Notes on a recent claim that a `mceliece348864` distinguisher uses only 2^{529} operations

Classic McEliece Team

17 April 2025

Contact address: `authorcontact-mceliece-merged@box.cr.yp.to`

There is a long history of enthusiastic papers on "algebraic" attacks. For example, a 2002 paper [8] says that its attacks are "an important threat for ciphers" such as AES and may have cost "subexponential when the number of rounds grows", which "would be already an important breakthrough". Papers investigating algebraic attacks against the McEliece cryptosystem include [13] from 2010 (describing itself as a "breakthrough in code-based cryptography"); [10]; [12]; [17]; [11]; [9]; [1]; and [16], a 2024 paper "The syzygy distinguisher".

This report analyzes specifically what the claims in [16] mean for Classic McEliece, assuming the claims are correct. The conclusion of this report is that [16] does not affect any of the Classic McEliece security claims. The Classic McEliece security analysis avoids relying on the indistinguishability property targeted by [16], namely indistinguishability of public keys from uniform random matrices. Furthermore, the distinguisher in [16] has vastly higher cost than any of the security levels claimed by Classic McEliece, so it would still not affect the system even if it could somehow be converted into a key-recovery attack.

The rest of this report is organized into the following sections:

1. This distinguisher has cost far above 2^{256}.
2. Classic McEliece avoids relying on this indistinguishability property.
3. Notes on the conjectures in [16].
4. Further comments.
5. Summary of impact.

1 This distinguisher has cost far above 2^{256}

The syzygy distinguisher is, according to the estimates in [16], much slower than a ludicrously expensive key-recovery attack already explained in the Classic McEliece documentation.

Specifically, an attacker can recover private keys in at most 2^{256} key-generation operations by searching through all 256-bit seeds. This attack is noted in, e.g., the Classic McEliece security guide [6, page 24].

A side effect of recovering a private key is that the public key is distinguished from a uniform random string of the same length with probability very close to 1 (to be precise, probability

$\geq 1-2^{b-B}$, where b and B are the number of bits in private keys and public keys respectively). A much more important effect of recovering a private key is being able to break every ciphertext encapsulated by that public key.

The reason this attack is not of concern is that searching through guesses for 256-bit secrets is infeasible, even with a future quantum computer. As [6, page 3] puts it, "guessing 256-bit secrets" is "not a real-world threat".

For comparison, [16, page 26] considers the smallest proposed Classic McEliece parameters $(3488, 64)$, and conjectures high probability for distinguishing public keys from uniform random matrices using an algorithm with "complexity around 2^{529}" (the main bottleneck being linear algebra acting on vectors of length "around 2^{223}"). Whether or not this conjecture is correct, complexity 2^{529} is much slower than searching through 2^{256} seeds. Unfortunately, the extremely high cost of this distinguisher is not mentioned in [16, pages 1–25].

There are many previous papers on infeasible methods to detect the structure of McEliece keys. The Classic McEliece security guide [6, Section 3.1] says the following:

> For all parameters of interest, the fastest attacks known against the OW-CPA problem treat G as a general matrix with no evident structure. . . . There are also much slower "structural attacks" known that try to exploit the secret structure of G, for example by recovering Goppa-code parameters for G.

Distinguishers costing far above 2^{256}, whether or not they recover keys, are consistent with this statement and with the rest of the security analysis. [6, Section 3.6] already mentions various examples of very slow key-recovery attacks.

The "fastest attacks" comment quoted above is referring to ISD, which takes about 2^{150} bit operations for $(3488, 64)$. See [2] for precise quantification and high assurance regarding the number of bit operations used by many ISD variants.

2 Classic McEliece avoids relying on this indistinguishability property

Beyond being extremely slow, the syzygy distinguisher is attacking a problem that appears in some papers but that the Classic McEliece security analysis explicitly avoids relying on.

A top-down view of the security analysis in the security guide is as follows:

- The Classic McEliece security goal is IND-CCA2. See [6, Section 1]. (We recommend using separate modules for generic transformations that aim at anything beyond an IND-CCA2 KEM; see [6, Section 6].)

- IND-CCA2 attacks much faster than QROM IND-CCA2 attacks would be surprising for the selected hash function (SHAKE256). See [6, Section 5.3.3]. This lets reviewers focus on QROM IND-CCA2 attacks.

- QROM IND-CCA2 security follows tightly from the hypothesis of OW-CPA security (i.e., one-wayness) of the underlying PKE. See [6, Section 5]. After checking this, reviewers can focus on OW-CPA security.
- OW-CPA security of the underlying PKE follows tightly from the hypothesis of OW-CPA security of the original McEliece PKE. See [6, Section 4]. This lets reviewers focus on McEliece OW-CPA security.
- [6, Section 3] reviews the state of the art in McEliece OW-CPA attacks.

This provides a clear structure for Classic McEliece security reviews. The only ways that something can possibly go wrong are some sort of disaster involving SHAKE256, a mistake in the tight reductions, or a better OW-CPA attack against the original McEliece system.

The literature sometimes mentions the following trivial implication: McEliece OW-CPA security follows from OW-CPA security of a uniform random matrix if the McEliece public key is indistinguishable from a uniform random matrix. Of course, finding a distinguishing attack faster than the target OW-CPA security level would render this implication vacuous.

The Classic McEliece security guide notes this motivation for studying distinguishing attacks, but it also notes that this is "not a two-way implication—perhaps there are distinguishers even if OW-CPA is secure". A fast high-probability distinguisher would not change any of the Classic McEliece security claims.

As a (real) pre-quantum analogy for such (hypothetical) distinguishers, consider the following trivial implication: if RSA public keys are indistinguishable from uniform random integers of the same size then OW-CPA for an RSA public key follows from OW-CPA for the same system with a uniform random integer. This is vacuous, since RSA keys are efficiently distinguishable from uniform. OW-CPA is nevertheless a reasonable pre-quantum hypothesis for RSA. The distinguisher does not threaten any of the security properties that are normally claimed for RSA.

In other words, for an algorithm to have an effect on the Classic McEliece security analysis, it would have to be much faster than the [16] distinguisher *and* would have to be a key-recovery attack or something else that breaks OW-CPA.

3 Notes on the conjectures in [16]

The paper [16] claims that the cost of its distinguisher is "subexponential in the error-correcting capability, hence better than that of generic decoding algorithms".

Quantitatively, with the usual choice of parameters (n, t) where t is proportional to $n/\log n$, the cost of ISD is exponential in $n/\log n$ (for almost all matrices, under amply tested assumptions), while [16] claims a distinguishing cost exponential in $n(\log\log n)^3/(\log n)^2$.

(Note that exponential in $n(\log\log n)^3/(\log n)^2$ is slightly subexponential in $n/\log n$. Once n is large enough, the $(\log\log n)^3$ in the numerator is outweighed by the extra $\log n$ in the denominator.)

If this claim is correct then presumably a wider range of algebraic attacks should also try iterating XL to compute higher-order syzygies. However, the question of whether the claim is correct needs further investigation. In particular, while it is clear that the stated distinguisher uses roughly the stated number of operations, it is unclear whether that distinguisher has a non-negligible success probability.

The largest distinguishing experiments reported in the paper (see [16, page 25, bottom]) are for $t = 3$ and n below field size 64. The paper reports that the distinguisher becomes less reliable for n below 56 and completely unreliable for n below 50. The paper explains this as being correlated with n crossing a line in [16, Heuristic 1]. Unfortunately, [16, abstract] does not mention that its claims depend on a new heuristic formulated in [16].

The evidence presented for the heuristic consists of other experiments for $n = 56$ reported in [16, page 34]. Those experiments also detect failures. [16, Heuristic 1] says "high probability" without quantification; there is no model or evidence of how the failure probability scales with n and t.

A testable heuristic would need to start with a quantified model of the failure probability. Tests would then need to be carried out for a parameter range wide enough that severe failures of the model can reasonably be expected to have been detected. Such tests would not be easy for such a slow algorithm.

We emphasize that this report's conclusion—namely, that [16] has no effect on the Classic McEliece security claims—does not rely in any way on the possibility of the distinguisher in [16] having extremely low success probability for large n and t.

4 Further comments

Four specific comments in [16] might seem to indicate that the paper's claims contradict Classic McEliece security claims. We address those comments here. Unfortunately, those comments have not been withdrawn, even though we pointed out all of the following facts in August 2024 in reply to [16]; see [7] and [15].

4.1 Security-proof structure

[16] claims that "a security proof" for the McEliece cryptosystem relies on (1) the assumption of indistinguishability of public keys from uniform random matrices and (2) the hardness of decoding random codes, which the paper says "stands on a firm theoretical ground: the decoding problem for generic linear codes is known to be NP-hard".

It is true that the literature sometimes assumes this indistinguishability property, and sometimes claims connections between NP-hardness and security. However, the Classic McEliece security analysis does not claim any such connections, and does not assume this indistinguishability property. The analysis is instead explicitly founded upon a simpler assumption, OW-CPA, which is a major target of McEliece cryptanalysis. See [6].

4.2 Previous analyses

[16] claims to disprove the belief that "we could possibly have reached the intrinsic complexity of cryptanalysis of this system", and claims that this is "the first time an analysis of the McEliece cryptosystem breaks the exponential barrier".

"Reaction attacks" from the turn of the century (also called "sloppy Alice" attacks; see [14] and [18]) are counterexamples to the "first time" claim: they take polynomial time (much faster than [16]) and are analyses of the McEliece cryptosystem. NTS-KEM and PALOMA are two examples of newer McEliece-based KEMs that had IND-CCA2 efficiently broken by members of the Classic McEliece team using variants of these attacks; see [4] and [3].

Why do these polynomial-time attacks not challenge the strength of the McEliece cryptosystem? Answer: The security property that McEliece targeted in the first place is one-wayness, subsequently named OW-CPA security. This property of the McEliece system has retained practically the same security level against known attacks for half a century despite extensive attempts to break it; see [6]. The way that Classic McEliece obtains QROM IND-CCA2 security by assuming just OW-CPA security (see Section 2) is an important, explicit feature of Classic McEliece, a feature that the attacked systems failed to provide.

For the same reason, a hypothetical fast distinguisher between public keys and uniform random matrices would not be a break of the McEliece system, and would not be a break of Classic McEliece. Even if the asymptotic claims in [16] are correct, they are not reducing the asymptotic OW-CPA security of the McEliece system, and they are not reducing the asymptotic QROM IND-CCA2 security of Classic McEliece. It is important to be clear about which problems are being attacked.

4.3 Previous distinguishers

[16] claims that "previous distinguishers" have "strong regime limitations", making those distinguishers inapplicable to "the codes used in Classic McEliece".

The Classic McEliece documentation has always noted various examples of attacks finding the codes used in Classic McEliece. Again, these key-recovery attacks imply distinguishers (and, more importantly, OW-CPA attacks), so these distinguishers are also applicable to these codes. The reason that the documentation dismisses these attacks is explicitly that the attacks are much slower than the state-of-the-art OW-CPA attacks for "all parameters of interest".

4.4 Security parameters

[16] says "Observe [5, §3.4] that the security parameter in the Classic McEliece system, based on the complexity of generic decoding algorithms, is linear in t ... our complexity is subexponential in the security parameter". Here "[5]" is [6], the Classic McEliece security guide.

There is no "security parameter" in [6]. The Classic McEliece security analysis starts with cryptosystem parameters (n etc.) selected to meet size constraints, and then maps those parameters to attack costs. The Classic McEliece rationale [5, Section 5.2] explains that this parameter-selection mechanism is much more robust than mapping backwards from "security parameters" to cryptosystem parameters, "partly because of small attack improvements over many years and partly because of different models for the costs of operations inside attacks".

Regarding attack costs, the section [6, Section 3.4] cited in [16] is explicitly on "Asymptotic costs of information-set decoding". That section is within [6, Section 3], which is explicitly on "OW-CPA security", i.e., one-wayness. The paper [16] does not improve the asymptotic costs of information-set decoding, and does not attack OW-CPA.

5 Summary of impact

The claims in [16] do not contradict any of the Classic McEliece security claims. There are two independent reasons for this, one quantitative and one structural:

- The costs claimed in [16] are 2^{529} for the smallest proposed Classic McEliece parameters. This is much more expensive than a brute-force search through 256-bit seeds, and much more expensive than ISD.
- The costs are for an algorithm that is merely distinguishing public keys from random, not attacking OW-CPA. The indistinguishability assumption targeted in [16] is not used in the Classic McEliece security analysis; it is even explicitly disclaimed by the Classic McEliece security analysis.

As noted above, even a fast public-key distinguisher would not violate the Classic McEliece security claims. A distinguisher above the target OW-CPA security level is even weaker: it doesn't even render vacuous the aforementioned trivial implication. A hypothetical extension of the distinguisher to an OW-CPA attack would certainly not make it faster.

It is unfortunate that, instead of prominently admitting these limitations, [16] incorrectly suggests that it (1) attacks a problem that Classic McEliece relies upon and (2) is faster than the best previous attacks against Classic McEliece. We promptly responded when [16] appeared, but no errata were issued. Some third parties are now citing [16] as supposedly significant attack progress. In fact, [16] is just the latest in a very long line of failures to break the McEliece system.

References

[1] Magali Bardet, Rocco Mora, and Jean-Pierre Tillich. Polynomial time key-recovery attack on high rate random alternant codes. *IEEE Transactions on Information Theory*, 70(6):4492–4511, 2024. URL: https://arxiv.org/abs/2304.14757.

[2] Daniel J. Bernstein and Tung Chou. CryptAttackTester: high-assurance attack analysis. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology—CRYPTO 2024—44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2024, Proceedings, Part VI*, volume 14925 of *Lecture Notes in Computer Science*, pages 141–182. Springer, 2024. URL: https://cat.cr.yp.to.

[3] Daniel J. Bernstein and Tanja Lange. CCA2 and partial key recovery attack on PALOMA (implementation and specification), 2024. URL: https://groups.google.com/g/kpqc-bulletin/c/h2A4UL9wni0/m/Zd327i3XAgAJ.

[4] Tung Chou. An IND-CCA2 attack against the 1st- and 2nd-round versions of NTS-KEM. In Diana Maimut, Andrei-George Oprina, and Damien Sauveron, editors, *Innovative Security Solutions for Information Technology and Communications—13th International Conference, SecITC 2020, Bucharest, Romania, November 19–20, 2020, Revised Selected Papers*, volume 12596 of *Lecture Notes in Computer Science*, pages 165–184. Springer, 2020. URL: https://tungchou.github.io/papers/ntskem_cca2.pdf.

[5] Classic McEliece Team. Classic McEliece: conservative code-based cryptography: design rationale, 2022. URL: https://classic.mceliece.org/mceliece-rationale-20221023.pdf.

[6] Classic McEliece Team. Classic McEliece: conservative code-based cryptography: guide for security reviewers, 2022. URL: https://classic.mceliece.org/mceliece-security-20221023.pdf.

[7] Classic McEliece Team. Re: Structural analysis of McEliece asymptotically better than generic decoding, 2024. URL: https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/Qj20WzX4dhE/m/74Qbc5fkCgAJ.

[8] Nicolas T. Courtois and Josef Pieprzyk. Cryptanalysis of block ciphers with overdefined systems of equations. In Yuliang Zheng, editor, *Advances in Cryptology—ASIACRYPT 2002, 8th International Conference on the Theory and Application of Cryptology and Information Security, Queenstown, New Zealand, December 1–5, 2002, Proceedings*, volume 2501 of *Lecture Notes in Computer Science*, pages 267–287. Springer, 2002. URL: https://iacr.org/archive/asiacrypt2002/25010267/25010267.pdf.

[9] Alain Couvreur, Rocco Mora, and Jean-Pierre Tillich. A new approach based on quadratic forms to attack the McEliece cryptosystem. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology—ASIACRYPT 2023—29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4—8, 2023, Proceedings, Part IV*, volume 14441 of *Lecture Notes*

in Computer Science, pages 3–38. Springer, 2023. URL: https://eprint.iacr.org/2023/950.

[10] Alain Couvreur, Ayoub Otmani, and Jean-Pierre Tillich. Polynomial time attack on wild McEliece over quadratic extensions. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology—EUROCRYPT 2014—33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11–15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 17–39. Springer, 2014. URL: https://eprint.iacr.org/2014/112.

[11] Freja Elbro and Christian Majenz. An algebraic attack against McEliece-like cryptosystems based on BCH codes. In *IEEE Information Theory Workshop, ITW 2023, Saint-Malo, France, April 23–28, 2023*, pages 70–75. IEEE, 2023. URL: https://eprint.iacr.org/2022/1715.

[12] Jean-Charles Faugère, Ayoub Otmani, Ludovic Perret, Frédéric de Portzamparc, and Jean-Pierre Tillich. Structural cryptanalysis of McEliece schemes with compact keys. *Designs, Codes and Cryptography*, 79(1):87–112, 2016. URL: https://eprint.iacr.org/2014/210.

[13] Jean-Charles Faugère, Ayoub Otmani, Ludovic Perret, and Jean-Pierre Tillich. A distinguisher for high-rate McEliece cryptosystems, 2010. URL: https://eprint.iacr.org/archive/2010/331/1275979644.pdf.

[14] Chris Hall, Ian Goldberg, and Bruce Schneier. Reaction attacks against several public-key cryptosystems. In Vijay Varadharajan and Yi Mu, editors, *Information and Communication Security, Second International Conference, ICICS'99, Sydney, Australia, November 9–11, 1999, Proceedings*, volume 1726 of *Lecture Notes in Computer Science*, pages 2–12. Springer, 1999. URL: https://cypherpunks.ca/~iang/pubs/paper-reaction-attacks.pdf.

[15] Hugues Randriambololona. Re: Structural analysis of McEliece asymptotically better than generic decoding, 2024. URL: https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/Qj20WzX4dhE/m/mNTpSSSqBgAJ.

[16] Hugues Randriambololona. The syzygy distinguisher, 2024. URL: https://eprint.iacr.org/archive/2024/1193/1722424045.pdf.

[17] Mohamed Ahmed Saeed. *Algebraic approach for code equivalence. (Approche algébrique sur l'équivalence de codes)*. PhD thesis, Normandy University, France, 2017. URL: https://tel.archives-ouvertes.fr/tel-01678829.

[18] Eric R. Verheul, Jeroen M. Doumen, and Henk C. A. van Tilborg. Sloppy Alice attacks! Adaptive chosen ciphertext attacks on the McEliece public-key cryptosystem. In Mario Blaum, Patrick G. Farrell, and Henk C. A. van Tilborg, editors, *Information, coding and mathematics*, volume 687 of *Kluwer International Series in Engineering and Computer Science*, pages 99–119. Kluwer, 2002.