

A year of lattice cryptanalysis

Daniel J. Bernstein

Some context: Do we keep an ECC layer?

- 👍 For 10 years I've been recommending PQ rollout as ECC+PQ (even for hash-based signatures).

Some context: Do we keep an ECC layer?

- 👍 For 10 years I've been recommending PQ rollout as ECC+PQ (even for hash-based signatures).
- 🗨 In 2024, Cisco said it was pushing “pure ML-KEM” because that's what NSA was “willing to buy”.
- 🗨 2025 NSA: “Our interactions with vendors suggests that this won't be a problem in most cases.”

Some context: Do we keep an ECC layer?

- 👍 For 10 years I've been recommending PQ rollout as ECC+PQ (even for hash-based signatures).
- 👎 In 2024, Cisco said it was pushing “pure ML-KEM” because that's what NSA was “willing to buy”.
- 👎 2025 NSA: “Our interactions with vendors suggests that this won't be a problem in most cases.”
- 👍 2026: 82 people filed unambiguous objections to IETF issuing an RFC on solo ML-KEM in TLS.

Some context: Do we keep an ECC layer?

- 👍 For 10 years I've been recommending PQ rollout as ECC+PQ (even for hash-based signatures).
- 🗨 In 2024, Cisco said it was pushing “pure ML-KEM” because that's what NSA was “willing to buy”.
- 🗨 2025 NSA: “Our interactions with vendors suggests that this won't be a problem in most cases.”
- 👍 2026: 82 people filed unambiguous objections to IETF issuing an RFC on solo ML-KEM in TLS.
- 🗨 But many votes for solo PQ: e.g., 6 NSA employees, 4 Cisco employees, 3 GCHQ employees, etc.

Confidence in lattices?

20260303 Daniel Apon: “full decade . . . process.
ML-KEM was fully vetted through this process”

Confidence in lattices?

20260303 [Daniel Apon](#): “full decade . . . process. ML-KEM was fully vetted through this process”

20261003 [Matthew Green](#): “What I’m interested in, in principle, is any result that calls into question our confidence about LWE (and very explicitly: module LWE) cryptosystems at current proposed ‘safe’ parameter sizes. . . . [such as reducing] the estimated cost of ML-KEM-768 by at least ~ 20 bits versus the FIPS 203 analysis . . . [or] the best known exponent for SVP/BDD at cryptographically relevant approximation factors . . . [or] anything subexponential for LWE with polynomial modulus.”

Let's review the lattice-cryptanalysis papers
that have appeared on eprint
within the past 365 days.

Caveat:

proofs often have big gaps;
proofs are rarely computer-checked;
experiments often have big gaps.

Sometimes papers are wrong.

20251012: MLWE

2025/1904, Léo Ducas, Lynn Engelberts, Paola de Perthuis, “Predicting Module-Lattice Reduction”:

“Is module-lattice reduction better than unstructured lattice reduction? . . . For power-of-two cyclotomic fields, we . . . conclude that module-BKZ requires a blocksize larger than its unstructured counterpart by $d - 1 + o(1)$.”

20251012: MLWE

2025/1904, Léo Ducas, Lynn Engelberts, Paola de Perthuis, “Predicting Module-Lattice Reduction”:

“Is module-lattice reduction better than unstructured lattice reduction? . . . For power-of-two cyclotomic fields, we . . . conclude that module-BKZ requires a blocksize larger than its unstructured counterpart by $d - 1 + o(1)$.”

— Seems to justify the Kyber [submission](#): “The best known attacks against the underlying MLWE problem in Kyber do not make use of the structure in the lattice. We therefore analyze the hardness of the MLWE problem as an LWE problem.”

Non-power-of-2 cyclotomics

2025/1904 continues: “for other cyclotomic fields, especially those whose conductor has one or more small odd prime factors, we predict a rather substantial gain . . . module-BKZ over the third cyclotomic field is predicted to gain 20 more dimensions on the blocksize at NIST security level 1”.

Non-power-of-2 cyclotomics

2025/1904 continues: “for other cyclotomic fields, especially those whose conductor has one or more small odd prime factors, we predict a rather substantial gain . . . module-BKZ over the third cyclotomic field is predicted to gain 20 more dimensions on the blocksize at NIST security level 1”.

— Sounds like several bits of security loss for some lattice KEMs, such as NTRU+.

20251013: hybrid attacks

2025/1910, Alexander Karenin, Elena Kirshanova, Alexander May, Julian Nowakowski, “Fast Slicer for Batch-CVP: Making Lattice Hybrid Attacks Practical”:

“We study the practicality of a primal hybrid lattice attack on LWE. Despite significant research efforts, the state-of-the-art in practical LWE record computations so far is the plain primal attack . . . Building on an idea by Espitau and Kirchner, Bernstein proposed in 2023 an LWE hybrid attack that asymptotically outperforms the primal attack. . . . The practical implications of this attack however remain widely unclear.”

Hybrid attacks, continued

“... we provide an efficient Slicer implementation that includes all required algorithmic ingredients like locality sensitive hashing. ...

Hybrid attacks, continued

“... we provide an efficient Slicer implementation that includes all required algorithmic ingredients like locality sensitive hashing. ...

We choose two ternary LWE secrets with different entropies $H(D)$ as used in NTRU, and the centered binomial distribution as used in Kyber. For all three distributions in all tested LWE dimensions $n \in [160, 210]$, our Slicer-based implementation practically demonstrates measurable speedups over the primal attack, up to a factor of 5.”

Review of primal attacks

Write $R = \mathbb{Z}[x]/(x^n + 1)$. Goal: find secrets $a, e \in R$ given $G \in R/q$ and $A = aG + e \in R/q$.

Review of primal attacks

Write $R = \mathbb{Z}[x]/(x^n + 1)$. Goal: find secrets $a, e \in R$ given $G \in R/q$ and $A = aG + e \in R/q$.

Standard primal attack: Pick β, κ . Define L as the set of all $(\alpha, \epsilon, \gamma) \in R \times \mathbb{Z}^\kappa \times \mathbb{Z}$ such that $\text{First}_\kappa(\gamma A - \alpha G) \equiv \epsilon \pmod{q}$. Use BKZ- β to reduce L , hope that it outputs $\pm(a, \text{First}_\kappa(e), 1)$.

Review of primal attacks

Write $R = \mathbb{Z}[x]/(x^n + 1)$. Goal: find secrets $a, e \in R$ given $G \in R/q$ and $A = aG + e \in R/q$.

Standard primal attack: Pick β, κ . Define L as the set of all $(\alpha, \epsilon, \gamma) \in R \times \mathbb{Z}^\kappa \times \mathbb{Z}$ such that $\text{First}_\kappa(\gamma A - \alpha G) \equiv \epsilon \pmod{q}$. Use BKZ- β to reduce L , hope that it outputs $\pm(a, \text{First}_\kappa(e), 1)$.

There's a standard heuristic choice of β, κ depending on n, q , and size of a, e coeffs.

Typical examples end up with, e.g., $\beta \approx 0.8n$.

How do hybrids do better?

Simplest case of 2007 Nick Howgrave-Graham:

Define L' as the set of all $(\omega, \epsilon) \in R \times \mathbb{Z}^\kappa$ such that $\text{First}_\sigma(\omega) = 0$ and $\text{First}_\kappa(-\omega G) \equiv \epsilon \pmod{q}$.

How do hybrids do better?

Simplest case of 2007 Nick Howgrave-Graham:

Define L' as the set of all $(\omega, \epsilon) \in R \times \mathbb{Z}^\kappa$ such that $\text{First}_\sigma(\omega) = 0$ and $\text{First}_\kappa(-\omega G) \equiv \epsilon \pmod{q}$.

Try guessing $u = \text{First}_\sigma a$. If guess is correct then $(a - u, \text{First}_\kappa(e + uG - A))$ is in L' and is close to the known vector $(0, \text{First}_\kappa(uG - A))$.

How do hybrids do better?

Simplest case of 2007 Nick Howgrave-Graham:

Define L' as the set of all $(\omega, \epsilon) \in R \times \mathbb{Z}^\kappa$ such that $\text{First}_\sigma(\omega) = 0$ and $\text{First}_\kappa(-\omega G) \equiv \epsilon \pmod{q}$.

Try guessing $u = \text{First}_\sigma a$. If guess is correct then $(a - u, \text{First}_\kappa(e + uG - A))$ is in L' and is close to the known vector $(0, \text{First}_\kappa(uG - A))$.

Apply a close-vector algorithm using BKZ- β .

$\text{Dim } L'$ is only $n + \kappa - \sigma$. ($\text{Dim } L$ was $n + \kappa + 1$.)

Larger $\sigma \rightarrow$ smaller dim $\rightarrow$ smaller $\beta \rightarrow$ saves time.

But we have to repeat for each guess!

But we have to repeat for each guess!

Or do we? — My [2023 paper](#) says no:

But we have to repeat for each guess!

Or do we? — My [2023 paper](#) says no:

1. Idea noted briefly by 2007 Howgrave-Graham, analyzed by 2020 Thomas Espitau–Paul Kirchner: Apply BKZ- β to reduce L' . Then use a CVP- β algorithm to find a vector in β -dimensional projection of L' closest to projection of target.

But we have to repeat for each guess!

Or do we? — My [2023 paper](#) says no:

1. Idea noted briefly by 2007 Howgrave-Graham, analyzed by 2020 Thomas Espitau–Paul Kirchner: Apply BKZ- β to reduce L' . Then use a CVP- β algorithm to find a vector in β -dimensional projection of L' closest to projection of target.
2. CVP algorithms from 2016 Thijs Laarhoven, 2019 Emmanouil Doulgerakis–Laarhoven–Benne de Weger, and 2020 Ducas–Laarhoven–Wessel van Woerden handle many targets at low cost after one lattice precomputation.

Does the reduced exponent matter?

$\leq 5\times$ speedup in experiments.

Does the reduced exponent matter?

$\leq 5\times$ speedup in experiments. — True. However:

- These attacks contradict the NewHope paper claiming that hybrid attacks are “induced by the fact that secrets are ternary”.

Does the reduced exponent matter?

$\leq 5\times$ speedup in experiments. — True. However:

- These attacks contradict the NewHope paper claiming that hybrid attacks are “induced by the fact that secrets are ternary”.
- The savings should be $\approx$ exponential in the dimension. >10 bits for Kyber-768?

Does the reduced exponent matter?

$\leq 5\times$ speedup in experiments. — True. However:

- These attacks contradict the NewHope paper claiming that hybrid attacks are “induced by the fact that secrets are ternary”.
- The savings should be $\approx$ exponential in the dimension. >10 bits for Kyber-768?
- Some papers (e.g., [2003](#), [2007](#), [2016](#), [2021](#), [2022](#)) plausibly say: collisions and multi-level collisions speed up u search. So try this too.

Does the reduced exponent matter?

$\leq 5\times$ speedup in experiments. — True. However:

- These attacks contradict the NewHope paper claiming that hybrid attacks are “induced by the fact that secrets are ternary”.
- The savings should be $\approx$ exponential in the dimension. >10 bits for Kyber-768?
- Some papers (e.g., 2003, 2007, 2016, 2021, 2022) plausibly say: collisions and multi-level collisions speed up u search. So try this too.
- Also try dual attacks, which supposedly beat primal attacks for batch close vectors.

20251030: hybrid attacks

2025/1990, Ludo N. Pulles, Paul Vié,
“Accelerating the Primal Hybrid Attack
against Sparse LWE using GPUs”

— Less scalable type of hybrid
(no slicer; only codimension-1 colattices),
focusing on a “sparse” case with very low noise.

But more GPU software speeds up experiments.
Would be good to see more advanced GPU attacks.

20251202: tuple lattice sieving

[2025/2189](#), Lynn Engelberts, Yanlin Chen, Amin Shiraz Gilani, Maya-Iggy van Hoof, Stacey Jeffery, Ronald de Wolf, “An Improved Quantum Algorithm for 3-Tuple Lattice Sieving”:

“we improve the quantum time complexity of 3-tuple sieving from $2^{0.3098d}$ to $2^{0.2846d}$... This is the fastest known quantum algorithm for SVP when total memory is limited to $2^{0.1887d}$.”

Review of tuple lattice sieving

2011 Bernstein–Lange (in application with de Weger): “One of the major goals of this project is to unify the ideas of sieving algorithms with the ideas of traditional lattice-basis reduction and enumeration. The literature describes the algorithms as being completely different, but we observe that they all follow the same basic strategy: generate new vectors by adding and subtracting old vectors; keep small vectors; discard large vectors.”

Review of tuple lattice sieving, part 2

“Traditional reduction algorithms discard as quickly as possible, replacing a few vectors with small combinations of those vectors; sieving discards as slowly as possible, accumulating a large list of vectors to improve the reduction of new vectors. One can imagine many intermediate possibilities, and it seems likely that the optimum combination will be somewhere in between.”

Review of tuple lattice sieving, part 3

“An initial step towards this goal will be to design and analyze a generalization of sieving with tunable memory consumption. Sieving is frequently criticized as requiring large amounts of storage (and large amounts of time for cache misses etc.), while traditional reduction and enumeration fit into far less storage; reducing the storage for sieving will help reduce the separation between sieving and other algorithms, and may already improve performance.”

Review of tuple lattice sieving, part 4

“The reason that sieving uses large amounts of storage is that it reduces vectors in only a very simple way, subtracting the closest vector from a large list; for comparison, one of the first steps in traditional lattice-basis reduction is to subtract a reasonably close *combination* of vectors from a short list. An intermediate possibility is to find a close combination of *two* vectors from a medium-size list; the algorithmic challenge here is to balance closeness with efficiency.”

Today this is known as “3-tuple lattice sieving”.

Analysis and optimization

This text wasn't public but went to Laarhoven.
2016 Shi Bai–Laarhoven–Damien Stehlé (crediting
idea to Bernstein–Lange) conjectured exponent
 ≈ 0.5661 for naive 3-tuples.

Analysis and optimization

This text wasn't public but went to Laarhoven.
2016 Shi Bai–Laarhoven–Damien Stehlé (crediting
idea to Bernstein–Lange) conjectured exponent
 ≈ 0.5661 for naive 3-tuples. Also:
 ≈ 0.4812 for improved 3-tuples.

Analysis and optimization

This text wasn't public but went to Laarhoven.
2016 Shi Bai–Laarhoven–Damien Stehlé (crediting
idea to Bernstein–Lange) conjectured exponent
 ≈ 0.5661 for naive 3-tuples. Also:
 ≈ 0.4812 for improved 3-tuples.
 ≈ 0.3717 : 2017 Gottfried Herold–Kirshanova.

Analysis and optimization

This text wasn't public but went to Laarhoven.
2016 Shi Bai–Laarhoven–Damien Stehlé (crediting
idea to Bernstein–Lange) conjectured exponent
 ≈ 0.5661 for naive 3-tuples. Also:
 ≈ 0.4812 for improved 3-tuples.
 ≈ 0.3717 : 2017 Gottfried Herold–Kirshanova.
 ≈ 0.3588 : 2017 Herold–Kirshanova–Laarhoven.

Analysis and optimization

This text wasn't public but went to Laarhoven.
2016 Shi Bai–Laarhoven–Damien Stehlé (crediting
idea to Bernstein–Lange) conjectured exponent
 ≈ 0.5661 for naive 3-tuples. Also:
 ≈ 0.4812 for improved 3-tuples.
 ≈ 0.3717 : 2017 Gottfried Herold–Kirshanova.
 ≈ 0.3588 : 2017 Herold–Kirshanova–Laarhoven.
 ≈ 0.3383 : 2023 Johanna Loyer–André Chailloux.

Analysis and optimization

This text wasn't public but went to Laarhoven.
2016 Shi Bai–Laarhoven–Damien Stehlé (crediting
idea to Bernstein–Lange) conjectured exponent
 ≈ 0.5661 for naive 3-tuples. Also:
 ≈ 0.4812 for improved 3-tuples.
 ≈ 0.3717 : 2017 Gottfried Herold–Kirshanova.
 ≈ 0.3588 : 2017 Herold–Kirshanova–Laarhoven.
 ≈ 0.3383 : 2023 Johanna Loyer–André Chailloux.

Also some analysis of k -tuples with $k \geq 4$,
and some papers analyzing the quantum case.

How low will this go?

So far tuple lattice sieving trades time for memory, contradicting claims about minimum memory.

How low will this go?

So far tuple lattice sieving trades time for memory, contradicting claims about minimum memory.

Maybe it ends up being *faster* in less memory.
(This is also what happened to hybrid attacks.)

How low will this go?

So far tuple lattice sieving trades time for memory, contradicting claims about minimum memory.

Maybe it ends up being *faster* in less memory.
(This is also what happened to hybrid attacks.)

There's a **lower bound** ≈ 0.2925 for a class of near-neighbor 2-tuple sieving algorithms.

The bound does not apply to 3-tuples etc.

How low will this go?

So far tuple lattice sieving trades time for memory, contradicting claims about minimum memory.

Maybe it ends up being *faster* in less memory.
(This is also what happened to hybrid attacks.)

There's a **lower bound** ≈ 0.2925 for a class of near-neighbor 2-tuple sieving algorithms.

The bound does not apply to 3-tuples etc.

2-tuples: ≈ 0.2925 is 70% of naive ≈ 0.4150 .

3-tuples: ≈ 0.3383 is only 60% of naive ≈ 0.5661 .

Increasing k (for k -tuples) places more emphasis on improving combinatorial enumeration techniques.

20251203: MLWE et al.

2025/2195, Paola de Perthuis, Filip Trenkić,
“Refined Modelling of the Primal Attack, and
Variants Against Module-Learning With Errors”:

Addresses a questionable independence assumption regarding vector norms. Also: “Asymptotically, we show the module primal attack over a degree d number field K has a reduced cost, resulting in a subexponential gain, whenever the discriminant Δ_K satisfies $|\Delta_K| < d^d$, one such case being non-power-two cyclotomics.”

20251203: MLWE et al.

2025/2195, Paola de Perthuis, Filip Trenkić,
“Refined Modelling of the Primal Attack, and
Variants Against Module-Learning With Errors”:

Addresses a questionable independence assumption regarding vector norms. Also: “Asymptotically, we show the module primal attack over a degree d number field K has a reduced cost, resulting in a subexponential gain, whenever the discriminant Δ_K satisfies $|\Delta_K| < d^d$, one such case being non-power-two cyclotomics.”

- Didn't Ducas–Engelberts–de Perthuis say that?
- Sort of. Old: module-BKZ. New: module-primal.

20251210: enumeration

2025/2227, Yuanmi Chen, Zhao Chen, Tingting Guo, Chao Sun, Weiqiang Wen, Yu Yu,
“Time Memory Trade-off For Enumeration”:

“We incorporate a meet-in-the-middle strategy into the enumeration algorithm, enabling a tunable time-memory trade-off. The algorithm can achieve a minimum asymptotic time complexity of $n^{n/(4e)+o(n)}$, which, in return, demands memory of the same order.”

20251210: enumeration

2025/2227, Yuanmi Chen, Zhao Chen, Tingting Guo, Chao Sun, Weiqiang Wen, Yu Yu,
“Time Memory Trade-off For Enumeration”:

“We incorporate a meet-in-the-middle strategy into the enumeration algorithm, enabling a tunable time-memory trade-off. The algorithm can achieve a minimum asymptotic time complexity of $n^{n/(4e)+o(n)}$, which, in return, demands memory of the same order.”

— I'll say more about enumeration later.

20260211: provability

[2026/225](#), Maiara F. Bollauf, Amaury Pouly, Yixin Shen, “Solving SIS in any norm via Gaussian sampling”:

“In particular, we show that its performance is at least 50 bits faster than the version of Wagner’s algorithm of Ducas, Engelberts, and Loyer (Crypto 2025) for solving SIS with the infinity norm across all security levels of Dilithium, but still far from heuristic estimates underlying the security of this scheme.”

20260211: provability

2026/225, Maiara F. Bollauf, Amaury Pouly, Yixin Shen, “Solving SIS in any norm via Gaussian sampling”:

“In particular, we show that its performance is at least 50 bits faster than the version of Wagner’s algorithm of Ducas, Engelberts, and Loyer (Crypto 2025) for solving SIS with the infinity norm across all security levels of Dilithium, but still far from heuristic estimates underlying the security of this scheme.”

— Sacrifices less speed than before for provability.

20260216: MLWE

2026/279, Tabitha Ogilvie, “On the Concrete Hardness Gap Between MLWE and LWE”:

“Our results demonstrate that the widely assumed equivalence between LWE and MLWE in power-of-two cyclotomics does not hold, with real world consequences for deployed schemes.”

20261004 revision says 0.5, 1.4, 2.5 bits less security for ML-KEM-512, ML-KEM-768, ML-KEM-1024 than previous estimates.

20260223: MLWE

[2026/366](#), Jianhua Hou, Haodong Jiang,
“Careful with the Ring: Enhanced Hybrid Decoding
Attacks against Module/Ring-LWE”:
— Similar ideas to [2026/279](#)
(and merged for Crypto 2026).

20260327: a few more bits

[2026/607](#), Xiaohan Zhang, Zijian Zhou, Longjiang Qu, “Refined Approx-SVP Rank Reduction Conditions and Adaptive Lattice Reduction for MSIS Security Estimation”:

“Applied to Dilithium, our analysis indicates that, when compact rank reduction behavior is integrated with the D4f technique, the estimated concrete security of Dilithium is reduced by 3.65-6.09 bits relative to the conservative Core-SVP baseline, providing a more realistic concrete assessment.”

20260518: quantum dual attacks

[2026/979](#), Cong Ling, Hao Yan, Nicholas Zhao,
“Improved Dual Attack and Trapdoor Sampling via
Quantum Rejection Sampling”:

“Substituting this sampler into the dual attack framework results in reduced overall attack-cost estimates. Compared with Pouly and Shen’s modern dual attack under the same parameter choices, our estimates reduce the attack cost by 9, 4, and 13 bits for Kyber-512, Kyber-768, and Kyber-1024, respectively.”

20260518: quantum dual attacks

[2026/979](#), Cong Ling, Hao Yan, Nicholas Zhao,
“Improved Dual Attack and Trapdoor Sampling via
Quantum Rejection Sampling”:

“Substituting this sampler into the dual attack framework results in reduced overall attack-cost estimates. Compared with Pouly and Shen’s modern dual attack under the same parameter choices, our estimates reduce the attack cost by 9, 4, and 13 bits for Kyber-512, Kyber-768, and Kyber-1024, respectively.”

— That’s a misleading comparison.

Review of dual attacks

Dual of G means set of linear maps $\varphi : R \rightarrow \mathbb{Z}$ such that $\varphi(G) \equiv \dots \equiv \varphi(x^{n-1}G) \equiv 0 \pmod{q}$;
i.e., $\varphi(\text{any multiple of } G \text{ in } R) \equiv 0 \pmod{q}$.

If $A = aG + e$ in R/q then $\varphi(A) \equiv \varphi(e) \pmod{q}$.
If φ is short then $\varphi(e) \in \mathbb{Z}$ is small;
try to measure bias of $\varphi(e) \bmod q$.

The dual-attack debate

Can use many vectors φ to amplify the bias,
but non-independence makes this hard to analyze.

See, e.g., [2023](#) Ducas–Pulles: “Does the Dual-Sieve
Attack on Learning with Errors even Work?”

The dual-attack debate

Can use many vectors φ to amplify the bias,
but non-independence makes this hard to analyze.

See, e.g., [2023](#) Ducas–Pulles: “Does the Dual-Sieve Attack on Learning with Errors even Work?”

Various subsequent papers say some dual attacks work: e.g., [saying](#) “the security levels for Kyber-512/768/1024 are 3.5/11.9/12.3 bits below the NIST requirements”.

The dual-attack debate

Can use many vectors φ to amplify the bias,
but non-independence makes this hard to analyze.

See, e.g., [2023](#) Ducas–Pulles: “Does the Dual-Sieve Attack on Learning with Errors even Work?”

Various subsequent papers say some dual attacks work: e.g., [saying](#) “the security levels for Kyber-512/768/1024 are 3.5/11.9/12.3 bits below the NIST requirements”.

Some papers claim worse-but-provable speeds.
Comparisons should include conjectured speeds too.

20260518: quantum dual attacks

2026/984, Clémence Cheviguard, André Schrottenloher, Yixin Shen, “Quantum algorithm for Discrete Gaussian Sampling”:

“Our sampler outputs a quantum state which can either be measured to get the desired distribution or be used directly as such in other quantum algorithms. By doing so, we derive two versions of quantum dual attacks that improve upon the previous ones in [Pouly & Shen, EUROCRYPT 2024].”

20260525: dual attacks

2026/1048, Yechen Li, Qunxiong Zheng, “Unified Dual Attack Analyses: Covariance-Based Score Distribution Prediction for LWE”

“These theoretical gaps motivate us to develop a unified estimation of the expectation and variance of the score in dual attacks.”

20260525: dual attacks

2026/1048, Yechen Li, Qunxiong Zheng, “Unified Dual Attack Analyses: Covariance-Based Score Distribution Prediction for LWE”

“These theoretical gaps motivate us to develop a unified estimation of the expectation and variance of the score in dual attacks.”

— A whole paper devoted to improving *one* aspect of the *analysis* of *one* type of lattice attack.

20260526: dual attacks

[2026/1060](#), Lei Bi, Yijian Liu, Xianhui Lu, Junjie Luo, Kunpeng Wang, “An Improved Hybrid Dual Attack on LWE with Sparse Secrets and its Application to FHE”:

“We demonstrate that, for typical FHE parameters, our attack does not rely on the problematic independence heuristic and lies outside the contradictory regime. Finally, we compare our attack with previous hybrid attacks, showing consistent and significant improvements across all evaluated cases.”

20260702: lower probability

[2026/1364](#), Divesh Aggarwal, Haoxiang Jin, “Time vs Success Probability Tradeoff for SVP and BDD with Implications to LWE and SIS”

20260702: lower probability

[2026/1364](#), Divesh Aggarwal, Haoxiang Jin, “Time vs Success Probability Tradeoff for SVP and BDD with Implications to LWE and SIS”

— I don't see anything new here: it's basically standard enumeration (picking a random node in the enumeration tree) after standard sieving as preprocessing. But note that fast low-probability attacks can make loose “security proofs” vacuous.

20260717: another few bits

[2026/1465](#), Marc Stevens, Michael Yonli, “On Reduction Probability Models in Lattice Sieving”:

“In lattice sieving, the sphere model assumes that sieving elements are distributed close to or on a sphere with uniformly distributed direction. It is extensively used to predict how lattice sieving behaves. . . . We show that the reduction probability advantages of our models over the sphere model range from $\times 1.5$ up to $\times 8$.”

20260803: provability

2026/1587, Yiming Gao, Yansong Feng, Honggang Hu, “Solving the Shortest Vector Problem in $2^{0.7314n+o(n)}$ Time via Discrete Gaussian Sampling on Superlattices”:

“This improves the $2^{n+o(n)}$ classical bound of Aggarwal, Dadush, Regev, and Stephens-Davidowitz (ADRS, STOC 2015).”

20260803: provability

2026/1587, Yiming Gao, Yansong Feng, Honggang Hu, “Solving the Shortest Vector Problem in $2^{0.7314n+o(n)}$ Time via Discrete Gaussian Sampling on Superlattices”:

“This improves the $2^{n+o(n)}$ classical bound of Aggarwal, Dadush, Regev, and Stephens-Davidowitz (ADRS, STOC 2015).”

— But it's worse than algorithms conjectured to have exponent $\approx 0.2925n + o(n)$.

20260803: quantum DCP

[2026/1591](#), Daniel R. Simon, “A Polynomial-Time Quantum Algorithm for the Dihedral Coset Problem”:

“We present a polynomial-time quantum algorithm for the Dihedral Coset Problem (DCP). . . . efficiently solve, for example, SVP with a $\sqrt{n}$ polylog(n) approximation factor, or LWE instances with $\alpha = \sqrt{n}$ polylog(n).”

20260803: quantum DCP

[2026/1591](#), Daniel R. Simon, “A Polynomial-Time Quantum Algorithm for the Dihedral Coset Problem”:

“We present a polynomial-time quantum algorithm for the Dihedral Coset Problem (DCP). . . . efficiently solve, for example, SVP with a $\sqrt{n}$ polylog(n) approximation factor, or LWE instances with $\alpha = \sqrt{n}$ polylog(n).”

— This paper has been disputed.

20261007 paper update continues the debate.

20260804: provability

2026/1597, Minki Hhan, “Solving the Shortest Vector Problem in $2^{0.6039n+o(n)}$ Time via Mid-point Hessian”:

“... improving the previous best algorithm running in $2^{n+o(n)}$ time and space of Aggarwal, Dadush, Regev, and Stephens-Davidowitz [STOC'15]”

— Again ignoring the *conjectured* $2^{\approx 0.2925n+o(n)}$ run time of other methods.

20260804: provability

2026/1599, Jiseung Kim, “One Discrete Gaussian Sample in $2^{n/2+o(n)}$ Time”:

“... an exact-SVP algorithm in $2^{0.7315n+o(n)}$ time”

20260815: quantum DCP

[2026/1693](#), Aparna Gupte, Seyoon Ragavan, Mark Zhandry, “The ePrint:2026/1591 Quantum Algorithm Does Not Solve DCP”:

“... either the algorithm has expected runtime $\Omega(2^{n/9})$ for at least one of the two secrets, or its eventual distinguishing advantage between the two secrets is at most $\text{poly}(n)2^{-n/9}$ ”

20260818: low-memory searches

2026/1720, Abul Kalam, Sudeshna Karmakar, Santanu Sarkar, “Improved Polynomial-Memory Algorithms via Nested Collision Search: Applications to Small Max-Norm LWE and Subset-Sum”:

“For uniformly random ternary secrets of length n , our best algorithm reduces the polynomial-memory runtime from $2^{0.926n}$ to $2^{0.8595n}$ ”

20260831: provability

2026/1844, Yiming Gao, Yansong Feng, Honggang Hu, “Discrete Gaussian Sampling Meets BDGL Decoding: Solving the Shortest Vector Problem in $2^{0.5596n+o(n)}$ Time”:

“As a result, we obtain a randomized classical algorithm for SVP that runs in $2^{0.5596n+o(n)}$ time and uses $2^{n/2+o(n)}$ space.”

20260902: provability

2026/1859, Minki Hhan, “Finding a Shortest Vector and More in $2^{n/2+o(n)}$ Time using q -ary Coset Difference Tree”:

“Thus, this algorithm solves a closest vector problem on such random instances in time and space $2^{n/2+o(n)}$.”

20260903: parallelization

2026/1876, Renya Hashimoto, Junji Shikata, Atsushi Takayasu, Yuntao Wang, “AH-BKZ: A Lattice Reduction Algorithm with Asynchronous Hybrid Processing”:

“Our experiments show that, compared with ProPnjBKZ-based preprocessing, AH-BKZ reduces the overall runtime by up to 23.1% and by 18.6% on average among successful runs for dimensions 124-148 under comparable peak memory constraints.”

20260915: ideals

[2026/2042](#), Hongyuan Qu, Guangwu Xu, “On Solving Ideal-SVP and Ideal-BDD with Group Representations”

— Useful only for ideals stabilized by a nontrivial automorphism.

20260918: enumeration

2026/2084, Yansong Feng, Yiming Gao, Jiaqi Liu
“Faster SVP in Polynomial Space”:

“Kannan’s algorithm, as analyzed by Hanrot and Stehlé in 2007, solves the exact Euclidean shortest vector problem in polynomial space and $n^{\frac{n}{2e}+o(n)}$ time. In the classical setting with polynomial space, we obtain the first improvement on this bound via a randomized algorithm that runs in $n^{\frac{n}{4e}+o(n)}$ time.”

20260918: enumeration

2026/2084, Yansong Feng, Yiming Gao, Jiaqi Liu
“Faster SVP in Polynomial Space”:

“Kannan’s algorithm, as analyzed by Hanrot and Stehlé in 2007, solves the exact Euclidean shortest vector problem in polynomial space and $n^{\frac{n}{2e}+o(n)}$ time. In the classical setting with polynomial space, we obtain the first improvement on this bound via a randomized algorithm that runs in $n^{\frac{n}{4e}+o(n)}$ time.”

— Wait, was this really the first improvement?

Flashback to 2020

[2020/707](#), Martin R. Albrecht, Shi Bai, Pierre-Alain Fouque, Paul Kirchner, Damien Stehlé, and Weiqiang Wen, “Faster Enumeration-based Lattice Reduction: Root Hermite Factor $k^{1/(2k)}$ in Time $k^{k/8+o(k)}$ ”:

“We give a lattice reduction algorithm that achieves root Hermite factor $k^{1/(2k)}$ in time $k^{k/8+o(k)}$ and polynomial memory.”

Flashback to 2020, continued

“This improves on the previously best known enumeration-based algorithms which achieve the same quality, but in time $k^{k/(2e)+o(k)}$. A cost of $k^{k/8+o(k)}$ was previously mentioned as potentially achievable (Hanrot-Stehlé’10) or as a heuristic lower bound (Nguyen’10) for enumeration algorithms.”

What does this mean?

Exponents for polynomial-space enumeration:

0.1839, more precisely $1/2e$, claimed in 2007.

0.1250, more precisely $1/8$, claimed in 2020,
matching supposed “lower bound ... for
enumeration algorithms” from 2010.

0.0920, more precisely $1/4e$, claimed in 2026.

What does this mean?

Exponents for polynomial-space enumeration:

0.1839, more precisely $1/2e$, claimed in 2007.

0.1250, more precisely $1/8$, claimed in 2020,
matching supposed “lower bound . . . for
enumeration algorithms” from 2010.

0.0920, more precisely $1/4e$, claimed in 2026.

Can ideas be combined with tuple lattice sieving?

What does this mean?

Exponents for polynomial-space enumeration:

0.1839, more precisely $1/2e$, claimed in 2007.

0.1250, more precisely $1/8$, claimed in 2020,
matching supposed “lower bound ... for
enumeration algorithms” from 2010.

0.0920, more precisely $1/4e$, claimed in 2026.

Can ideas be combined with tuple lattice sieving?

Will there be more enumeration improvements?

e.g. can the representation ideas for $1/4e$ be
combined with the block-size ideas for $1/8$?

20260920: provability

[2026/2117](#), Hongyuan Qu, Chengliang Tian, Geng Wang, Guangwu Xu, “Provable dual attack on LWE via lattice projection”:

“We rigorously analyze the parameter conditions for the success of our framework, and by enumerating all feasible parameters we achieve an improvement of 26.1–43.7 bits over the original provable dual attack framework.”

20260923: provability

[2026/2166](#), Jiale Li, Li-Ping Wang, Huaxiong Wang, “A Provable Correctness Analysis of the MATZOV-Style FFT Dual Attack on ML-KEM”:

“In this work, we give a rigorous correctness analysis of a MATZOV-style FFT-based dual attack for the ML-KEM parameter sets.”

— Again at the expense of claimed speed.

20260923: provability

[2026/2186](#), Jiale Li, Li-Ping Wang, Huaxiong Wang, “Provable FFT-Accelerated Dual Attack on LWE Using the General Dual Lattice”:

“In this work, we provide a provable correctness analysis of the Guo–Johansson-style FFT-based dual attack over the general dual lattice.”

20260926: MLWE

2026/2228, Henry Bambury, Seungki Kim, Changmin Lee, Phong Q. Nguyen, “Adelic reduction of module lattices”:

— Presents another LLL generalization to number fields, improving upon 2019 Lee–Alice Pellet–Mary–Stehlé–Alexandre Wallet.

20260928: quantum DCP

[2026/2249](#), Mathias Boucher, Pierre-Alain Fouque, Yixin Shen, “Cyclotomic Cosets: Hidden Subgroup and Quantum Sieving Algorithm for Prime-Power Moduli”

“... we emphasize that our result does not, by itself, yield a quasi-polynomial-time algorithm for standard LWE, because the currently known reduction produces only a limited number of approximate CCP states.”

[2026/2261](#), Xuyuan Han, Yiming Gao, Honggang Hu, “A Provable Subexponential-Time Algorithm for LWE from $n + o(n)$ Samples via Wagner-Style Gaussian Sampling”:

“For polynomial prime modulus $q = n^{\kappa+o(1)}$ with any fixed $\kappa > 1/2$, the algorithm handles independent discrete Gaussian errors drawn from $D_{\mathbb{Z}, s_e}$, with Gaussian parameter s_e satisfying $\ln(2 + s_e) = (\ln n)^{a+o(1)}$ for any fixed $0 \leq a < 1$.”

The importance of sample counts

2014 Martin R. Albrecht–Carlos Cid–Jean-Charles Faugère–Ludovic Perret: algebraic attacks break constant-error LWE in subexponential time for $\Theta(n \log \log n)$ samples; exponential but competitive with sieving for (e.g.) $7n$ samples with binary errors.

This doesn't sound relevant to, e.g.,
NTRU with n samples, or
the GAM/LPR cryptosystem with $2n$ samples.

The importance of sample counts, part 2

2015 Paul Kirchner–Pierre-Alain Fouque claimed subexponential break of bounded-error LWE for n samples using an improvement of BKW.

The importance of sample counts, part 2

2015 Paul Kirchner–Pierre-Alain Fouque claimed subexponential break of bounded-error LWE for n samples using an improvement of BKW.

2015 Herold–Kirshanova–May: previous claim is flawed; a fix uses $\Theta(n \log n)$ samples.

The importance of sample counts, part 3

2016 Lyubashevsky recommending bounded-error designs: “For efficiency, want s, e, e_1, e_2 to be as small as possible. But [AG ‘11] says that if they are too small, then (Ring)-LWE is easy. But ... the attack in [AG ‘11] requires many linear equations – in the cryptosystem, we only have $2n$ equations. So, is it safe to take very small (say 0/1) coefficients if q is not too large? We thought so. And later, some evidence appeared: [MP ‘13] says that it is safe to use smaller LWE coefficients if there are few samples. [DM ‘13, MP ‘13] say that taking secret/errors from a non-Gaussian distribution is OK.”

The importance of sample counts, part 4

2025 Léo Ducas–Lynn Engelberts–Johanna Loyer: a variant of BKW solves ∞ -SIS for $n + o(n)$ variables with norm bound $q/\text{polylog}(n)$ in subexponential time, but still doesn't give “a subexponential-time algorithm for LWE with narrow secrets”.

What 2026/2261 claims

2026/2261 is another BKW variant. Says it breaks bounded-error LWE in subexponential time with only $n + o(n)$ samples.

What 2026/2261 claims

2026/2261 is another BKW variant. Says it breaks bounded-error LWE in subexponential time with only $n + o(n)$ samples.

Paper says it doesn't apply to RLWE etc.

— Why?

— Because the paper's analysis considers *uniform random matrices*. RLWE matrices aren't random.

— But there's no reason to expect those matrices to be stronger than random matrices.

20260930: provability

[2026/2271](#), Yang Zhou, Zhedong Wang, Mingsheng Wang, “Exact SVP on Haar-Random Lattices at the Heuristic Sieving Exponents”:

— Reports proof of ≈ 0.2925 for a modified 2-tuple sieving algorithm for almost all lattices.

20260930: provability

[2026/2271](#), Yang Zhou, Zhedong Wang, Mingsheng Wang, “Exact SVP on Haar-Random Lattices at the Heuristic Sieving Exponents”:

— Reports proof of ≈ 0.2925 for a modified 2-tuple sieving algorithm for almost all lattices.

This description also applies to [20260926](#) paper from Emmanouil Doulgerakis, Thijs Laarhoven.

2026/2289, Zhili Chen, Kaijie Jiang, Anyu Wang, Chuanqi Zhang, “A Subexponential Algorithm for Binary LWE with Sublinear Samples”:

“Fix any constant $a > 1/2$, let $q = n^{a+o(1)}$ be an odd prime, and let $m = \lceil n / \log \log n \rceil$. We give a randomized algorithm that distinguishes $(A, A^T s + e) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$ from a uniformly random pair, for every fixed secret–error pair satisfying $\|s\|_2^2 + \|e\|_2^2 = O(m + n)$. Its expected running time and space are $\exp(O(n / \log \log n))$.”

20261007: it never ends

2026/2386: similar to 2026/2289.

2026/2388: exponent ≈ 0.6078 for γ -SVP for a large constant γ , but conjecturally exponent ≈ 0.3066 .

Worse than what 2026/2271 says for almost all lattices, but maybe better for occasional lattices.

Conclusions

Answers if the papers are correct

“Subexponential for LWE with polynomial modulus”? — Yes, for typical case of bounded errors and common sample counts.

Answers if the papers are correct

“Subexponential for LWE with polynomial modulus”? — Yes, for typical case of bounded errors and common sample counts.

Reduced “exponent for SVP/BDD” at relevant approx factors? — Yes for low memory; for medium memory; for high memory with bounded errors etc.

Answers if the papers are correct

“Subexponential for LWE with polynomial modulus”? — Yes, for typical case of bounded errors and common sample counts.

Reduced “exponent for SVP/BDD” at relevant approx factors? — Yes for low memory; for medium memory; for high memory with bounded errors etc.

20-bit security loss in ML-KEM-768 compared to “the FIPS 203 analysis”? — Unclear question. FIPS 203 doesn't contain an analysis.

Answers if the papers are correct

“Subexponential for LWE with polynomial modulus”? — Yes, for typical case of bounded errors and common sample counts.

Reduced “exponent for SVP/BDD” at relevant approx factors? — Yes for low memory; for medium memory; for high memory with bounded errors etc.

20-bit security loss in ML-KEM-768 compared to “the FIPS 203 analysis”? — Unclear question. FIPS 203 doesn't contain an analysis.

ML-KEM was “fully vetted”? — Obviously not.

Recommendations

Keep ECC as a backup layer. Reduces risks from cryptanalysis and from [software bugs](#).

Recommendations

Keep ECC as a backup layer. Reduces risks from cryptanalysis and from [software bugs](#).

Use the largest parameters that the application can afford. Encourage design teams to support much larger parameters. Remember that [eprint 2010/613](#) *didn't* sound crazy when it said its dim-256 proposal would take “about 2^{150} operations” to break.

Recommendations

Keep ECC as a backup layer. Reduces risks from cryptanalysis and from [software bugs](#).

Use the largest parameters that the application can afford. Encourage design teams to support much larger parameters. Remember that [eprint 2010/613](#) *didn't* sound crazy when it said its dim-256 proposal would take “about 2^{150} operations” to break.

Keep investing in alternatives to reduce the damage in case of even more severe security losses.