WORKSHOP ON ATTACKS IN CRYPTOGRAPHY 8 · UCSB

On AI & Cryptanalysis

HAWK “Guessing Game” is not Polynomial-Time / Other Recent AI Adventures

mjos.fi

https://mjos.fi/doc/20260816-wac8-ai-cryptanalysis.pdf

Markku-Juhani O. Saarinen
Tampere University • Finland
markku-juhani.saarinen@tuni.fi

TLP:CLEAR

CONTENTS

001 Introduction

002 Rise and fall of HAWK

003 McEliece

004 Some Conclusions

SECTION

001

INTRODUCTION

AI CRYPTANALYSIS -VS- ML CRYPTANALYSIS

ML Cryptanalysis trains a Deep Neural Network (DNN), Transformer, or other ML model directly with *plaintext / ciphertext / key material samples*, and then challenges the model with a key recovery task or a distinguishing task.

Examples: Gohr's 2019 SPECK attack [12], Salsa LWE attacks from 2022 [23].

DEFINITION (CASUAL CURRENT USE)

AI Cryptanalysis uses Large Language Models (LLMs) and related methods to analyze *cryptographic algorithms* to discover and to implement cryptanalytic attacks, methods, and related abstract theory.

Examples: Anthropic's 2026 HAWK attack [21], some results in this talk.

FRONTIER MODELS ARE REALLY GOOD NOW

Around mid-2025, I was struck by respected cybersecurity professionals admitting that frontier models are better at finding vulnerabilities than they are.

Many contributing factors for radical improvements in cryptographic abilities:

- Chain-of-Thought (intermediate reasoning), and tool use to do math.
- Math is ideal for Reinforcement Learning with Verifiable Rewards (RLVR).
- Mixture of Experts (MoE) or Model Routing, often with cryptography now.

  *Trained on, among other things, **all** IACR ePrints, arXiv, GitHub repositories, etc. – the models have extremely broad knowledge of math and crypto.*
- Quite obvious: Exponential spending on data centers & model training.

IN CASE YOU HAVE BEEN HIDING UNDER A ROCK

YES, EVERYONE CODES WITH "AI" NOW

Case: This year we did Python & Rust implementations of three new (NIST MPTC candidate) schemes with Monash PhD students (Nikai, Yini, and Arzaki):

- LoTRS: Practical Post-Quantum Structured Threshold Ring Signatures from Lattices
- Lemur: Scalable Post-Quantum Synchronized Multi-Signatures
- LaTSEC: Lattice-Based Threshold Signature from Expander Codes

> **CRYPTO CODE DEVELOPMENT GOES FROM MONTHS → DAYS**
>
> Each was 1–2 weeks of semi-autonomous prompting of GPT and Opus.
> I estimate that a similar outcome would have taken $5\times$ longer without AI.

ADD ANY AMOUNT OF TESTING AND VERIFICATION

Base level – note how there are *two* implementations of every scheme:

1. Python “executable pseudocode” – a deterministic, unoptimized, readable model that maps directly to the paper. Used to create unit-level test vectors.
2. An optimized, native Rust implementation that may use machine intrinsics or Rayon multi-threading. Checked to reproduce the same test vectors.

Plus a lot of unit tests! I’d expect that even casual coding can now move to full formal models and logical equivalence checking – previously “too laborious”.

Hardware too: I recently vibe-coded a vector processor (Verilog/FPGA). A key enabler was the formal Sail model of RISC-V ISA. https://karucore.com/

LET'S USE AI TO RUN CRYPTANALYSIS EXPERIMENTS

```
read 2026-1318.pdf, check for correctness, and write a Sage prototype!
```

- In reality, you can't reliably "one-shot" anything very advanced. It's an interactive process where you drive the activities; ask for code, tests, etc.
- I'm a terminal guy and use CLI coding harnesses almost exclusively: Claude Code, Codex, Antigravity, OpenCode – the CLIs are very similar.
- A sandboxed environment with the necessary tools for the models to use; Full SageMath, Rust toolchain, Lean, etc + a big stash of research PDFs.
- I've found it **extremely** useful to have different (vendor) models cross-check each other. Just like in "pair programming" – they have different training.

WHAT IS ALL OF THIS TALK OF “LEAN” ?

```
open Complexity OpenAI -- NP-hard N^c-GapCVP for fixed c < 1/2
theorem hardnessBelowHalf (c : Real) (hc : 0 < c)
   (h : c < (1 : Real) / 2) :
   GapCVP.NPHardPromise (integerTargetGapPromiseAt c hc.le) := by
  let s := Classical.choice (exists_axesExponentSchedule hc h)
  let R := concreteIntegerTargetUniformReductionCertificate c hc.le s
  exact isNPHard_of_integerTargetUniformReductionCertificate c hc.le R
```

The rise of AI in Mathematics happily coincides with the rise of general machine-checkable math proofs and especially **Lean**: https://lean-lang.org/

Lean itself is not AI – it is a programming language with a carefully written logic kernel and libraries that formalize a lot of mathematics. Standard crypto things like Game-based proofs are still evolving (see HOPSCOTCH [8] and VCVio [22].)

SECTION

002

RISE AND FALL OF HAWK

HAWK – A PQC SIGNATURE SCHEME

HAWK [3, 7] was a NIST PQC Signature standardization (“On-Ramp”) candidate until July 29, 2026. Many saw HAWK as a successor algorithm to Falcon [9] (FIPS 206 / FN-DSA), although it is *based on different security foundations*.

Attractive: Fast signing & verification, 555-byte signatures, no floating point.

> QUOTE FROM NIST IR 8610, MAY 2026 [1]
>
> *NIST selected HAWK for the third round due to its strong performance and because its implementation only needs integer arithmetic operations. NIST encourages further analysis of HAWK’s security assumptions, particularly the smLIP problem within the specific structure of cyclotomic number fields.*

STILL A NEW HARD PROBLEM

HAWK SECURITY ASSUMPTIONS

HAWK was *inspired* by the Lattice Isomorphism Problem (LIP). Over time, HAWK's proofs became more rigorous, but the concrete hardness foundation became less clean.

Layer	2022 ASIACRYPT Hawk	2025 2nd Round Hawk v1.1
Key recovery	Rank-2 search module-LIP (smLIP).	More specific smLIP in a power-of-two cyclotomic CM field.
Unforgeability	Introduced: One more short vector problem (omSVP) game.	Patched – more public symmetries declared trivial.
Parameters	BKZ experiments on unusual-SVP and approximate-CVP	Extended – dimensions-for-free and alternative models.

CRYPTANALYSIS OF HAWK: A GUESSING GAME

Paper: IACR ePrint *"Cryptanalysis of HAWK: a Guessing Game"* came out on 2026.06.25 [18].

Claim: HAWK secret key can be recovered in probabilistic polynomial time. Analysis assumes four *"very plausible"* number-theoretic heuristics.

Despite the title, the abstract stated *"At the time of writing, we do not claim that HAWK is broken, as we have not yet verified these heuristics experimentally."*

→ **Let's try to do an AI-assisted implementation!**

PERHAPS FEW HUNDRED PROMPTS OVER 48 HOURS

IMPLEMENTING “GUESSING GAME”

OpenAI and Anthropic try to make it sound as if research science is done with single prompts, like *“Solve this Conjecture. Don’t even think of giving up!”*

However, I approached the task more like my coding or hardware projects:

1. First extract pseudocode and implementation requirements; plan what is needed to have SageMath scripts that model the **GuessingGame** algorithm.
2. Proceed with implementation in small steps, cross-checking findings with unit tests. Used sanity checks such as “Toy Hawk” reference code.
3. Maintain documentation throughout. My final goal was a fast, parallel implementation once the basic attack mechanism was validated with Sage.

STRUCTURE OF THE "GUESSING GAME" ALGORITHM

The GuessingGame attack algorithm has a randomized loop-until-succeed structure with early aborts ("filtering"). Some fairly advanced tools/steps are used:

- The reduction of HAWK to the principal ideal problem in a quaternion algebra from EUROCRYPT 2025 [5].
- CM-order principal-ideal computation with Lenstra-Silverberg [17].

Both of these subroutines run in classical probabilistic polynomial time (PPT).

Since the rest of the steps in the loop are also polynomial, the overall complexity is polynomial **iff the number of iterations is polynomially bounded**.

The probability of passing stages of the algorithm is argued with heuristics.

IMPLEMENTING “GUESSING GAME”

After two days, we had modeled much of the algorithm in SageMath using Opus 4.8 and GPT-5.5, and were testing the success probability of Heuristics 3 and 4.

The paper claims a polynomial $O(n^{3+o(1)})$ number of Lenstra–Silverberg calls. This is related to h_F, the number of ideal classes in the CM field $F = K(\sqrt{\alpha})$.

n	Sample size	Experimental success	Geomean h_F
8	2467	17/2467 = 0.0069	242
16	51	0/51 = 0.0000	87330

Experimentally, doubling dimension n raised h_F by a factor $87330/242 \approx 2^{8.5}$. Absent extraordinary cancellation, our own heuristic is $h_F = 2^{\Theta(n \log n)} = n^{\Theta(n)}$.

CULPRIT FOUND

GUESSING GAME – HEURISTIC 4

After a while, the issue was triaged to Heuristic 4 and its "$O\left(n^{1+o(1)}\right)$ ideals".

HEURISTIC (NUMBER 4 IN THE PAPER)

Let $K = \mathbb{Q}(\zeta_{2^l} + \zeta_{2^l}^{-1})$ and suppose that F is a field isomorphic to $K(\sqrt{\alpha})$. During the main algorithm [18, Algorithm 2], one forms an ideal $\mathfrak{a} \subseteq \mathcal{O}_F$ that depends on both q' and α, and $N_{F/K}(\mathfrak{a}) = q'\mathcal{O}_K$. In [18, Section 9] the authors show that, with overwhelming probability, there are $O\left(n^{1+o(1)}\right)$ ideals of $\mathcal{O}_F$ with relative norm $q'\mathcal{O}_K$. Given that one of these ideals is principal, the probability that $\mathfrak{a}$ itself is principal is $\Omega\left(n^{-1-o(1)}\right)$.

ANTHROPIC'S ATTACK ON HAWK

On **2026.07.28**: Z. Straznickas and S. A. Weis from Anthropic released: *"HAWK-n Key Recovery Reduces to SVP in Dimension n/2 + 1"* [21].

A polynomial reduction from HAWK-n to poly(n) exact Shortest Vector Problem (SVP) oracle calls in dimension $n/2+1$ (here n is the ring degree). With SVP, it is still exponential, but the HAWK v1.1's BKZ analysis attacked a rank-$2n$ lattice.

> "I KNOW THIS!"
>
> The attack combines two well-known cyclotomic involutions to expose a previously hidden HAWK automorphism as a shortest vector in a public rank-n cocycle lattice, then applies Ducas reduction [6] and van Gent–Pulles descent [10]: *Most ingredients were well known to the HAWK Team*.

RECAP

HAWK TIMELINE

- **2022.12.08**: Ducas, Postlethwaite, Pulles, van Woerden. *Hawk: Module LIP Makes Lattice Signatures Fast, Compact and Simple* [7]. ASIACRYPT '22.
- **2026.05.14**: Hawk selected to 3rd round of NIST PQC Signature On-Ramp.
- **2026.06.25**: Nelson, Limbrey, Ling, Mendelsohn. *Cryptanalysis of HAWK: a Guessing Game* [18]. Authors conjecture a classical polynomial-time attack.
- **2026.07.05**: Saarinen. *HAWK "Guessing Game" is not Polynomial-Time* [19]. AI-Assisted implementation attempt/analysis. (Also: This talk submitted.)
- **2026.07.28**: Straznickas, Weis. *HAWK-n Key Recovery Reduces to SVP in Dimension n/2 + 1* [21]. – Anthropic's AI cryptanalysis announcement.
- **2026.07.29**: HAWK Withdrawn from the NIST Competition by its authors.

WHILE ON THE TOPIC OF SVP

CAN WE MAKE SIMON'S DCP ALGORITHM WORK?

In early August 2026, Daniel R. Simon published a *"Polynomial-Time Quantum Algorithm for the Dihedral Coset Problem"* as an IACR ePrint [20].

Q: Can you formalize it?

A: The algorithm, but not the proof (at least Lemmas 3 and 4 look wrong.)

Q: Can the algorithm be made to work anyway?

A: Nope, or even modified slightly – see below.

Q: is poly-time quantum DCP possible at all?

A: Genuinely open.

Aparna Gupte, Seyoon Ragavan, and Mark Zhandry: *"The ePrint:2026/1591 Quantum Algorithm Does Not Solve DCP"* [13]. This is an impossibility result.

SECTION

003

MCELIECE

IN SO MANY WAYS

2026: THE YEAR OF MCELIECE

Classic McEliece was standardized in ISO/IEC 18033-2:2006/Amd 2:2026 [16].

Let's look at some algebraic ("structural") attacks (two here at CRYPTO 2026*):

- Pfaffian modeling (Bardet-Lemoine-Tillich [2]* – best paper).
- Higher-Order Vanishing / distinguisher (Hemmert-Wiemers [15]*).
- Higher-Order Vanishing / key recovery (Hemmert [14]).
- Rank-2 "Subexponential" (Briaud-Lemoine-Randriambololona-Tillich [4]).
- Holdout / PIR "Quasipolynomial" (Ghoshal-Ishai-Jain-Sun [11]).

Note: Most of these attacks are above 2^{256} against Classic McEliece parameters. AFAIK, the CM team has not yet commented on the latest Holdout / PIR attack.

FAVOURABLE REGIMES FOR STRUCTURAL ATTACKS

OBSERVATIONS

- These attacks are all conditional on parameter ranges (or even keys!)
- There is no single "input size" security parameter for asymptotics.
- A distinguisher does not directly imply decryption or OW-CPA break.

Attack	**Type**	**Indicative Conditions** *(very rough)*	**Asymptotics**
Pfaffian [2]	key recovery	Very high rate; square defect visible	poly in (m, r)
HOV [15]	distinguisher	Random HOV kernel dies; Goppa not	–
HOV [14]	key recovery	Clean Frobenius-split HOV kernel	–
Rank-2 [4]	key recovery	Rank–2 relation at low solving degree	subexp $2^{o(r)}$
Holdout [11]	distinguisher	Kernel survives past Hermite threshold	qpoly $n^{O(\log n)}$

DOES “HOLDOUT / PIR” BREAK CLASSIC MCELIECE?

The main result of [11] is a very interesting public-key distinguisher. Distinguishers do not threaten the security Classic McEliece KEM (not assumed in proofs.)

The paper sketches decryption heuristics, but no real “end-to-end” pipeline even for toy parameters. The paper lists this among its Open Problems:

> *“**Can the decryption heuristic be proved or refuted?** [..] Further validation of the underlying conjectures, or ideally an unconditional proof of quasipolynomial decryption, are left for future work. We could not run an end-to-end test on binary Goppa codes [..]”*

Status: It’s a conjecture. I have not yet found super convincing evidence that [11] works against Classic McEliece, nor clear refuting evidence against it either.

TII CHALLENGES

The Technology Innovation Institute (UAE) ran a McEliece challenge in 2023-24.

TII offered cash prizes (the reward for TII-70 and above was $ 10,000). Too bad it has ended – the new AI algorithms would have paid for a lot of AI tokens :-)

Note: The TII work factor label is for a "dumb" brute force attack on (n, m, t): Goppa-polynomial search $\times$ support-set search $\times$ $O(n^3)$ candidate verification.

$$\lambda_{\mathrm{TII}} = \log_2 \left(I(t, 2^m) \cdot \binom{2^m}{n} \cdot n^3 \right)$$

Varied (n, m, t) helped motivate research into different parameter regimes.
https://github.com/ElenaKirshanova/tii_decoding_challenge

HOVER: A HIDDEN FIELD INSIDE THE PUBLIC KEY

HOVER (Higher-Order Vanishing Endomorphism Recovery) is an extended variant of the HOV Key Recovery attack [14, 15], discovered via our AI experiments.

Idea: HOV yields an m-dimensional matrix space from a public key. Useful combinations lie along m secret directions, permuted cyclically by squaring.

Rather than search for those directions, ask which binary matrices act on the coefficients without carrying any relation outside the relation space:

$$\mathcal{E} = \{B \in \mathbb{F}_2^{m\times m} : B \text{ sends every relation to a relation}\} \;\cong\; \mathbb{F}_{2^m}.$$

Linear algebra finds $\mathcal{E}$. When derivative images split and their kernels are distinct, a field generator's eigenlines are exactly the hidden directions.

HOVER: NEW TII KEY RECOVERY RECORDS

HOVER turns out to be quite effective (against specific low-degree parameters.) Native multi-threaded implementation. Verified secret keys with official scripts.

Challenge	m	r	n	Wall Time	Memory	System
TII-252	10	11	1008	6 min 15 s	2.7 GiB	12-core laptop
TII-246	10	11	1009	3 min 11 s	2.7 GiB	12-core laptop
TII-240	10	11	1010	3 min 06 s	2.7 GiB	12-core laptop
TII-213	9	10	496	6 h 45 min	144 GiB	28-vCPU server
TII-129	9	9	509	4 h 30 min	80 GiB	28-vCPU server

Evidence (secret keys): https://github.com/mjosaarinen/tii-solved

Note: TII-83 was the “highest” challenge broken by the May 2024 deadline.

SECTION

004

SOME CONCLUSIONS

REACTIONS: CONG LING ON NIST PQC FORUM

Dear Colleagues,

This is a historic moment in cryptology: AI has beaten humans at cryptanalysis. We would like to share a few thoughts with the PQC community based on our own experience.

There are other dimension-reduction attacks on HAWK in literature. At Eurocrypt 2026, we proposed a dimension-halving algorithm for quaternion Ideal-SVP, which potentially implies that HAWK-n key recovery can be reduced to (cyclotomic) Ideal-SVP in dimension n.

Later, we proposed another guessing attack based on advanced number theory, under a few heuristic assumptions. However, one of these heuristics was soon found to be invalid by colleagues with the assistance of AI.

Apparently, we were outpaced by AI in both cases. A couple of lessons we have learned are: (a) humans can make mistakes; and (b) humans are slower than AI. Beyond cryptanalysis itself, the use of AI to check mathematical proofs is becoming an increasingly serious issue.

One can even imagine a crisis in cryptography: in the worst-case scenario, all human-designed cryptosystems are broken by AI. Whether or not that happens, a new era – Post-AI Cryptography – has arrived.

THE PROOF CHECKING BURDEN

Peter Scholze on the Liquid Tensor Experiment (June 2021):

> *Question: So, besides the authors of course, who understands the proof now?*
>
> *Answer: I guess the computer does, as does Johan Commelin.*

This was the situation with advanced mathematics already **before** advanced AI.

2026: Full versions of some conference papers are 80 pages with 50-page proof appendices that reviewers were not required to read. 30 pages are published.

Near Future: AI allows more proof coverage – which should be a *good thing*. I'd predict conferences that already require implementation artifacts also to require machine-checkable proof artifacts (if the proof is a key contribution.)

REMINISCENT OF USG’S EFFORTS IN 1990S TO “CONTROL CRYPTOGRAPHY PROFILERATION”

Claude

Hello,

Thank you for submitting your application for the Cyber Verification Program. Upon reviewing the details of your submission, we have adjusted the safeguards applied to your account to allow for the cyber use cases you described.

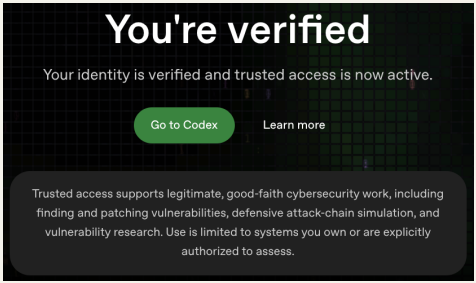

.. but very often, when working on cryptography research ..

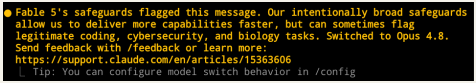

MY OWN MODEST CONCLUSIONS

The way mathematics & engineering is done has changed radically in a year.
Learning and adaptation required for many. Curriculum rethink at Universities.

AI can't find bugs that aren't there – security improves.
Software patching is currently at an all-time high (cryptography not immune).

Powerful, open mathematics and cryptography models are needed.
Open science needs to resist corporate and government control.

AI helps prove or disprove statements that humans have not been able to.
Some of the undiscovered theory and security proofs to defend against "Post-AI cryptanalysis" may turn out to be very elegant. But some of it may be so complex it is basically beyond human comprehension. Yet those proofs are necessary!

END OF TALK!

QUESTIONS

https://mjos.fi/doc/20260816-wac8-ai-cryptanalysis.pdf

REFERENCES (1)

[1] Gorjan Alagic, Maxime Bros, Pierre Ciadoux, Quynh Dang, Thinh Dang, John Kelsey, Jacob Lichtinger, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Hamilton Silberg, Daniel Smith-Tone, and Noah Waller. "Status Report on the Second Round of the Additional Digital Signature Schemes for the NIST Post-Quantum Cryptography Standardization Process". Interagency or Internal Report NIST IR 8610. National Institute of Standards and Technology, May 2026, p. 38. DOI 10.6028/NIST.IR.8610.

[2] Magali Bardet, Axel Lemoine, and Jean-Pierre Tillich. "An Attack on the CFS Scheme and on TII McEliece Challenges". Cryptology ePrint Archive, Paper 2026/430. 2026. URL: https://eprint.iacr.org/2026/430.

[3] Joppe W. Bos, Olivier Bronchain, Léo Ducas, Serge Fehr, Yu-Hsuan Huang, Thomas Pornin, Eamonn W. Postlethwaite, Thomas Prest, Ludo N. Pulles, and Wessel van Woerden. "HAWK version 1.1 (February 5, 2025)". Feb. 2025. URL: https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/hawk-spec-round2-web.pdf.

[4] Pierre Briaud, Axel Lemoine, Hugues Randriambololona, and Jean-Pierre Tillich. "A Heuristic Subexponential Attack on the McEliece Cryptosystem". Cryptology ePrint Archive, Paper 2026/1232. 2026. URL: https://eprint.iacr.org/2026/1232.

REFERENCES (2)

[5] Clémence Chevignard, Guilhem Mureau, Thomas Espitau, Alice Pellet-Mary, Heorhii Pliatsok, and Alexandre Wallet. “A Reduction from Hawk to the Principal Ideal Problem in a Quaternion Algebra”. In: *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part II*. Ed. by Serge Fehr and Pierre-Alain Fouque. Vol. 15602. Lecture Notes in Computer Science. Springer, 2025, pp. 154–183. DOI 10.1007/978-3-031-91124-8_6. URL: https://eprint.iacr.org/2025/287.

[6] Léo Ducas. “Provable lattice reduction of $\mathbb{Z}^n$ with blocksize $n/2$”. In: *Des. Codes Cryptogr.* 92.4 (2024), pp. 909–916. DOI 10.1007/S10623-023-01320-7.

[7] Léo Ducas, Eamonn W. Postlethwaite, Ludo N. Pulles, and Wessel P. J. van Woerden. “Hawk: Module LIP Makes Lattice Signatures Fast, Compact and Simple”. In: *Advances in Cryptology - ASIACRYPT 2022 - 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5-9, 2022, Proceedings, Part IV*. Ed. by Shweta Agrawal and Dongdai Lin. Vol. 13794. Lecture Notes in Computer Science. Springer, 2022, pp. 65–94. DOI 10.1007/978-3-031-22972-5_3. URL: https://eprint.iacr.org/2022/1155.

REFERENCES (3)

[8] Stefan Dziembowski, Grzegorz Fabiański, Daniele Micciancio, and Rafał Stefański. “Game Hopping in Lean”. 2026. arXiv: 2608.06261 [cs.CR]. URL: https://arxiv.org/abs/2608.06261.

[9] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. “Falcon: Fast-Fourier Lattice-based Compact Signatures over NTRU (Specification v1.2 — 01/10/2020)”. Oct. 2020. URL: https://falcon-sign.info/falcon.pdf.

[10] Daniël M. H. van Gent and Ludo N. Pulles. “HAWK: Having Automorphisms Weakens Key”. In: *IACR Commun. Cryptol.* 2.2 (2025), p. 20. DOI 10.62056/A3QJP2W9P.

[11] Ashrujit Ghoshal, Yuval Ishai, Aayush Jain, and Nuozhou Sun. “Quasipolynomial Cryptanalysis of the McEliece Cryptosystem (or: PIR Meets McEliece)”. Cryptology ePrint Archive, Paper 2026/1630. 2026. URL: https://eprint.iacr.org/2026/1630.

[12] Aron Gohr. “Improving Attacks on Round-Reduced Speck32/64 Using Deep Learning”. In: *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11693. Lecture Notes in Computer Science. Springer, 2019, pp. 150–179. DOI 10.1007/978-3-030-26951-7_6. URL: https://eprint.iacr.org/2019/037.

REFERENCES (4)

[13] Aparna Gupte, Seyoon Ragavan, and Mark Zhandry. “The ePrint:2026/1591 Quantum Algorithm Does Not Solve DCP”. Cryptology ePrint Archive, Paper 2026/1693. 2026. URL: https://eprint.iacr.org/2026/1693.

[14] Tobias Hemmert. “Key Recovery for the McEliece Cryptosystem Using Higher-Order Vanishing”. Cryptology ePrint Archive, Paper 2026/1339. 2026. URL: https://eprint.iacr.org/2026/1339.

[15] Tobias Hemmert and Andreas Wiemers. “Distinguishing Goppa Codes Using Higher-Order Vanishing”. Cryptology ePrint Archive, Paper 2025/1661. 2025. URL: https://eprint.iacr.org/2025/1661.

[16] ISO. “Information technology – Security techniques – Encryption algorithms – Part 2: Asymmetric ciphers – Amendment 2”. en. Standard ISO/IEC 18033-2:2006/Amd 2:2026. International Organization for Standardization, 2026. URL: https://www.iso.org/standard/86890.html.

[17] Hendrik W. Lenstra Jr. and Alice Silverberg. “Testing Isomorphism of Lattices over CM-Orders”. In: *SIAM J. Comput.* 48.4 (2019), pp. 1300–1334. DOI 10.1137/17M115390X. URL: http://arxiv.org/abs/1706.07373.

REFERENCES (5)

[18] Ben Nelson, Joshua Limbrey, Cong Ling, and Andrew Mendelsohn. “Cryptanalysis of HAWK: a Guessing Game”. Cryptology ePrint Archive, Paper 2026/1318. June 2026. URL: https://eprint.iacr.org/2026/1318.

[19] Markku-Juhani O. Saarinen. “HAWK “Guessing Game” is not Polynomial-Time”. Cryptology ePrint Archive, Paper 2026/1377. 2026. URL: https://eprint.iacr.org/2026/1377.

[20] Daniel R. Simon. “A Polynomial-Time Quantum Algorithm for the Dihedral Coset Problem”. Cryptology ePrint Archive, Paper 2026/1591. 2026. URL: https://eprint.iacr.org/2026/1591.

[21] Zygimantas Straznickas and Stephen A. Weis. “HAWK-n Key Recovery Reduces to SVP in Dimension $n/2 + 1$”. Cryptology ePrint Archive, Paper 2026/1593. 2026. URL: https://eprint.iacr.org/2026/1593.

[22] Devon Tuma, Quang Dao, James Waters, Alexander Hicks, and Nicholas Hopper. “VCVio: Verified Cryptography in Lean via Oracle Effects and Handlers”. Cryptology ePrint Archive, Paper 2026/899. 2026. URL: https://eprint.iacr.org/2026/899.

REFERENCES (6)

[23] Emily Wenger, Mingjie Chen, Francois Charton, and Kristin Lauter. “SALSA: attacking lattice cryptography with transformers”. In: *Proceedings of the 36th International Conference on Neural Information Processing Systems*. NeurIPS ’22. New Orleans, LA, USA: Curran Associates Inc., 2022. DOI 9781713871088. URL: https://eprint.iacr.org/2022/935.