-1}$ and $v(xX)$ corresponding to $\mathbf{x}B$ if $\|\mathbf{x}B\| \leq N^{h-1}/\sqrt{n}$. This gives the following condition on X: $\kappa_1 2^{\frac{n-1}{4}} N^{\frac{h-1}{2}} X^{\frac{n-1}{2}} \leq N^{h-1}/\sqrt{n}$. The bound given in (7) is naturally derived from this last inequality. □

The bound X of Th. 4 is never larger than that of Cor. 2. However, if one selects $c \geq (3/2)^n$, then the two bounds are asymptotically equivalent. This is why Alg. 2 uses $c = (3/2)^n$.

3.2 Running time: proof of Theorem 3

The original matrix B had entries whose bit-size was $O(h \log N)$. Let $\beta = \frac{N^h X^{\delta-1}}{X^{h\delta-\delta}}$ be the ratio between the maximal diagonal coefficient and the minimal diagonal coefficient of $\tilde{B}$. If B is size-reduced, the entries of the new matrix $\tilde{B} = \lfloor cB/X^{h\delta-\delta}\rfloor$ are upper bounded by $c\beta$.

By Lemma 2, we know that if $h \geq 2$, then $\beta \geq N^{1-1/\delta}$, and if further $X \geq \Omega(N^{\frac{h-1}{n-1}})$ and $h\delta = O(\log N)$, then $\beta = N^{O(1)}$. Hence, the bit-size of $\tilde{B}$'s entries is $\leq \log c + O(\log N)$. And the dimension of $\tilde{B}$ is the same as B, *i.e.* $h\delta$. It follows that the running time of L^2 in Step 7 is $O(\delta^6 h^6(\log c + \log N) + \delta^5 h^5(\log c + \log N)^2)$ without fast integer arithmetic, which is $O(\delta^5 h^5(\log c + \log N)^2$ because $\delta < (\log N)/2 - 1$, and is $O((h\delta)^{5+\varepsilon}(\log c + \log N) + (h\delta)^{\omega+1+\varepsilon}(\log c + \log N)^{1+\varepsilon})$ for any $\varepsilon > 0$ using fast integer arithmetic and $\tilde{L}^1$ in Step 7, where $\omega \leq 2.376$ is the matrix multiplication complexity constant.

This leads to our main result (Th. 3), a variant of Coppersmith's algorithm with improved complexity upper bound. More precisely, as in Coppersmith's algorithm, one can easily prove that the number of loops performed in Alg. 2 is at most constant. Indeed, when $c = (3/2)^n$, then $\kappa_1^{\frac{-2}{n-1}}$ converges to 1. This means that the bound X achieved by Th. 4 is asymptotically equivalent to the one achieved by Cor. 2, which completes the proof of Th. 3, because $\log c = O(\log N)$ when $c = (3/2)^n$.

Remark: Surprisingly, Lemma 2 also allows to prove that the $\tilde{L}^1$algorithm, when carefully analyzed using the balancedness of the Gram-Schmidt norms, already achieves the complexity bound $O(\log^{6+\varepsilon} N)$ given in Th. 3. Indeed, using Th. 6 from [NSV11] which gives the $\tilde{L}^1$complexity upper bound $O(n^{3+\varepsilon}\tau) = O(\log^{3+\varepsilon} N\tau)$ where τ is the total number of iterations, and combining it with

[DV94] applied to Coppersmith's matrix (Lemma 2), which gives $\tau = O(n^2 \log N) = O(\log^3 N)$, allows to retrieve the above complexity $O(\log^{6+\varepsilon} N)$. However, we propose in this paper a direct improvement of Coppersmith's method based on elementary tools and which can therefore be easily implemented on usual computer algebra systems (*e.g.* Sage, Magma, NTL) with immediate practical impact on cryptanalyses. Furthermore, we are not aware of any implementation of the $\tilde{L}^1$ algorithm for the time being, which makes a practical comparison tricky.

In the sequel, we present a method that allows to speed up the exhaustive search which is performed to reach Coppersmith's bound $N^{1/\delta}$.

4   Chaining LLL

As recalled in Section 2.2, in order to find all solutions which are close to the bound $N^{1/\delta}$, one should not use a very large lattice dimension (*i.e.* $n = O(\log N)$). Instead, it is better to use a lattice of reasonable dimension and to perform exhaustive search on the most significant bits of x until finding all solutions. Namely, we consider polynomials $f_t(x) = f(X \cdot t + x)$ where $0 \leqslant t < \frac{N^{1/\delta}}{X}$ and $X = \lfloor 2^{\frac{-1}{2}} N^{\frac{h-1}{n-1}} (n+1)^{-\frac{1}{n-1}} \rfloor$. Thus, an initial solution x_0 that can be written $x_0 = X \cdot t_0 + x_0'$ is obtained by finding the solution x_0' of the polynomial f_{t_0}. In this case, this solution satisfies $|x_0'| < X$ and it has a correct size for LLL to find it using a lattice of dimension n. For each polynomial f_t, one runs LLL on a certain matrix (Step 4 of Alg. 2).

In Section 4.1, we describe a method that allows to take advantage of the LLL performed for the case $t = i$ to reduce (in practice) the complexity of the LLL performed for the case $t = i+1$. The method is based on a hidden relationship between Coppersmith's lattices. Thereafter, in Section 4.2 we combine this improvement with the rounding approach described in Section 3.

4.1   Exploiting Relations Between Consecutive Lattices

The following proposition discloses a connection between the lattice used for the case $t = i$ and the next lattice used for $t = i + 1$. This connection is based on the well-known *Pascal matrix* $P = (p_{s,t})$ defined as the $n \times n$ lower-triangular matrix whose non-zero coefficients are the binomials: $p_{s,t} = \binom{s}{t}$ for $0 \leq t \leq s \leq n - 1$.

Proposition 1. *Let B be a basis of the n-dimensional lattice used by Coppersmith's algorithm to find all small roots of the polynomial $f_i(x) = f(X \cdot i + x)$, where X is the small-root bound. Then $B \cdot P$ is a basis of the "next" lattice used for the polynomial $f_{i+1}(x)$.*

Proof. Because all lattice bases are related by some unimodular matrix, it suffices to prove the statement for a special basis B. We thus only consider the special basis $B = B_i$ formed by the n shifted polynomials constructed from $f_i(x)$ and written in the basis $\mathcal{B} = (1, xX^{-1}, (xX^{-1})^2, \ldots, (xX^{-1})^{n-1})$. For the case $t = i + 1$, one tries to solve the polynomial

$$f_{i+1}(x) = f(X \cdot (i+1) + x) = f(X \cdot i + x + X) = f_i(x + X).$$

Therefore, the shifted polynomials constructed from f_{i+1} are the same as for the case $t = i$, but written in the different basis $\mathcal{B}' = (1, xX^{-1} + 1, (xX^{-1} + 1)^2, \ldots, (xX^{-1} + 1)^{n-1})$. Yet, we need to return to the original representation of the polynomials, *i.e.* in the basis $\mathcal{B}$. To this end, we use the following property regarding the lower triangular Pascal matrix P:

$$\mathcal{B}'^T = P \cdot \mathcal{B}^T.$$

As a consequence, left-multiplying each side of this equality by the matrix B_i proves that the matrix $B_i \cdot P$ is a basis of the lattice used for finding small roots of the polynomial $f_{i+1}(x)$. □

The proposition allows us to use different matrices to tackle the polynomial $f_{i+1}(x)$ than the one initially used by Coppersmith's method. In particular, we can use a matrix of the form $B^R \cdot P$ where B^R is an LLL-reduced basis of the previous lattice used to solve $f_i(x)$: intuitively, it might be faster to LLL-reduce such matrices than the initial Coppersmith's matrix. Although we are unable to prove the lattice reduction will be faster, we can show that the vectors of such a matrix are not much longer than that of B^R:

Corollary 5 *Let B_i^R be the LLL-reduced matrix used for solving f_t for $t = i$ and P be the Pascal matrix. The matrix*

$$B_{i+1} = B_i^R \cdot P$$

spans the same lattice used for solving the case $t = i + 1$. This matrix consists of vectors $\mathbf{b_{i+1},j}$ whose norms are close to vector norms of the LLL-reduced matrix B_i^R. Namely, for all $1 \leqslant j \leqslant n$ we have:

$$||\mathbf{b_{i+1},j}|| < \sqrt{n} \cdot 2^{n-1} \cdot ||\mathbf{b_i^R},j|| \, .$$

In particular, for the case $i = t_0$ the first vector of B_{i+1} has a norm bounded by $2^{n-1} \cdot N^{h-1}$.

Proof. The previous proposition immediately gives the first statement. Since the matrix B_{i+1} is the product of B_i^R with a matrix P composed of relatively short elements, the elements in B_{i+1} remain close to those in the reduced matrix B_i^R. Indeed, the largest element in P is $\binom{n-1}{\frac{n}{2}-1}$ which is less than 2^{n-2}. More precisely, the maximal norm of column vectors in P is reached in the $(\frac{n}{2} - 1)$-th column and is smaller than $\sqrt{2^{2 \cdot (\frac{n-2}{2}-1)} + 2^{2 \cdot (\frac{n-2}{2})} + \cdots + 2^{2 \cdot (n-2)}} < \sqrt{2^{2n-3}} < 2^{n-1}$. Therefore the norm of each row vector of B_{i+1} is at most enlarged by a factor $\sqrt{n} \cdot 2^{n-1}$ compared to the norm of the corresponding vector in B_i^R, *i.e.* for all $1 \leqslant j \leqslant n$ we have $||\mathbf{b_{i+1},j}|| < \sqrt{n} \cdot 2^{n-1} \cdot ||\mathbf{b_i^R},j||$. In particular, for $i = t_0$, since the first vector of B_i^R has a norm bounded by $N^{h-1}/\sqrt{n}$, the norm of the first vector of $B_i^R \cdot P$ is bounded by $2^{n-1} \cdot N^{h-1}$ which is relatively close to $N^{h-1}/\sqrt{n}$. $\square$

Cor. 5 shows us that vectors of B_{i+1} are relatively close to the ones in the LLL-reduced matrix B_i^R. Thus, we intuitively expect the LLL-reduction of B_{i+1} to be less costly than the one of the original Coppersmith's matrix. However, our bounds are too weak to rigorously prove this. Yet, one can use this property iteratively to elaborate a new method which *chains* all LLL reductions as follows. First, one LLL-reduces B_0 for the case $t = 0$. This gives a reduced matrix B_0^R. Then, one iterates this process by performing LLL reduction on $B_{i+1} = B_i^R \cdot P$ (for $i \geq 0$) to obtain B_{i+1}^R and so forth until all solutions are found (each time by solving the polynomial corresponding to the first vector of B_i^R).

In the sequel, we study this *chaining method* by performing similar roundings as in Section 3 before each call of LLL reduction.

4.2 Rounding and Chaining LLL

During the exhaustive search described in Section 4.1, we perform the LLL algorithm on the matrix $B_{i+1} = B_i^R \cdot P$ for $0 \leqslant i < N^{1/\delta}/X$, where B_i^R is LLL-reduced. It is worth noticing that the structure of B_i^R and thereby of B_{i+1}, is different from the original Coppersmith's matrix B_0 (in particular, it is not triangular anymore). Yet, we are able to show that under certain conditions on B_{i+1} verified experimentally, one can combine the rounding technique of Section 3 with the chaining technique of Section 4.1. Indeed, we show that during the chaining loop, one can size-reduce B_{i+1} and then round its elements for all $i \geq 0$ as follows:

$$\tilde{B}_{i+1} = \left\lfloor cB_{i+1} / \min_{1 \leq i \leq n} ||\mathbf{b}_i^\star|| \right\rfloor , \tag{8}$$

where $\mathbf{b}_i^\star$ are Gram-Schmidt vectors of B_{i+1} and c is a rational that will be determined later. Then, one applies LLL on the rounded matrix $\tilde{B}_{i+1}$ as performed in Section 3. We obtain an LLL-reduced matrix $\tilde{B}_{i+1}^R$ and a unimodular matrix $\tilde{U}_{i+1}$ such that $\tilde{U}_{i+1} \cdot \tilde{B}_{i+1} = \tilde{B}_{i+1}^R$. Then one shows that by applying $\tilde{U}_{i+1}$ on B_{i+1}, the first vector of this matrix $\tilde{U}_{i+1} \cdot B_{i+1}$ is a short vector that allows to find the solutions provided that they are smaller than a bound X that will be determined latter. For the sake of clarity, in the sequel we denote by B the matrix B_{i+1}, and by $\mathbf{x}B$, the first vector of matrix $\tilde{U}_{i+1} \cdot B_{i+1}$. We would like to exhibit an upper-bound on $\|\mathbf{x}B\|$. To this end, we will need, as in Section 3, to upper-bound the value $\|\tilde{B}^{-1}\|_2$. This is done in the following lemma:

Lemma 5. *Let $B = (b_{i,j})$ be an $n \times n$ non-singular integral matrix and $\alpha \geq 1$ such that $n\alpha\|B^{-1}\|_2 < 1$. Then the matrix $\tilde{B} = \lfloor B/\alpha \rceil$ is invertible with:*

$$\|\tilde{B}^{-1}\|_2 \leq \frac{\alpha\|B^{-1}\|_2}{1 - n\alpha\|B^{-1}\|_2} \, .$$

As one can see, this value depends on $\|B^{-1}\|_2$ which is given in Lemma 6. Both Lemmas 5 and 6 are proven in Annex.

Lemma 6. *Let B be an $n \times n$ non-singular size-reduced matrix, with Gram-Schmidt vectors $\mathbf{b}_i^\star$. Then:*

$$\|B^{-1}\|_2 \leq \frac{\sqrt{n}(3/2)^{n-1}}{\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\|} \, .$$

One can now give an upper-bound on $\|\mathbf{x}B\|$:

Corollary 6 *Let $B = (b_{i,j})$ be an $n \times n$ size-reduced non-singular matrix over $\mathbb{Z}$. Let $\alpha \geq 1$ such that $n^2\alpha\|B^{-1}\|_2 < 1$. Then $\tilde{B} = \lfloor cB/\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\| \rceil = \lfloor B/\alpha \rceil$ is non-singular. And if $\mathbf{x}\tilde{B}$ is the first vector of an LLL-reduced basis of $\tilde{B}$, then:*

$$0 \; < \; \|\mathbf{x}B\| \; < \; \kappa_2 \; 2^{\frac{n-1}{4}} \; \det(B)^{\frac{1}{n}},$$

where $\kappa_2 = \dfrac{c^{\frac{n+1}{n}}}{(c - n^{3/2}(3/2)^{n-1})(c - n^{5/2}(3/2)^{n-1})^{1/n}} \, .$

Sketch of proof. The proof follows Lemma 3 proof. The major differences being that the bounds on $\|B^{-1}\|_2$ and $\|\tilde{B}^{-1}\|_2$ differ (see Lemmas 5 and 6), and that $\det(\tilde{B})$ is not straightforwardly predictable anymore since matrix B is no more triangular. For a complete proof, see Annex F. $\qquad\qquad\square$

Again, if $\|\mathbf{x}B\|$ is sufficiently short, then it corresponds to a polynomial of the form $v(xX)$ for some $v(x) \in \mathbb{Z}[x]$ satisfying Lemma 1. In particular, for the case $t = t_0$, solving this polynomial equation would allow to retrieve the solution x_0. Note that the condition $n^2\alpha\|B^{-1}\|_2 < 1$ specified in Cor. 6 gives a condition on the rational c. Indeed, since $\alpha = \min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\|/c$ and using Lemma 6, one gets:

$$n^2\alpha\|B^{-1}\|_2 \leqslant n^2 \frac{\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\|}{c} \frac{\sqrt{n}(3/2)^{n-1}}{\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\|} \leqslant \frac{n^{5/2}(3/2)^{n-1}}{c} < 1 \, ,$$

that is c should be such that $c > n^{5/2}(3/2)^{n-1}$. The whole chaining and rounding algorithm is depicted in Algorithm 3. Note that in practice, we do not need to perform Step 8 of Alg. 3 and that $\min_{1 \leq i \leq n} \|\mathbf{b}_{t+1}{}_i^\star\|$ can be estimated instead of being computed in Step 9 (see Section 4.3 for more details).

In the following, we give a small-root bound X on the solution x_0' sufficient to guarantee success:

Algorithm 3 Coppersmith's Method with Chaining and Rounding

Input: Two integers $N \geq 1$ and $h \geq 2$, a univariate degree-δ monic polynomial $f(x) \in \mathbb{Z}[x]$ with coefficients in $\{0, \ldots, N-1\}$ and $2 < \delta + 1 < (\log N)/2$.
Output: All $x_0 \in \mathbb{Z}$ s.t. $|x_0| \leq N^{1/\delta}$ and $f(x_0) \equiv 0 \mod N$.

1: Perform Step 1 and Steps 3 to 7 of Alg. 2. Step 7 returns $\tilde{B}_0^R$ and $\tilde{U}_0$ such that $\tilde{U}_0 \cdot \tilde{B}_0 = \tilde{B}_0^R$.
2: Let $n = h\delta$, X the bound given in Th. 7, $c = n^{\frac{5}{2}}(\frac{3}{2})^n$, $t = 0$, P is the $n \times n$ lower triangular Pascal matrix.
3: Compute the matrix $\tilde{U}_0 \cdot B_0$, where B_0 is the matrix computed in Step 5 of Alg. 2.
4: The first vector of $\tilde{U}_0 \cdot B_0$ corresponds to a polynomial of the form $v(xX)$ for some $v(x) \in \mathbb{Z}[x]$.
5: Compute and output all roots $x_0 \in \mathbb{Z}$ of $v(x)$ satisfying $f(x_0) \equiv 0 \pmod{N}$ and $|x_0| \leq X$.
6: **while** $Xt < N^{1/\delta}$ **do**
7:      Compute the matrix $B_{t+1} = \tilde{U}_t \cdot B_t \cdot P$.
8:      Size-reduce B_{t+1}.
9:      Compute the matrix $\tilde{B}_{t+1} = \lfloor cB_{t+1} / \min_{1 \leq i \leq n} \|\mathbf{b_{t+1}}_i^\star\| \rfloor$ obtained by rounding B_{t+1}.
10:     Run L^2 algorithm on matrix $\tilde{B}_{t+1}$ which returns $\tilde{B}_{t+1}^R$ and $\tilde{U}_{t+1}$ s.t. $\tilde{U}_{t+1} \cdot \tilde{B}_{t+1} = \tilde{B}_{t+1}^R$.
11:     Compute the matrix $\tilde{U}_{t+1} \cdot B_{t+1}$.
12:     The first vector of $\tilde{U}_{t+1} \cdot B_{t+1}$ corresponds to a polynomial of the form $v(xX)$.
13:     Compute all the roots x_0' of the polynomial $v(x) \in \mathbb{Z}[x]$ over $\mathbb{Z}$.
14:     Output $x_0 = x_0' + Xt$ for each root x_0' which satisfies $f(x_0' + Xt) \equiv 0 \pmod{N}$ and $|x_0'| \leq X$.
15:     $t \leftarrow t + 1$.
16: **end while**

Theorem 7. *Given as input two integers $N \geq 1$ and $h \geq 2$, a rational $c > n^{5/2}(3/2)^{n-1}$, and a univariate degree-δ monic polynomial $f(x) \in \mathbb{Z}[x]$ with coefficients in $\{0, \ldots, N-1\}$, one loop of Alg. 3, corresponding to $t < N^{1/\delta}/X$, outputs all $x_0 = Xt + x_0' \in \mathbb{Z}$ s.t. $|x_0'| \leq X$ and $f(x_0) = 0$ mod N, and $n = h\delta$, where*

$$X = \left\lfloor \frac{N^{\frac{h-1}{n-1}} \kappa_2^{\frac{-2}{n-1}}}{\sqrt{2}\, n^{1/(n-1)}} \right\rfloor, \tag{9}$$

and κ_2 is the value defined in Cor. 6.

Proof. During the chaining-LLL method, all matrices B satisfy the property $\det(B)^{1/n} = N^{\frac{h-1}{2}} X^{\frac{n-1}{2}}$, since the Pascal matrix has determinant equal to 1. Furthermore, Lemma 1 is satisfied with $M = N^{h-1}$ and $v(xX)$ corresponding to $\mathbf{x}B$ if $\|\mathbf{x}B\| \leqslant N^{h-1}/\sqrt{n}$. Therefore, using Corollary 6, one gets the condition: $\kappa_2\, 2^{\frac{n-1}{4}}\, N^{\frac{h-1}{2}} X^{\frac{n-1}{2}} \leq N^{h-1}/\sqrt{n}$, and the bound given in (9) follows. $\qquad\square$

The bound X of Th. 7 is never larger than that of Cor. 2. However, if one selects $c > n^{5/2}(3/2)^{n-1}$, then the two bounds are asymptotically equivalent. This is why Alg. 3 uses $c = n^{5/2}(3/2)^n$.

4.3 Complexity Analysis: A Heuristic Approach

The complexity of Alg. 3 relies on the complexity of the LLL-reduction performed in Step 10. The cost of this reduction depends on the size of coefficients in matrix $B = \tilde{B}_{t+1}$, which itself depends on the value $\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\|$. The exact knowledge of this value does not seem straightforward to obtain without computing the Gram-Schmidt matrix explicitly. However, experiments show that the Gram-Schmidt curve is roughly decreasing, *i.e.* $\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\| \approx \|\mathbf{b}_n^\star\|$ and is roughly symmetric: *i.e.* $\log \|\mathbf{b}_i^\star\| - \log \|\mathbf{b}_{n/2}^\star\| \approx \log \|\mathbf{b}_{n/2}^\star\| - \log \|\mathbf{b}_{n-i+1}^\star\|$. If we assume these two experimental facts, we deduce that $\|\mathbf{b}_{n/2}^\star\| \approx |\det(B)|^{1/n}$. By duality, this means that $\|\mathbf{b}_n^\star\| \approx |\det(B)|^{2/n}/\|\mathbf{b}_1^\star\|$. Furthermore, from the definition of the GSO, we know that $\|\mathbf{b}_1^\star\| = \|\mathbf{b}_1\|$, where $\mathbf{b}_1$ is the first vector of matrix B. Therefore we have:

$$\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\| \approx \|\mathbf{b}_n^\star\| \approx |\det(B)|^{2/n}\|\mathbf{b}_1^\star\|^{-1} = N^{h-1}X^{n-1}\|\mathbf{b}_1\|^{-1}, \tag{10}$$

Thus, we need an estimation on $\|\mathbf{b}_1\|$. Since in practice matrix $B = B_{i+1} = \tilde{U}_i \cdot B_i \cdot P$ is already nearly size-reduced, one can skip Step 8 of Alg. 3. Therefore, vector $\mathbf{b}_1$ is the first vector of matrix $\tilde{U}_i \cdot B_i \cdot P$. Using Cor. 6, one deduces that the first vector of matrix $\tilde{U}_i \cdot B_i$ is roughly as short as the first vector of an LLL-reduced matrix. From the well-known experimental behavior of LLL [NS06], we can model the first vector of the LLL-reduced basis as a "random" vector of norm $\approx 1.02^n |\det(B)|^{1/n}$ (where 1.02 has to be replaced by a smaller constant for dimension $n \leq 60$) . Since the Pascal matrix P has a norm smaller than 2^{n-1} (see proof of Cor. 5), one gets the bound $\|\mathbf{b}_1\| \leqslant \sqrt{n} 2^{n-1} 1.02^n |\det(B)|^{1/n}$. Therefore, we deduce that: $\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\| \approx |\det(B)|^{1/n}/(\sqrt{n} 2^{n-1} 1.02^n)$. In practice, we conjecture that $\min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\| > |\det(B)|^{1/n}/\beta^n$, where $\beta < 2$ (see Fig. 5 in Sec. 5).

This discussion leads to the following heuristic approach regarding the method: firstly, one should rather use the estimation (10) in Step 9 of Alg. 3, instead of explicitly computing the Gram-Schmidt matrix; secondly, one can skip Step 8 of Alg. 3. This heuristic version of Algorithm 3 is the one we used during our experiments, all these assumptions were always verified.

To conclude our analysis, it suffices to reduce a rounded matrix such that $\max_{1 \leq i \leq n} \|\tilde{\mathbf{b}}_i^\star\| \leq c \max_{1 \leq i \leq n} \|\mathbf{b}_i^\star\| / \min_{1 \leq i \leq n} \|\mathbf{b}_i^\star\| \leqslant c\beta^{2n}$, instead of $\max_{1 \leq i \leq n} \|\tilde{\mathbf{b}}_i^\star\| \leq \beta^n |\det(B)|^{1/n}$. This means that we are trading entries of size $O(n)$. Therefore, by considering $n = O(\log N)$, we obtain the same complexity as in Theorem 3 but in a heuristic way. However, even if both asymptotic complexities are identical, in practice for reasonable dimensions the speed-up brought by using Alg. 3 rather than Alg. 2 is considerable (see Section 5). Indeed, the LLL-reduction of matrix $\tilde{U}_i \cdot B_i \cdot P$ (Step 10 of Alg. 3) performs surprisingly faster than expected. This comes from the fact that for reasonable dimensions, the Gram-Schmidt curve of this matrix remains quite close to the one of matrix $\tilde{U}_i \cdot B_i$, where $\tilde{U}_i \cdot B_i$ turns out to be LLL-reduced (or nearly). Besides, the overall running-time of Alg. 3 is approximately the time spent to perform one LLL-reduction, multiplied by the number of executed loops, i.e. by $N^{1/\delta}/X$.

5   Experiments

We implemented Coppersmith's algorithm (Alg. 1) and our improvements (Algs. 2 and 3) using Shoup's NTL library [Sho]. However, for the LLL reduction, we used the fplll implementation [CPS08] by Cadé et al., which includes the L^2 algorithm [NS09]: fplll is much faster than NTL for Coppersmith's matrices. It should be stressed that fplll is a wrapper which actually implements several variants of LLL, together with several heuristics: L^2 is only used as a last resort when heuristic variants fail. This means that there might be a discrepancy between the practical running time and the theoretical complexity upper bound of LLL routines. Our test machine is a 2.93-GHz Intel Core 2 Duo processor E7500 running on Fedora. Running times are given in seconds.

Like in [CNS99], we used the case $\delta = 3$, and N an RSA-type modulus: the exact polynomial congruence is derived from RSA encryption with public exponent δ. Then, one loop of Alg. 1, with $n = 3h$, can find all roots x_0 as long as $|x_0'| \leq X = \lfloor 2^{-1/2} N^{\frac{h-1}{n-1}} n^{-\frac{1}{n-1}} \rfloor$. For a fixed h, the rounding strategy (Alg. 2) gives a worse bound than X, but the difference can be made arbitrarily small by increasing the parameter c: in our experiments, we therefore chose the smallest value of c such that $\kappa_1^{\frac{-2}{n-1}}$ and $\kappa_2^{\frac{-2}{n-1}}$ are larger than 0.90, so that the new bound is never less than the old bound X by more than 10%, which is essentially the same. However, we note that the value c can be taken smaller in practice: indeed, our theoretical analysis was a worst-case analysis. For instance, it has been proved in [VT98] that if T is a random $n \times n$ lower-triangular matrix with unit diagonal and subdiagonal coefficients normally distributed, then $(\||T^{-1}\|_2)^{1/n}$ converges to $1.3057\ldots$ And experimentally, when subdiagonal coefficients are uniformly distributed over $[-1/2, +1/2]$, then we have $(\||T^{-1}\|_\infty)^{1/n} \leq 1.1$ with high probability. This means that the constants of Lemma 7 (and therefore the implicit $3/2$ in formulas for c) are better in practice.

Furthermore, it is worth noticing that since the value α is not significant in itself, in order to increase the efficiency, one can round matrices at negligible cost by taking $\alpha := 2^{\lfloor \log_2(\alpha) \rfloor}$ and performing shifts of $\lfloor \log_2(\alpha) \rfloor$ bits. In the same vein, one can increment t by 2 instead of 1 in Step 9 of Alg. 1 or in Step 12 of Alg. 2, and one can multiply the matrix $\tilde{U}_i \cdot B_i$ by P^2 instead of P in Step 7 of Alg. 3. This comes from the fact that if $0 < x'_0 < X$ (resp. $-X < x'_0 < 0$), then $x'_0 - X$ (resp. $x'_0 + X$) is also a valid solution. This refinement allows to divide by 2 the global timing of Algs. 1 and 2. However, it seems to be much less relevant when applied to Alg. 3.

Figures 1 and 2 summary our limited experiments respectively comparing one loop of Algs. 1 and 2, and one loop of Algs. 1 and 3 in practice. They provide the bit-length of X and the corresponding running times of the lattice reduction only, because the cost of solving a univariate equation over $\mathbb{Z}$ turns out to be much less in practice. Running times are given as averages over 5 samples. For a typical case where $\lceil \log N \rceil = 2048$, the whole Alg. 1 would perform in $((\lfloor 2048/3 - 666 \rfloor)/2) \times 6431.2 \approx 6.7$ years and the new Alg. 3 would perform in $(\lfloor 2048/3 - 666 \rfloor) \times 15.52 \approx 11.8$ days, which is about 207 times faster (see Fig. 2 and 6).

**Fig. 1.** Bounds and running time of rounding method for cubic congruences

Size of N	Data type	Parameter h				
		10	15	20	25	30
	Size of X	318	324	328	331	332
	$T_{original}$	2.54	30.48	216.3	793.4	3720.8
1024	$T_{rounded}$	0.68	4.49	18.22	48.17	175.9
	Speed-up	3.74	6.79	11.87	16.47	21.16
	Size of X	634	650	658	663	666
	$T_{original}$	13.47	150.7	865.7	3078	10146.7
2048	$T_{rounding}$	3.14	17.79	63.3	166.4	379.8
	Speed-up	4.29	8.40	13.67	18.50	26.72
	Size of X	1270	1302	1318	1327	1333
	$T_{original}$	41.45	582.6	3162	11968	42053
4096	$T_{rounded}$	7.07	43.25	157.5	449.8	1301.5
	Speed-up	5.86	13.47	20.07	26.61	32.31

Fig. 2. Bounds and running time of rounding plus chaining method for cubic congruences

Size of N	Data type	Parameter h				
		10	15	20	25	30
	Size of X	316	323	327	330	332
	$T_{original}$	2.14	23.55	161.55	646.37	1955.1
1024	T_{rc}	0.04	0.42	1.71	5.56	12.71
	$Speed-up_{rc}$	53.5	56.07	94.47	116.25	153.83
	Size of X	633	649	657	663	666
	$T_{original}$	8.21	95.12	641.22	2299.5	6431.2
2048	T_{rc}	0.07	0.55	2.39	7.75	15.52
	$Speed-up_{rc}$	117.28	172.95	268.29	296.71	414.38
	Size of X	1270	1302	1318	1327	1333
	$T_{original}$	27.64	378.62	2226	8303.2	25813
4096	T_{rc}	0.11	0.87	3.73	11.72	29.65
	$Speed-up_{rc}$	251.27	435.19	596.78	708.46	870.6

Fig. 3. Speed-up of rounding method

Fig. 4. Speed-up of rounding plus chaining method

From Figure 3, we see that we already get significant speedups (say, larger than 10) even for small values of h and typical sizes of N, by using the rounding method (Alg. 2). The speedup grows when $\log N$ or h grows: for fixed N, the speedup grows roughly a bit less than quadratically in h, whereas the theoretical analysis gives a speedup linear in h. From Figure 4, we see that we can obtain more speedups as the sizes of N or h increase, by using the rounding and chaining method (Alg. 3). Hence, our improvement is practical and allows to get much closer to the asymptotic small-root bound.

Furthermore, we verify the assumption on the value $\min_{1\leq i\leq n}\|b_i^\star\|$ for matrix B. We write $\max_{1\leq i\leq n}\|\mathbf{b}_i^\star\| \approx \beta_1^n \mathrm{vol}(L)^{1/n}$ and $\min_{1\leq i\leq n}\|\mathbf{b}_i^\star\| \approx \beta_2^n \mathrm{vol}(L)^{1/n}$. In this paper, we have assumed that $\beta_1 = 1/\beta_2$. We summarize the results of our experiments for $\lceil \log N \rceil = 512$ with dimensions $30, 60, 90, 120, 150$ in Table 5. We can see that $\beta_1 \times \beta_2 \approx 1$ and that $\beta_1 \leq 2$. This means our assumptions are reasonable.

Fig. 5. Beta values for $\lceil \log N \rceil = 512$

Data	Parameter h				
type	10	20	30	40	50
β_1	1.7582	1.8751	1.9093	1.9218	1.9435
β_2	0.5460	0.5271	0.5155	0.5091	0.5077
product	0.9600	0.9883	0.9842	0.9785	0.9867

Fig. 6. Timings comparisons for the total method

	log N		
	1024	2048	4096
Original	5.8 days	6.7 years	1757782 years
Alg. 3	1.8 hours	11.8 days	4038 years
Speed-up	77	207	435

6 Other Small-Root Algorithms

We now discuss whether our rounding method can similarly speed up other small-root algorithms (see the surveys [May10,Ngu09]), which are based on the same main ideas where LLL reduction plays a crucial role. In theory, the rounding method provides a speedup for any triangular matrix whose diagonal coefficients are all large. However, in order to have a large speedup, we need the minimal diagonal coefficient to be much larger than the ratio between the maximal diagonal coefficient and the minimal diagonal coefficient. In Coppersmith's algorithm, the smallest diagonal coefficient was $N^{h-O(1)}$, while the gap was $N^{O(1)}$, which translated into a polynomial speedup. It turns out that other small-root algorithms do not share the same features: we only get a (small) constant speedup. We leave it as an open problem to obtain polynomial (non-constant) speedups for these other small-root algorithms: this might be useful to make practical attacks on certain fully-homomorphic encryption schemes (see [CH11]). See Appendix G for a further discussion on these generalizations.

References

[BD00] Dan Boneh and Glenn Durfee. Cryptanalysis of RSA with private key d less than $N^{0.292}$. *IEEE Transactions on Information Theory*, 46(4):1339, 2000.

[BDHG99] Dan Boneh, Glenn Durfee, and Nick Howgrave-Graham. Factoring $n = p^r q$ for large r. In *Advances in Cryptology - Proc. CRYPTO '99*, volume 1666 of *Lecture Notes in Computer Science*, pages 326–337. Springer, 1999.

[Buc94] Johannes Buchmann. Reducing lattice bases by means of approximations. In *Algorithmic Number Theory – Proc. ANTS-I*, volume 877 of *Lecture Notes in Computer Science*, pages 160–168. Springer, 1994.

[CH11] Henry Cohn and Nadia Heninger. Approximate common divisors via lattices. *IACR Cryptology ePrint Archive*, 2011:437, 2011.

[CNS99] Christophe Coupé, Phong Q. Nguyen, and Jacques Stern. The effectiveness of lattice attacks against low-exponent RSA. In *Public Key Cryptography – Proc. PKC '99*, volume 1560 of *Lecture Notes in Computer Science*, pages 204–218. Springer, 1999.

[Coh93] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993.

[Cop96a] Don Coppersmith. Finding a small root of a bivariate integer equation; factoring with high bits known. In *Advances in Cryptology - Proc. EUROCRYPT '96*, volume 1070 of *Lecture Notes in Computer Science*, pages 178–189. Springer, 1996.

[Cop96b] Don Coppersmith. Finding a small root of a univariate modular equation. In *Advances in Cryptology - Proc. EUROCRYPT '96*, volume 1070 of *Lecture Notes in Computer Science*, pages 155–165. Springer, 1996.

[Cop97]   Don Coppersmith. Small solutions to polynomial equations, and low exponent RSA vulnerabilities. *J. Cryptology*, 10(4):233–260, 1997. Journal version of [Cop96b,Cop96a].

[Cor07]   Jean-Sébastien Coron. Finding small roots of bivariate integer polynomial equations: A direct approach. In *Advances in Cryptology – Proc. CRYPTO '07*, volume 4622 of *Lecture Notes in Computer Science*, pages 379–394. Springer, 2007.

[CPS08]   David Cadé, Xavier Pujol, and Damien Stehlé. FPLLL library, version 3.0. Available from `http://perso.ens-lyon.fr/damien.stehle`, Sep 2008.

[DV94]   Hervé Daudé and Brigitte Vallée. An upper bound on the average number of iterations of the lll algorithm. *Theor. Comput. Sci.*, 123(1):95–115, 1994.

[GAKK93]   S. K. Godunov, A. G. Antonov, O. P. Kiriljuk, and V. I. Kostin. *Guaranteed accuracy in numerical linear algebra*, volume 252 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1993. Translated and revised from the 1988 Russian original.

[GVL13]   Gene H. Golub and Charles F. Van Loan. *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, fourth edition, 2013.

[HG97]   Nick Howgrave-Graham. Finding small roots of univariate modular equations revisited. In *Cryptography and Coding – Proc. IMA '97*, volume 1355 of *Lecture Notes in Computer Science*, pages 131–142. Springer, 1997.

[HG01]   Nick Howgrave-Graham. Approximate integer common divisors. In *Proc. CaLC '01*, volume 2146 of *Lecture Notes in Computer Science*, pages 51–66. Springer, 2001.

[LLL82]   A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Ann.*, 261:513–534, 1982.

[May10]   A. May. Using LLL-reduction for solving RSA and factorization problems: A survey. 2010. In [NV10].

[Ngu09]   P. Q. Nguyen. Public-key cryptanalysis. In I. Luengo, editor, *Recent Trends in Cryptography*, volume 477 of *Contemporary Mathematics*. AMS–RSME, 2009.

[NJD11]   H. Najafi, M.E.D. Jafari, and M.-O. Damen. On adaptive lattice reduction over correlated fading channels. *Communications, IEEE Transactions on*, 59(5):1224–1227, 2011.

[NS06]   Phong Q. Nguyen and Damien Stehlé. LLL on the average. In *Algorithmic Number Theory – Proc. ANTS*, LNCS, pages 238–256. Springer, 2006.

[NS09]   P. Q. Nguyen and D. Stehlé. An LLL algorithm with quadratic complexity. *SIAM J. of Computing*, 39(3):874–903, 2009.

[NSV11]   Andrew Novocin, Damien Stehlé, and Gilles Villard. An LLL-reduction algorithm with quasi-linear time complexity: extended abstract. In *Proc. STOC '11*, pages 403–412. ACM, 2011.

[NV10]   P. Q. Nguyen and B. Vallée, editors. *The LLL Algorithm: Survey and Applications*. Information Security and Cryptography. Springer, 2010.

[Sch82]   Arnold Schönhage. *The fundamental theorem of algebra in terms of computational complexity - preliminary report*. Universität Tübingen, 1982.

[Sho]   Victor Shoup. Number Theory C++ Library (NTL) version 5.4.1. Available at `http://www.shoup.net/ntl/`.

[Sho02]   Victor Shoup. OAEP reconsidered. *J. Cryptology*, 15(4):223–249, 2002.

[VT98]   D. Viswanath and L. N. Trefethen. Condition numbers of random triangular matrices. *SIAM J. Matrix Anal. Appl.*, 19(2):564–581 (electronic), 1998.

A   Lemma 7

An upper bound on inverses of triangular matrices can be derived from the following elementary lemma. This upper bound will be used in Appendices B and D.

Lemma 7. *Let $t > 0$ and $T = (t_{i,j})$ be an $n \times n$ lower-triangular matrix over $\mathbb{R}$, with unit diagonal (i.e. $t_{i,i} = 1$ for $1 \le i \le n$), and such that $|t_{i,j}| \le t$ for $1 \le j < i \le n$. Then $\|T^{-1}\|_\infty \le (1+t)^{n-1}$.*

Proof. Let $S = T^{-1}$. Then for $1 \le i, j \le n$: $\sum_{k=j}^{n} s_{i,k} t_{k,j} = \delta_{i,j}$, where $\delta_{i,j}$ is Kronecker's symbol. Therefore $s_{i,j} = \delta_{i,j} - \sum_{k=j+1}^{n} s_{i,k} t_{k,j}$, which implies that S is lower-triangular and for $1 \le j < i \le n$:

$$|s_{i,j}| \le t \left(1 + \sum_{k=j+1}^{i-1} |s_{i,k}| \right). \tag{11}$$

Let us prove that for all $j < i$, $|s_{i,j}| \le t(1+t)^{i-j-1}$, by induction over $i - j$. Since $|t_{i,j}| \le t$ for $1 \le j < i \le n$: $|s_{i,i-1}| \le t|s_{i,i}| = t$, which starts the induction for $i - j = 1$. Now, assume by induction that $|s_{i,k}| \le t(1+t)^{i-k-1}$ for all k s.t. $i - k < i - j$ for some $1 \le j < i \le n$. Then (11) implies:

$$
\begin{aligned}
|s_{i,j}| &\le t \left(1 + \sum_{k=j+1}^{i-1} t(1+t)^{i-k-1} \right) &= t \left(1 + t \sum_{k=0}^{i-j-2} (1+t)^k \right) \\
&= t \left(1 + t \frac{(1+t)^{i-j-1} - 1}{1 + t - 1} \right) &= t(1+t)^{i-j-1}
\end{aligned}
$$

which completes the induction. Hence, we have:

$$\|S\|_\infty \le 1 + \sum_{i=2}^{n} t(1+t)^{i-2} = 1 + t \sum_{i=0}^{n-2} (1+t)^i = (1+t)^{n-1}.$$

$\square$

B   Proof of Lemma 4

The matrix $\tilde{B}$ is lower-triangular like B. Because B is size-reduced, the entries of $\tilde{B}$ satisfy, for $1 \le j < i \le n$:

$$\frac{|\tilde{b}_{i,j}|}{\tilde{b}_{j,j}} < \frac{|b_{i,j}|/2^k}{b_{j,j}/2^k - 1} \le \frac{1}{2} \times \frac{1}{1 - 2^k/b_{j,j}} \le \frac{1}{2} \times \frac{1}{1 - 1/c}$$

This means that $\tilde{B}$ is almost size-reduced. Let Δ be the $n \times n$ diagonal matrix whose i-th diagonal entry is $1/\tilde{b}_{i,i}$. Then $T = \Delta\tilde{B}$ satisfies the conditions of Lemma 7 with $t = 1/(2(1 - 1/c))$. Therefore we get:

$$\|T^{-1}\|_\infty \le \left(\frac{1}{2} \times \frac{1}{1 - 1/c} + 1 \right)^{n-1} = \left(\frac{3c - 2}{2c - 2} \right)^{n-1}.$$

Furthermore, we have $\|\tilde{B}^{-1}\|_2 \le \sqrt{n} \|\tilde{B}^{-1}\|_\infty \le \sqrt{n} \|T^{-1}\|_\infty \|\Delta\|_\infty$. Therefore we deduce that:

$$\|\tilde{B}^{-1}\|_2 \le \sqrt{n} \left(\frac{3c - 2}{2c - 2} \right)^{n-1} \times \frac{1}{\min_{1 \le i \le n} \tilde{b}_{i,i}} \le \sqrt{n} \left(\frac{3c - 2}{2c - 2} \right)^{n-1} / \lfloor c \rfloor.$$

C   Perturbation bounds for matrix operators

The following lemma gives some perturbation bounds for matrix operators. Namely, Lemma 8.a which gives a bound on determinants for perturbed matrices is used in Appendix F and the bound on inverses for perturbed matrices in Lemma 8.b is used in Appendix E.

Lemma 8. *Let A be an $n \times n$ non-singular matrix and ΔA be a "small" $n \times n$ perturbation:*

a) *Godunov* et al. *[GAKK93] showed that if $n\|A^{-1}\|_2\|\Delta A\|_2 < 1$:*

$$\frac{|\det(A + \Delta A) - \det A|}{|\det A|} \le \frac{n\|A^{-1}\|_2\|\Delta A\|_2}{1 - n\|A^{-1}\|_2\|\Delta A\|_2} \tag{12}$$

b) *Theorem 2.3.4 in the book [GVL13] by Golub and Loan states that if we have $\rho = \|A^{-1}\Delta A\|_2$ < 1, then $A + \Delta A$ is non-singular with:*

$$\|(A + \Delta A)^{-1} - A^{-1}\|_2 \le \frac{\|\Delta A\|_2\|A^{-1}\|_2^2}{1 - \rho} \tag{13}$$

D   Proof of Lemma 6

Proof. Using the Gram-Schmidt factorization, we have $B = \mu DQ$. Therefore $B^{-1} = Q^{-1}D^{-1}\mu^{-1}$ and $\|B^{-1}\|_2 \le \|D^{-1}\|_2\|\mu^{-1}\|_2$. Because B is size-reduced, we can apply Lemma 7 on μ and $t = 1/2$, which proves that $\|\mu^{-1}\|_\infty \le (3/2)^{n-1}$. Hence:

$$\|B^{-1}\|_2 \le \|D^{-1}\|_2\sqrt{n}(3/2)^{n-1},$$

where $\|D^{-1}\|_2 = 1/\min_{1 \le i \le n}\|\mathbf{b}_i^\star\|$. $\qquad\square$

E   Proof of Lemma 5

Proof. Let again $\bar{B} = \alpha\tilde{B}$, which implies that $\|B - \bar{B}\|_2 < n\alpha$. Since $n\alpha\|B^{-1}\|_2 < 1$, we have $\rho = \|B^{-1}\Delta B\|_2 < 1$, where $\Delta B = \bar{B} - B$. Then Lemma 8.b implies that $\bar{B}$ is invertible with:

$$\|\bar{B}^{-1} - B^{-1}\|_2 \le \frac{\|\Delta B\|_2\|B^{-1}\|_2^2}{1 - \rho} \le \frac{n\alpha\|B^{-1}\|_2^2}{1 - n\alpha\|B^{-1}\|_2}$$

Hence:

$$\|\bar{B}^{-1}\|_2 \le \|\bar{B}^{-1} - B^{-1}\|_2 + \|B^{-1}\|_2 \le \|B^{-1}\|_2 + \frac{n\alpha\|B^{-1}\|_2^2}{1 - n\alpha\|B^{-1}\|_2} = \frac{\|B^{-1}\|_2}{1 - n\alpha\|B^{-1}\|_2}$$

Since $\bar{B} = \alpha\tilde{B}$, we have $\tilde{B}^{-1} = \alpha\bar{B}^{-1}$. Therefore $\|\tilde{B}^{-1}\|_2 = \alpha\|\bar{B}^{-1}\|_2$, which concludes the proof. $\qquad\square$

F   Proof of Corollary 6

Proof. We have $\tilde{B} = \lfloor B/\alpha \rceil$. Define the matrix $\bar{B} = \alpha\tilde{B}$ whose entries are $\bar{b}_{i,j} = \alpha\tilde{b}_{i,j}$. Since $n^2\alpha\|B^{-1}\|_2 < 1$, Lemma 5 implies that $\tilde{B}$ is non-singular. Furthermore, we have:

$$\|\mathbf{x}B\| \le \|\mathbf{x}(B - \bar{B})\| + \|\mathbf{x}\bar{B}\| \le \|\mathbf{x}\| \times \|B - \bar{B}\|_2 + \alpha\|\mathbf{x}\tilde{B}\| < n\|\mathbf{x}\|\alpha + \alpha\|\mathbf{x}\tilde{B}\|.$$

Let $\tilde{\mathbf{v}} = \mathbf{x}\tilde{B}$. Then $\|\mathbf{x}\| \leq \|\tilde{\mathbf{v}}\|\|\tilde{B}^{-1}\|_2$, and we obtain:

$$\|\mathbf{x}B\| < \left(n\|\tilde{B}^{-1}\|_2 + 1\right)\alpha\|\tilde{\mathbf{v}}\|.$$

Since $\tilde{\mathbf{v}} = \mathbf{x}\tilde{B}$ is the first vector of an LLL-reduced basis of $\tilde{B}$, and $\tilde{B}$ is non-singular, $\mathbf{x}B \neq 0$ and we have:

$$\alpha\|\tilde{\mathbf{v}}\| \leq \alpha 2^{\frac{n-1}{4}}\det(\tilde{B})^{\frac{1}{n}} = 2^{\frac{n-1}{4}}\det(\bar{B})^{\frac{1}{n}}$$

The condition $n^2\alpha\|B^{-1}\|_2 < 1$ implies that $n\|B^{-1}\|_2\|\bar{B} - B\|_2 < 1$. Since $\bar{B} = B + (\bar{B} - B)$, Lemma 8.a implies:

$$\frac{|\det(\bar{B}) - \det B|}{|\det B|} \leq \frac{n\|B^{-1}\|_2\|\bar{B} - B\|_2}{1 - n\|B^{-1}\|_2\|\bar{B} - B\|_2} \leq \frac{n^2\alpha\|B^{-1}\|_2}{1 - n^2\alpha\|B^{-1}\|_2}$$

It follows that:

$$|\det(\bar{B})| \leq |\det B|\left(1 + \frac{n^2\alpha\|B^{-1}\|_2}{1 - n^2\alpha\|B^{-1}\|_2}\right) = \frac{|\det B|}{1 - n^2\alpha\|B^{-1}\|_2}.$$

Therefore, we get the following inequality:

$$\|\mathbf{x}B\| < \left(n\|\tilde{B}^{-1}\|_2 + 1\right)2^{\frac{n-1}{4}}(1 - n^2\alpha\|B^{-1}\|_2)^{-1/n}\det(B)^{\frac{1}{n}}. \tag{14}$$

Using Lemmas 6 and 5, which respectively give the upper bounds

$$\|B^{-1}\|_2 \leq \frac{\sqrt{n}(3/2)^{n-1}}{\min_{1 \leq i \leq n}\|\mathbf{b}_i^\star\|} \quad \text{and} \quad \|\tilde{B}^{-1}\|_2 \leq \frac{\alpha\|B^{-1}\|_2}{1 - n\alpha\|B^{-1}\|_2},$$

and combining them in (14), one gets

$$\|\mathbf{x}B\| < \frac{\min_{1 \leq i \leq n}\|\mathbf{b}_i^\star\|^{\frac{n+1}{n}}2^{\frac{n-1}{4}}\det(B)^{\frac{1}{n}}}{(\min_{1 \leq i \leq n}\|\mathbf{b}_i^\star\| - n^{3/2}\alpha(3/2)^{n-1})(\min_{1 \leq i \leq n}\|\mathbf{b}_i^\star\| - n^{5/2}\alpha(3/2)^{n-1})^{1/n}}$$

Finally, replacing α by $\min_{1 \leq i \leq n}\|\mathbf{b}_i^\star\|/c$, allows to get the bound stated in Cor. 6. □

G   Approximate Gcd and Multivariate Equations

G.1   Gcd Generalization

Coppersmith's algorithm (Alg. 1) has been generalized by essentially Howgrave-Graham [HG01] and Boneh *et al.* [BDHG99] (see the surveys [May10,Ngu09]) as follows:

Theorem 8. *There is an algorithm which, given as input an integer N of unknown factorization, a rational α s.t. $0 < \alpha \leq 1$ and a monic polynomial $f(x) \in \mathbb{Z}[x]$ of degree δ and coefficients in $\{0, \ldots, N-1\}$, outputs all integers $x_0 \in \mathbb{Z}$ such that $\gcd(f(x_0), N) \geq N^\alpha$ and $|x_0| \leq N^{\alpha^2/\delta}$ in time polynomial in $\log N$, δ and the bit-size of α.*

Th. 1 is then the special case $\alpha = 1$ of Th. 8. The algorithm underlying Th. 8 is in fact very similar to Alg. 1: instead of applying Lemma 1 with $M = N^{h-1}$, one uses $M = p^{h-1}$ where $p \geq N^\alpha$ is some unknown divisor of N. And one considers the same family of polynomials $g_{i,j}(x) = x^j N^{h-1-i}f^i(x)$ but over different indices. Alg. 1 used $0 \leq i < h$ and $0 \leq j < \delta$. This time, we use two sets of indices: one with $0 \leq i < h-1$ and $0 \leq j < \delta$, and another with $i = h-1$ and $0 \leq j < \gamma$, where γ is chosen asymptotically to be $\lfloor \delta(h-1)(1/\alpha - 1) \rfloor$. Then the dimension is $n = (h-1)\delta + \gamma$. The maximal diagonal coefficient is still $N^{h-1}X^{\delta-1}$, and

the minimal diagonal coefficient is still $X^{h\delta-\delta}$, like in Lemma 2. However, the balance between these two coefficients has changed, because the bound X is much smaller than in Coppersmith's algorithm. Before, X was essentially $N^{(h-1)/(n-1)}$ whose order of magnitude is the same as $N^{1/\delta}$, but now, it is close to $N^{\alpha^2/\delta}$, so $X^{h\delta-\delta}$ is close to $N^{(h-1)\alpha^2}$. In other words, the ratio between the maximal and minimal diagonal coefficient is about $N^{(1-\alpha^2)(h-1)}$ which is no longer $N^{O(1)}$. We are thus trading an LLL reduction of a matrix with bit-size $\approx (h-1)\log N$, with one with bit-size $\approx (1-\alpha^2)(h-1)\log N$, which can only provide a small constant speedup at best, namely $1/(1-\alpha^2)^2$ for L^2 or close to $1/(1-\alpha^2)$ for $\tilde{L}^1$. In the gcd generalization, the input basis is much less reduced than in Coppersmith's algorithm.

G.2 Multivariate Equations

Coppersmith [Cop96b,Cop97] showed that his algorithm for finding small roots of univariate polynomial congruences can heuristically be extended to multivariate polynomial congruences: the most famous example is the Boneh-Durfee attack [BD00] on RSA with short secret exponent.

Not all these multivariate variants use triangular matrices, though they sometimes can be tweaked: some rely on lattices which are not full-rank, including the Boneh-Durfee attack [BD00]. However, when the matrix is triangular, there is a similar problem than for the gcd generalization: the diagonal coefficients are much more unbalanced than in the univariate congruence case, which means that the speedup of the rounding method is at most a small constant. And in the Boneh-Durfee attack, the coefficients which play the role of the diagonal coefficients are also unbalanced.

For instance, assume that one would like to find all small roots of $f(x,y) \equiv 0 \pmod N$ with $|x| \leq X$ and $|y| \leq Y$, where $f(x,y)$ has total degree δ and has at least one monic monomial $x^\alpha y^{\delta-\alpha}$ of maximal total degree. Then, for a given parameter h, the lower-triangular matrix has dimension $n = (h\delta+1)(h\delta+2)/2$ and diagonal coefficients $N^{h-v}X^{u_1+v\delta}Y^{u_2+v(\delta-\alpha)}$, where $u_1 + u_2 + \delta v \leq h\delta$ and $u_1, u_2, v \geq 0$ with $u_1 < \alpha$ or $u_2 < \delta - \alpha$. For typical choices of X and Y such that $XY < N^{1/\delta-\varepsilon}$, the ratio between the largest and smallest diagonal coefficient is no longer $N^{O(1)}$.