CONJURA · OPEN PROBLEM

Recovering the Ring-LWE Error Modulo Two Is as Hard as Recovering It

Search Ring-LWE with centred binomial noise over an NTT-friendly ring

Status. Statement: AI-written from Julien Duman, Kathrin Hövelmanns, Eike Kiltz, Vadim Lyubashevsky, Gregor Seiler, Dominique Unruh, *A Thorough Treatment of Highly-Efficient NTRU Instantiations*, Cryptology ePrint Archive 2021, not yet checked by a human. The direction in which recovering the whole error yields its parities is trivial and is not in question; the converse asserted here has no formal reduction in the source, which substitutes two heuristics — one that drifts off the prescribed error distribution, and one that reduces from a nonstandard first-is-errorless decision problem rather than from $\mathcal{R}$-LWE itself.

Category. *Public-Key Cryptography & Assumptions.*

Conjecture (informal). *In the Ring-LWE problem you are given a uniformly random ring element and a noisy product, and asked to recover the small error term. This paper needs a weaker-looking variant: recover only the error term reduced modulo two, that is, the parities of its coefficients and nothing else. The variant arises because the paper's fastest NTRU encryption scheme has a secret key that, by design, recovers only the parities of the error, so the message it can carry is exactly that parity vector; the scheme's one-wayness therefore rests on the parity-only problem rather than on full Ring-LWE. Recovering the parities is clearly no harder than recovering everything. The conjecture is the converse: an efficient algorithm that finds the parities can be turned into an efficient algorithm that finds the whole error, so that the two problems are equally hard up to polynomial factors. The paper says it has no such reduction and offers two heuristic arguments instead. The first repeatedly strips off known parities, halving the error each time, but each round leaves the error with a different, narrower distribution than the one the parity oracle was promised, so it is not a reduction. The second gives a genuine reduction, but only from a nonstandard decision problem in which the first coefficient carries no noise at all.*

1 The setting

The NTRU trapdoor function [HPS98] sets the public key to a quotient $\mathbf{h}$ of two small-coefficient polynomials and the ciphertext to $\mathbf{hr}+\mathbf{e}$, so its one-wayness combines two assumptions: that $\mathbf{h}$ is indistinguishable from uniform, and that the Ring-LWE problem of recovering $\mathbf{e}$ from $(\mathbf{h}, \mathbf{hr}+\mathbf{e})$ is hard [SSTX09, LPR10]. (Choosing the small polynomials wide enough removes the first assumption [SS11], at a substantial efficiency cost.)

The paper's fastest instantiation, over the NTT-friendly ring $\mathbb{Z}_q[X]/(X^d - X^{d/2} + 1)$ in the style of [LS19], is built so that the message need not be the whole of $\mathbf{e}$. Its secret key has the form $\mathbf{f} = 2\mathbf{f}' + 1$, and decryption computes $(\mathbf{cf} \bmod^{\pm} q) \bmod 2$, which equals $\mathbf{e} \bmod 2$ whenever the noise stays inside the decoding radius. Only the parity vector is recovered, so only the parity vector can be the message. Two things follow. First, the adversary controls only $\mathbf{e} \bmod 2$ and hence has little freedom to inflate the decryption error, which is why this scheme needs no worst-case-to-average-case correctness transformation at all. Second, its one-wayness rests on the parity-only problem $\mathcal{R}$-LWE2 rather than on $\mathcal{R}$-LWE, and this parity-only assumption is what appears in the paper's concrete security bounds for the resulting KEM.

The easy direction is immediate: an algorithm that outputs $\mathbf{e}$ also outputs $\mathbf{e} \bmod 2$, so $\mathcal{R}$-LWE2 is no harder than $\mathcal{R}$-LWE. The conjecture is the direction the scheme needs, and the paper states plainly that it has no formal reduction for it, offering two heuristics instead.

The first heuristic is an iterative stripping argument: given the correct $\mathbf{f} \equiv \mathbf{e} \pmod 2$, one forms $(2^{-1}\mathbf{h},\ 2^{-1}\mathbf{hr} + 2^{-1}(\mathbf{e}-\mathbf{f}))$, which is again an instance with uniform first component, and repeats until all of $\mathbf{e}$ is known. The obstruction is named explicitly: the new error is drawn from a strictly narrower distribution than ψ_2^d, so on the second and later rounds the parity oracle is being fed inputs it was never promised to handle. The paper's justification for calling it anyway is a belief about how lattice algorithms behave, not a proof. This is exactly why the conjecture below insists that $\mathcal{A}$'s guarantee holds only on-distribution: a reduction that dodged this restriction would be assuming the heuristic.

The second heuristic is a genuine reduction, but from a mod-

ified problem: if the decision version of $\mathcal{R}$-LWE in which the first coefficient carries no noise is hard, then $\mathcal{R}$-LWE2 is hard, since the reduction can inject its own noise into that coefficient and compare the returned parity with the parity it injected. Brakerski et al. [BLP13] study this "First-is-Errorless" variant of LWE and show it essentially as hard as the usual one, and Boudgoust et al. [BJRW21] give a non-tight reduction from Module-LWE to a stronger relative; but neither yields a reduction to the search problem $\mathcal{R}$-LWE2 with the fixed narrow distribution ψ_2^d, which is the object the scheme is built on.

Beyond this one scheme, the question is the natural search-to-partial-search question for Ring-LWE with fixed narrow noise [Reg09]: are the low-order bits of the error as hard as the error itself, when the standard re-randomisation tricks are unavailable because they widen or narrow the noise distribution?

Progress to date. The paper gives two heuristic arguments (page 20). (i) Iterated stripping: a correct parity answer lets one halve the error and recurse, so the full error is eventually recovered — but each round narrows the error distribution, so the oracle is queried off-distribution. (ii) A clean reduction from the decision Ring-LWE variant whose first coefficient is noiseless, together with the observation that current algorithms would solve that decision problem via its search version, where the two variants are equally hard. It cites [BLP13] for the hardness of the First-is-Errorless variant of LWE and [BJRW21] for a non-tight Module-LWE reduction to a stronger relative, concluding only that assuming equal concrete hardness is reasonable.

2 Notation and parameters

Throughout, $d = 2^i 3^j$ for integers $i, j \geq 1$, and $q > 4$ is a prime. We write

$$\mathcal{R} \;:=\; \mathbb{Z}_q[X]/\left(X^d - X^{d/2} + 1\right)$$

for the associated polynomial ring. Elements of $\mathcal{R}$ are denoted by bold lower-case letters and are identified with their coefficient vectors in $\mathbb{Z}_q^d$. Writing $\mathbf{x} \leftarrow \mathcal{R}$ means that $\mathbf{x}$ is sampled uniformly from $\mathcal{R}$.

For $h \in \mathbb{Z}_q$, let $h \bmod^{\pm} q$ be the unique integer in $\left\{-\frac{q-1}{2}, \ldots, \frac{q-1}{2}\right\}$ congruent to h modulo q; this map is applied coefficientwise to

elements of $\mathcal{R}$. For $\mathbf{e} \in \mathcal{R}$ we write $\mathbf{e} \bmod 2 \in \{0,1\}^d$ for the vector obtained by reducing each entry of $\mathbf{e} \bmod^{\pm} q$ modulo 2.

Let ψ_2^d be the centred binomial distribution on $\mathbb{Z}^d$ given by

$$\vec{a}_1 + \vec{a}_2 - \vec{b}_1 - \vec{b}_2, \qquad \vec{a}_1, \vec{a}_2, \vec{b}_1, \vec{b}_2 \leftarrow \{0,1\}^d$$

independently and uniformly, so that each coordinate takes the values $-2,-1,0,1,2$ with probabilities $\frac{1}{16}, \frac{4}{16}, \frac{6}{16}, \frac{4}{16}, \frac{1}{16}$. Since $q > 4$, a sample of ψ_2^d is regarded as an element of $\mathcal{R}$ through the representatives above, and its reduction modulo 2 in the sense of the previous paragraph agrees with coefficientwise reduction of the integer vector.

Algorithms are classical and probabilistic; $T(\mathcal{A})$ denotes the running time of $\mathcal{A}$, and P denotes a polynomial with non-negative coefficients.

- d: degree of the ring, of the form $2^i 3^j$; the error term has d coefficients.
- q: prime modulus of the ring $\mathcal{R} = \mathbb{Z}_q[X]/(X^d - X^{d/2} + 1)$.
- ψ_2^d: centred binomial distribution used for both the secret $\mathbf{r}$ and the error $\mathbf{e}$; coefficients lie in $\{-2, \ldots, 2\}$.
- ε: success probability of the assumed algorithm for the parity-only problem $\mathcal{R}$-LWE2.
- P: polynomial bounding both the reduction's running time and the reciprocal of its success probability.

3 The conjecture

Definition 1 (Search $\mathcal{R}$-LWE with binomial noise). For an algorithm $\mathcal{A}$ set

$$\mathrm{Adv}_{d,q}^{\mathcal{R}\text{-LWE}}(\mathcal{A}) \;:=\; \Pr\big[\, \mathcal{A}(\mathbf{h}, \mathbf{hr} + \mathbf{e}) = \mathbf{e} \,\big],$$

where the probability is over $\mathbf{h} \leftarrow \mathcal{R}$, over $\mathbf{r}, \mathbf{e} \leftarrow \psi_2^d$ sampled independently, and over the coins of $\mathcal{A}$.

Definition 2 ($\mathcal{R}$-LWE2 with binomial noise). For an algorithm $\mathcal{A}$ set

$$\mathrm{Adv}^{\mathcal{R}\text{-LWE2}}_{d,q}(\mathcal{A}) \;:=\; \Pr\big[\,\mathcal{A}(\mathbf{h}, \mathbf{h}\mathbf{r}+\mathbf{e}) = \mathbf{e} \bmod 2\,\big],$$

with $\mathbf{h}, \mathbf{r}, \mathbf{e}$ distributed exactly as in the previous definition (Definition 1).

Conjecture 1 (parity-only Ring-LWE is as hard as Ring-LWE). *There exist an oracle algorithm $\mathcal{B}$ and a polynomial P such that the following holds for every $d = 2^i 3^j$ with $i, j \geq 1$, every prime $q > 4$, every $\varepsilon \in (0,1]$, and every algorithm $\mathcal{A}$ with*

$$\mathrm{Adv}^{\mathcal{R}\text{-LWE2}}_{d,q}(\mathcal{A}) \;\geq\; \varepsilon.$$

The algorithm $\mathcal{B}$, given (d, q, ε) together with its own $\mathcal{R}$-LWE input and oracle access to $\mathcal{A}$, satisfies

$$\mathrm{Adv}^{\mathcal{R}\text{-LWE}}_{d,q}\big(\mathcal{B}^{\mathcal{A}}\big) \;\geq\; \frac{1}{P(d, \log q, 1/\varepsilon)} \qquad \text{and}$$
$$T\big(\mathcal{B}^{\mathcal{A}}\big) \;\leq\; P(d, \log q, 1/\varepsilon) \cdot \big(T(\mathcal{A}) + 1\big).$$

Here $\mathcal{B}$ may query $\mathcal{A}$ on inputs of its choice, but the assumed success probability of $\mathcal{A}$ is guaranteed only for inputs distributed exactly as in the definition of $\mathrm{Adv}^{\mathcal{R}\text{-LWE2}}_{d,q}$, i.e. with $\mathbf{h}$ uniform on $\mathcal{R}$ and $\mathbf{r}, \mathbf{e} \leftarrow \psi_2^d$; no guarantee is made about $\mathcal{A}$'s behaviour on inputs from any other distribution.

Remark 1 (the on-distribution restriction is part of the statement). The final clause of Conjecture 1 is what separates it from the paper's first heuristic. That heuristic calls the parity oracle on instances whose error is drawn from a strictly narrower distribution than ψ_2^d, and $\mathcal{A}$ is promised nothing there. A reduction that helped itself to $\mathcal{A}$'s success on such inputs would be assuming the heuristic rather than proving it.

4 Bibliography

[BJRW21] K. Boudgoust, C. Jeudy, A. Roux-Langlois, W. Wen. *On the hardness of module-LWE with binary secret.* CT-RSA, volume 12704 of Lecture Notes in Computer Science, pages 503–526, 2021.

[BLP13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, D. Stehle. *Classical hardness of learning with errors.* STOC, pages 575–584, 2013.

[HPS98] J. Hoffstein, J. Pipher, J. H. Silverman. *NTRU: A ring-based public key cryptosystem.* ANTS, pages 267–288, 1998.

[LPR10] V. Lyubashevsky, C. Peikert, O. Regev. *On ideal lattices and learning with errors over rings.* EUROCRYPT, pages 1–23, 2010.

[LS19] V. Lyubashevsky, G. Seiler. *NTTRU: truly fast NTRU using NTT.* IACR Trans. Cryptogr. Hardw. Embed. Syst., 2019(3):180–201, 2019.

[Reg09] O. Regev. *On lattices, learning with errors, random linear codes, and cryptography.* J. ACM, 56(6), 2009.

[SS11] D. Stehle, R. Steinfeld. *Making NTRU as secure as worst-case problems over ideal lattices.* EUROCRYPT, pages 27–47, 2011.

[SSTX09] D. Stehle, R. Steinfeld, K. Tanaka, K. Xagawa. *Efficient public key encryption based on ideal lattices.* ASIACRYPT, pages 617–635, 2009.