Appeal under Section 515 of Public Law 106-554

Daniel J. Bernstein, 2025-08-09

1 Executive summary

Table 5 of NIST IR 8545—a table placed prominently on page 6 of a 34-page NIST report, providing the report's only support for various statements in the text—is ludicrously inaccurate. The table purports to report speeds of an algorithm submitted in 2022 to NIST, namely round-4 Classic McEliece; but NIST actually substituted speeds of a third-party algorithm, much worse speeds than round-4 Classic McEliece.

Figure 1 is a copy of the speed information (in cycles) from the submitted documentation. It is straightforward to re-measure the submitted algorithm and to verify that the speeds are in line with Figure 1. Figure 2 here is a copy of NIST's falsified table (in kcycles). For, e.g., mceliece348864f keygen (aka keypair), notice the gigantic gap between 114189 kcycles in Figure 2 and 35977 median kcycles with low variability in Figure 1.

NIST IR 8545 was issued in March 2025 (https://web.archive.org/web/20250313083030/https://nvlpubs.nist.gov/nistpubs/ir/2025/NIST.IR.8545.pdf). I pointed out in March 2025 that the table was wrong. The NIST office responsible for this error should have promptly admitted error and issued an update of NIST IR 8545 to fix the table. Instead the office has been stonewalling, turning this into a story of NIST intentionally providing falsified data to the public.

This document is an appeal to the NIST Associate Director for Laboratory Programs. The statutory and regulatory grounds for this appeal are described later in this document. What NIST should do at this point is (1) update NIST IR 8545 to replace Table 5 with correct information, (2) apologize for the fabrications by NIST that produced the original table, and (3) apologize for NIST's series of failures to properly handle my requests for correction.

2 Background: Post-quantum cryptography

Billions of people rely on cryptography to protect data in transit through the Internet. However, an algorithm published by Peter Shor in the 1990s will break many widely deployed cryptographic standards—including various NIST standards—as soon as the attacker has a large quantum computer.

Public quantum computers are not yet large enough to run Shor's algorithm; see, e.g., the survey in https://sam-jaques.appspot.com/quantum_landscape_2024. However, attackers are recording encrypted data to attack later; see https://www.forbes.com/sites/andygreenberg/2013/06/20/leaked-nsa-doc-says-it-can-collect-and-keep-your-encrypted-data-as-long-as-it-takes-to-crack-it/. So these standards are already failing *today* to provide long-term confidentiality. Users often want long-term confidentiality, and confidentiality rights are legally protected for, e.g., medical records and lawyer-client communication, but various current NIST standards are failing to stop attacks.

The phrase "post-quantum cryptography" refers to cryptography under the assumption that the attacker has a quantum computer. I coined the phrase in 2003, and a large part of my work is on this topic.

3 Classic McEliece

Many proposals for post-quantum cryptography have been shown to not achieve their security goals. My paper https://cr.yp.to/papers.html#qrcsp includes a survey of many examples.

	operation	quartile	median	average	quartile
mceliece348864	keypair	35039714	56705880	60333686	67615011
mceliece348864f	keypair	35970884	35976620	35978769	35981416
mceliece460896	keypair	116209838	153266214	213425513	264539700
mceliece460896f	keypair	117267744	117297677	117301747	117331130
mceliece6688128	keypair	265554240	443746986	479441242	532990499
mceliece6688128f	keypair	274329761	274384229	274484625	274430338
mceliece6960119	keypair	241288202	316995472	432622560	468394597
mceliece6960119f	keypair	240198020	240226771	240328382	240254131
mceliece8192128	keypair	308008713	486195290	548932457	664466919
mceliece8192128f	keypair	306203040	306238935	306349035	306280509
mceliece348864	enc	34951	36457	37585	38980
mceliece460896	enc	69674	76086	81312	88956
mceliece6688128	enc	165296	171442	175788	185077
mceliece6960119	enc	139980	144678	147192	149592
mceliece8192128	enc	155174	156945	158068	159040
mceliece348864	dec	127036	127140	127668	127256
mceliece460896	dec	262919	263046	263634	263225
mceliece6688128	dec	305910	306212	306946	306925
mceliece6960119	dec	286353	286596	287218	287038
mceliece8192128	dec	309938	310097	310773	310475

Table 2: Time for complete cryptographic functions on an Intel Haswell CPU core. All times are expressed in CPU cycles. Statistics are computed across SUPERCOP's default 93 experiments. The `f` variants have different `keypair` algorithms but identical `enc` algorithms and identical `dec` algorithms.

For `6960119` and `6960119f`, there is a distinction between Simply Decoded Classic McEliece and Narrowly Decoded Classic McEliece. The official software implements Narrowly Decoded Classic McEliece.

Table 2 reports speeds of the `avx` implementations on an Intel Haswell CPU core. These software measurements were collected using `supercop-20220506` running on a computer named `hiphop`. The CPU on `hiphop` is an Intel Xeon E3-1220 v3 running at 3.10GHz. This CPU does not support hyperthreading. It does support Turbo Boost but `/sys/devices/system/cpu/intel_pstate/no_turbo` was set to 1, disabling Turbo Boost. `hiphop` has 32GB of RAM and runs Ubuntu 18.04. Benchmarks used `./do-part`, which ran on one core of the CPU. The compiler list was reduced to just `gcc -march=native -mtune=native -O3 -fomit-frame-pointer -fwrapv -fPIC -fPIE`.

For comparison, the `mceliece8192128` software originally submitted for round 1 took about 2 billion cycles for each key-generation attempt (and on average about 6 billion cycles for total key generation), slightly under 300000 cycles for encapsulation, and slightly over 450000 cycles for decapsulation.

10

Figure 1: Excerpt from the round-4 Classic McEliece submission to NIST, specifically from the submitted speed documentation (https://classic.mceliece.org/mceliece-impl-20221023.pdf), which accurately reports easily reproducible speed measurements.

parameter set from the corresponding submission. The "Level" columns indicate the security categories that the submission parameter sets claim to meet. BIKE and HQC each had one parameter set per security category, while Classic McEliece had two. The Classic McEliece f versions have faster key generation, while the non-f versions have simpler key generation.

In these benchmarks, BIKE is 6-10 times slower than HQC in key generation, 5-7 times slower than HQC in decapsulation, and about twice as fast as HQC in encapsulation. Key generation in Classic McEliece is an outlier, being three orders of magnitude more costly than HQC.

Parameter Set	Level	keygen	encaps	decaps
BIKE Level 1	I	637	111	1 428
BIKE Level 3	III	1 892	251	4 313
BIKE Level 5	V	4 535	505	10 382

Table 3. Performance of BIKE in thousands of cycles on x86_64 [1]

Parameter Set	Level	keygen	encaps	decaps
hqc-128	I	105	197	360
hqc-192	III	244	460	746
hqc-256	V	4246	844	1 410

Table 4. Performance of HQC in thousands of cycles on x86_64 [1]

Parameter Set	Level	keygen	encaps	decaps
mceliece348864	I	137 345	49	120
mceliece348864f		114 189	45	120
mceliece460896	III	430 364	91	232
mceliece460896f		313 600	92	231
mceliece6688128	V	674 012	196	273
mceliece6688128f		493 758	176	274
mceliece6960119	V	602 164	167	252
mceliece6960119f		404 166	169	253
mceliece8192128	V	686 110	203	269
mceliece8192128f		453 985	206	269

Table 5. Performance of Classic McEliece in thousands of cycles on x86_64 [1]

Figure 2: NIST's fake claims in NIST IR 8545 regarding the speed of the round-4 Classic McEliece submission. The numbers are in fact measurements of much slower third-party software called "liboqs". NIST falsely portrayed these as measurements of the submission, and has stonewalled in response to correction requests.

Some proposals are less risky. In particular, a cryptosystem introduced by McEliece in 1978—one of the earliest public-key cryptosystems—has maintained practically the same pre-quantum strength for half a century despite a long list of papers trying to find any weaknesses in it. After Shor's algorithm appeared, researchers also began searching for post-quantum weaknesses in McEliece's cryptosystem, and again the system has held up very well. Meanwhile newer options have much worse track records. See https://classic.mceliece.org/comparison.html and https://classic.mceliece.org/papers.html.

I am speaking for myself throughout this document, but I am also one member of a large team (https://classic.mceliece.org/people.html) that developed "Classic McEliece". The name refers to a guarantee that an important security property of Classic McEliece is as difficult to break as McEliece's original 1978 cryptosystem. This does not mean that Classic McEliece is identical to the original McEliece proposal: Classic McEliece incorporates various improvements that are guaranteed to not damage security. This includes a series of cost improvements.

As https://mceliece.org shows, there are already many users of Classic McEliece, and many further public recommendations of Classic McEliece. The basic reasons for interest in Classic McEliece are well documented:

- the security track record of the McEliece system, and
- the fact that Classic McEliece is the most efficient post-quantum option for a large class of cryptographic applications (called "static keys" in, e.g., NIST SP 800-56 and NIST SP 800-57).

See the quotes and links in https://blog.cr.yp.to/20250423-mceliece.html.

There is, however, also a similarly basic reason for interest in other options: namely, there is another large class of cryptographic applications for which other options outperform Classic McEliece.[1] If the resources available for cryptography inside an application do not suffice for Classic McEliece (this is a quantitative question about the application resources and about the costs of Classic McEliece), while the same resources suffice for a higher-risk option, then the application cannot use Classic McEliece and should use the higher-risk option—rather than doing nothing; doing nothing would be failing to protect user data against attack.

4 NIST's call for submissions of post-quantum algorithms

In 2016, NIST issued 81 FR 92787 (https://www.federalregister.gov/documents/2016/12/20/2016-30615/announcing-request-for-nominations-for-public-key-post-quantum-cryptographic-algorithms) soliciting proposals of post-quantum replacements for various NIST standards:

> This notice solicits nominations from any interested party for candidate algorithms to be considered for public-key post-quantum standards. The submission requirements and the minimum acceptability requirements of a ''complete and proper'' candidate algorithm submission, as well as the evaluation criteria that will be used to appraise the candidate algorithms, can be found at http://www.nist.gov/pqcrypto.

https://web.archive.org/web/20220119113311/https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf is a copy of NIST's statement of the requirements and evaluation criteria. Each submission package was required to "describe a collection of algorithms, also called a cryptosystem or cryptographic scheme"; to present further documentation, including "a statement regarding the algorithm's estimated computational efficiency and memory requirements"; and to include software, where in particular the "optimized implementation, targeting the Intel x64 processor (a 64-bit implementation), is intended to demonstrate the performance of the algorithm".

[1] Those options have larger "ciphertexts" but smaller "keys". This increases total costs for "static keys" used for many ciphertexts, but decreases total costs for, e.g., "one-time keys" used for just one ciphertext per key.

NIST listed evaluation criteria under three main headings: "Security", "Cost", and "Algorithm and implementation characteristics". Security was listed as "the most important factor", but "Cost" occupied a full page of criteria and obviously was not ignored.

The requirements said that "Submission packages must be received by NIST by November 30, 2017" but also said that "NIST expects to perform multiple rounds of evaluation, over a period of three to five years" with a possibility of modified submissions in subsequent rounds, including new software ("new reference and optimized implementations"). Ultimately NIST collected four rounds of submissions with due dates in 2017 (round 1), 2019 (round 2), 2020 (round 3), and 2022 (round 4), plus further "on-ramp" submissions not relevant here.

5 NIST's non-transparency

NIST claimed that it would carry out transparent evaluations:

> NIST will perform a thorough analysis of the submitted algorithms in a manner that is open and transparent to the public, as well as encourage the cryptographic community to also conduct analyses and evaluation. This combined analysis will inform NIST's decision on the subsequent development of post-quantum standards.

Unfortunately, NIST actually kept a wide range of project documents secret, often marking the documents "Not for public distribution". The reason this is now known is that I have two ongoing FOIA lawsuits that have been forcing NIST to disclose documents regarding this project. The court-monitored disclosure process is slow, but reviewing the results so far from those lawsuits (https://nist.pqcrypto.org/foia/index.html) already shows that NIST drastically limited the information that it released to the public.

Furthermore, for the information it *did* release, NIST delayed almost all releases until end-of-round reports: NIST IR 8240 (2019), NIST IR 8309 (2020), NIST IR 8413 (2022), and NIST IR 8545 (2025). Occasionally NIST released information earlier, but those reports were again and again stating NIST decisions based on NIST claims that the public had not had an opportunity to see or correct.

This pattern of secrecy does not remove the public interest in correcting NIST's errors once the errors *are* visible. Post-quantum cryptography is attracting more and more attention, and many parties are making further post-quantum decisions. For example, NIST IR 8545 includes the following statement: "After the ISO standardization process has been completed, NIST may consider developing a standard for Classic McEliece based on the ISO standard." What stops misinformation in NIST IR 8545 from influencing decisions about Classic McEliece by ISO, NIST, and other parties? The misinformation should be promptly corrected.

6 Performance-measurement failures

There are many sources of performance measurements. Sometimes measurements are botched. See, e.g., the survey of "benchmarking crimes" in https://arxiv.org/abs/1801.02381.

But surely the NIST Associate Director for Laboratory Programs will *presume* that NIST employees handling performance measurements are doing so competently. After all, NIST is the United States government's official "measurement laboratory", with a mission to "promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life".

It therefore seems necessary to point out that NIST's Cryptographic Technology Group is not entitled to any such presumption. The group's hiring criteria and stated goals are about topics other than performance

measurement: see, e.g., https://web.archive.org/web/20250729034746/https://csrc.nist.gov/Groups/Computer-Security-Division/Cryptographic-Technology. The group has repeatedly failed at basic measurement tasks.

6.1. Example: failing to report measurement variability. Consider again Figure 2 showing measurement tables from NIST IR 8545. Each measurement is presented as just *one number*. There's nothing measuring variability (for example, standard deviations or quartiles). This is an example of Benchmarking Crime B4: "No indication of significance of data". The benchmarking-crimes paper cited above explicitly lists "the lack of some indication of variance, such as a standard deviation or significance test" as a concrete example of Benchmarking Crime B4.

For comparison, one statement from https://web.archive.org/web/20250507112244/https://www.nist.gov/director/nist-information-quality-standards is the following: "NIST Policy 1800.00 requires that NIST measurement results be accompanied by quantitative statements of their uncertainty, and that a uniform approach to expressing measurement uncertainty be followed."

6.2. Example: conflating different CPUs. As another—less blatant but still indefensible—example, recall from Section 4 that NIST's call for submissions specified "the Intel x64 processor" as a performance target. This is an example of Benchmarking Crime F1: "Missing platform specification".

There are many different Intel x64 processors running at different speeds. The most obvious reason for speed variations is clock frequency: for example, a 2GHz CPU core carries out 2 billion cycles per second, while a 3GHz CPU core carries out 3 billion cycles per second. One can compensate for this particular variation by measuring cycle counts rather than seconds—but that is not enough: there are further variations, and often those variations are important.

Here is a concrete example. The "Kyber" submission, one of the submissions that NIST later selected for standardization (after revisions), provided "ref" software and much faster "avx2" software. Recall that NIST had asked for "reference and optimized implementations", with the optimized implementations being used "to demonstrate the performance of the algorithm". But Kyber's "avx2" software relies on features that Intel introduced with its Haswell microarchitecture in 2013. Older Intel x64 CPUs (and some newer but smaller Intel x64 CPUs such as Tremont CPUs) had no choice but to use the "ref" software, producing much higher cycle counts. Meanwhile various newer Intel CPUs use somewhat fewer cycles for the "avx2" software than Intel Haswell CPUs do.

So there are actually many different cycle counts for Kyber on different Intel x64 processors. Conflating the processors into "the" Intel x64 processor creates severe ambiguities in performance claims, complicates efforts to check those claims, and complicates comparisons—for example, should the Kyber submission have been penalized for relying on Haswell features?

In 2019, NIST addressed the ambiguity of "the Intel x64 processor" by designating Haswell as its Intel comparison platform. See https://web.archive.org/web/20220120140641/https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/BjLtcwXALbA/m/Bjj_77pzCAAJ, where NIST asked submission teams to provide "an AVX2 (Haswell) optimized implementation". But this fix took a long time and, as we'll see in Section 8.1, NIST has now reintroduced and weaponized the same ambiguity.

7 NIST's falsified cost data regarding round-4 Classic McEliece

NIST IR 8545 ("Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process") claims to be a report on the round-4 submissions ("This report describes the evaluation and selection process of these fourth-round candidates based on public feedback and internal review"), the submissions filed in 2022, including the round-4 Classic McEliece submission.

More specifically, NIST IR 8545 says "Tables 3 through 5 show representative benchmarks for key generations, encapsulations, and decapsulations of BIKE, HQC, and Classic McEliece, respectively. Each row is a specific

parameter set from the corresponding submission." Table 5 similarly claims to report the "Performance of Classic McEliece in thousands of cycles on x86_64" for three operations: "keygen", "encaps", and "decaps".

But that's not true. **The table consists of data falsified by NIST**, making the round-4 Classic McEliece submission sound more expensive than it actually is. I'll give an example in Section 7.1 below.

The correct data was, as required, provided as part of the submission. Software achieving that speed was also provided as part of the submission. Unfortunately, instead of checking and reporting the numbers from the submission, NIST substituted wrong numbers.

7.1. The mceliece348864f example. I'm using mceliece348864f as a running example in this document. For this example, https://classic.mceliece.org/mceliece-impl-20221023.pdf#page.10 (the speed table in the round-4 submission) reported median Haswell cycle counts of (e.g.) 35976620 for keygen, 36457 for enc, and 127140 for dec. The table also reported quartiles to measure variability, and the text accompanying the table explained in detail how to reproduce the measurements. See Figure 1 above.

NIST IR 8545 claims 114189 kcycles (instead of 35977 kcycles) for keygen, 45 kcycles (instead of 36 kcycles) for enc, and 120 kcycles (instead of 127 kcycles) for dec. See Figure 2 above. The keygen number is absurdly inaccurate, more than 3 times larger than the correct number.

8 How did this happen? Part 1

I eventually figured out *why* NIST IR 8545's numbers have such large deviations from reality. Section 8, Section 9, and Section 10 of this document explain the analysis and how to verify the final conclusions.

Figuring this out was unnecessarily difficult because of NIST's violations of (1) normal scientific procedures, (2) NIST's information-quality standards, and (3) the rules for this NIST project. In particular, if NIST's "analysis" had been carried out "in a manner that is open and transparent to the public" as NIST had promised (see Section 5), then there would have already been public records showing how and why NIST (incorrectly) threw away the (correct) cost evaluation from the submission. But NIST did not provide this transparency, and in particular did not publish these records. So I found myself in March 2025 spending time on detective work.

8.1. The CPU question. Recall from Section 6 that NIST had originally named "the Intel x64 processor" as a performance target, as if there were just one such processor; in fact, there are many, with different performance characteristics. NIST had fixed this in 2019 by specifying Haswell as a performance target.

But the "x86_64" in the table caption in NIST IR 8545 was regressing to NIST's original underspecification ("x86_64" and "x64" are synonyms in the CPU context). Could this unspecified "x86_64" be something much slower than a Haswell? I had seen nothing from NIST requesting public feedback on a proposal to switch to a different Intel target, but perhaps NIST was switching without a prior announcement.

The caption of Table 5 in NIST IR 8545 ends "[1]". The entry "[1]" in the bibliography is "Open quantum safe (OQS) algorithm performance visualizations. Available at https://openquantumsafe.org/benchmarking." This is the complete bibliography entry—in particular, notice that NIST failed to state the date when NIST accessed that web page.

Looking around that web site and various versions of the web site archived on https://archive.org, I found statements that (at every potentially relevant moment) the web site was providing numbers for an "Intel(R) Xeon(R) Platinum 8259CL CPU". That CPU has a Cascade Lake architecture, a *newer* Intel microarchitecture having generally *smaller* cycle counts than Haswell.

Quantitatively, Cascade Lake cores are almost identical to cores of Cascade Lake's predecessor Skylake (see https://en.wikichip.org/wiki/intel/microarchitectures/cascade_lake); Skylake, in turn, is two generations newer than Haswell. The following performance numbers for mceliece348864f on Skylake appear

in https://bench.cr.yp.to/results-kem/amd64-samba.html (along with quartiles within a few percent): 31314 kcycles for keygen, 30 kcycles for enc, and 118 kcycles for dec. For comparison, NIST IR 8545 claims 114189 kcycles for keygen, 45 kcycles for enc, and 120 kcycles for dec.

So, yes, NIST did deviate from its announced comparison platform, but in a way that makes NIST's claim of 114189 kcycles (3.65 times larger than 31314 kcycles) even harder to explain.

It is procedurally improper for NIST to have switched away from its announced Haswell comparison target, to have never announced the switch, and to have not given submitters any opportunity to accommodate the switch. But what's *really* disturbing about this switch is that NIST is now waving at unspecified, unquantified platform variations as supposedly justifying Table 5 of NIST IR 8545. In other words, not only has NIST reintroduced the ambiguity of its original reference to "the Intel x64 processor", but NIST is actively hiding behind this ambiguity, putting extra burdens upon people who want to check NIST's table. Content-wise, NIST's platform switch does *not* rescue NIST's table—again, it moves slightly in the *opposite* direction—but seeing this is extra work that would never have been necessary if NIST had stuck to Haswell.

8.2. Obsolete data. Meanwhile the benchmarking link that NIST had cited gave what looked like a much more straightforward explanation for NIST's wrong numbers: namely, the top of the linked page showed a prominent warning that the site was giving out-of-date numbers from a defunct measurement project. This warning was already on the page in April 2024 (https://web.archive.org/web/20240425050637/https://openquantumsafe.org/benchmarking/), almost a year before NIST IR 8545 was issued:

> **Disclaimer**
>
> These pages visualize measurements taken by the **now-defunct** OQS profiling project. This project is not currently maintained, and these measurements are not up to date.

(Boldface in original.)

For Classic McEliece, the round-2 submission in 2019 added, *inter alia*, mceliece348864f and much faster keygen than in round 1; see https://classic.mceliece.org/nist/mceliece-20190331.pdf#page.21. The round-3 submission in 2020 included, *inter alia*, much faster keygen than in round 2; see https://classic.mceliece.org/nist/mceliece-20201010.pdf#page.33. Less importantly, the round-4 submission added, *inter alia*, a tweak to mceliece348864f that marginally improved speeds; accordingly, marginally better numbers were provided later in 2022 to supplement the numbers shown in Figure 1.

At this point Occam's razor says that NIST took round-2 Classic McEliece speeds where it should instead have taken round-4 Classic McEliece speeds. Such sloppiness is unacceptable but at least understandable.

But Occam's razor turned out to be wrong in this case: what actually happened was more complicated. Astonishingly, NIST's numbers didn't come from *any* version of the submission. Instead **NIST took measurements of a slower third-party algorithm, and falsely reported those as measurements of the round-4 submission**. See Section 10 below.

9 How did this happen? Part 2

I had sent a message to NIST's public pqc-forum mailing list dated 21 Feb 2025 11:34:59 +0100, and a followup message dated 27 Mar 2025 19:48:13 +0100, regarding various benchmarking aspects of NIST's ongoing cryptographic standardization work. NIST IR 8545 was released in the interim, and my followup message (https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/hRTPcxHohPM/m/BOpE15i_BQAJ) commented that "some serious fixes seem necessary to the procedures that NIST uses for handling benchmark numbers".

In particular, my message looked at NIST's claim of 114189 kcycles for mceliece348864f keygen, went through the same initial detective work as in Section 8 above, and ended as follows:

> It's astonishing to see NIST issuing a report in 2025 with benchmarks of code that's six years out of date, and presenting those as benchmarks of the round-4 submission, especially when the source that NIST cites is a page that says at the top that it's presenting obsolete measurements from a defunct benchmarking project.
>
> If NIST had been carrying out its selection discussions on a public mailing list, then this 4x error could have been corrected as soon as it appeared. But FOIA results show that NIST's selection discussions are mostly carried out in secret. There's no evident way that on-ramp teams will be able to correct similar benchmarking errors.

Unfortunately, NIST didn't respond to anything I had written about NIST IR 8545.

I sent a much shorter message dated 25 Apr 2025 10:30:17 +0200 saying "the report's 114 189 cost figure is inflated by 4x compared to reality" and requesting an erratum.

NIST responded a week later (https://web.archive.org/web/20250507151442/https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/CSYRn8A-0zE/m/1PPIazqcAAAJ) with the following stonewalling:

> Thank you for your feedback on our Round 4 report, NIST IR 8545. As noted in the document, the benchmark figures presented in the tables are intended to be representative,

Representative of *what*?

NIST IR 8545 says it's presenting evaluations of the round-4 submissions. However, for Classic McEliece, the numbers in Table 5 of NIST IR 8545 are for much slower software than the round-4 submission. This is **data falsification by NIST**. NIST is obliged to correct the resulting misinformation.

> and NIST did not rely solely on those listed numbers when evaluating performance.

Even if other information is available, even if other information was used by NIST, NIST remains responsible for correcting the wrong information that NIST presented in NIST IR 8545.

> We recognize that performance can vary significantly depending on factors such as platform, application, computing environment, and network conditions.

This is a smokescreen: NIST deterring review by making the situation at hand sound more complicated than it actually is.

NIST is falsely claiming (e.g.) 114189 kcycles for round-4 mceliece348864f keygen on "x86_64". The ambiguity of "x86_64" is part of the same smokescreen, but investigation shows that this number is from an Intel Xeon Platinum 8259CL CPU; see Section 8. On that CPU, the submitted round-4 software is 4× faster than NIST's claim. The network conditions and application have no effect on these numbers.

> We wanted to keep the report short, so we did not cite every source that we reviewed.

This doesn't remove NIST's responsibility to correct the falsified data presented in NIST IR 8545.

> The performance data on SUPERCOP has been very useful to us over the years.

Here NIST is referring to a large cryptographic-software-testing project (https://bench.cr.yp.to) that I co-initiated many years ago and co-manage. There have been more than 100 releases of the complete test package, which is called SUPERCOP. The latest release contains 4601 "implementations" (submitted by hundreds of people) of 1430 cryptographic "primitives". The web pages show results for this release—in particular, cost measurements of the fastest "implementation" of each "primitive"—from 45 different computers, including 3 computers with Intel Haswell CPUs. Anyone with the same type of computer can re-run the tests to check the full results or to spot-check selected results. As of 8 August 2025, Google Scholar reports 398 publications that cite this project using its recommended citation (see https://scholar.google.com/scholar?cites=16138629620498175528&as_sdt=400005&sciodt=0,14&hl=en), and 711 publications mentioning https://bench.cr.yp.to (see https://scholar.google.com/scholar?hl=en&as_sdt=0%2C14&q=%22bench.cr.yp.to%22&btnG=).

Each version of the official Classic McEliece software was promptly integrated into SUPERCOP. In particular, the round-4 Classic McEliece software was integrated into SUPERCOP in October 2022. Simply looking at the mceliece348864f numbers on the web pages after that (or re-running SUPERCOP to reproduce the numbers) is another way that NIST could and should have caught the massive errors in Table 5 of NIST IR 8545.

Unfortunately, the massive errors are there. NIST now has a responsibility to fix those errors.

> We have verified that the performance figures given in the tables in our report accurately reflect the benchmarks we cited.

NIST provided no details to justify this claim.

In response (email dated 1 May 2025 17:37:37 +0200), I pointed out again that NIST's source already had a prominent disclaimer at the top in April 2024 saying that it was presenting out-of-date numbers from a defunct profiling project. I asked why NIST was taking numbers from this source in a March 2025 report and suppressing any mention of the disclaimer. I ended as follows: "It is, in any event, easy to test the fourth-round software and see that the NIST IR 8545 numbers are highly inaccurate. This is now my third time pointing out this particular error in NIST IR 8545; it's amazing that the error still hasn't been fixed."

NIST did not reply.

10 How did this happen? Part 3

A week later, I sent email (dated 7 May 2025 17:39:21 +0200) to info.quality@nist.gov with "a request under Section 515 of Public Law 106-554 for correction of information". The law forces NIST to respond to such requests.

NIST's reply took three months (it was sent by email dated 4 Aug 2025 20:01:45 +0000) and was mostly stonewalling. But there was a bit of progress, because NIST did specifically address the round-2-vs.-round-4 point: "NIST disagrees with this assertion. NIST is confident that the benchmarks cited from Reference [1] accurately reflect the fourth-round version of Classic McEliece."

The numbers in NIST IR 8545 remain wildly inaccurate. NIST's declaration of confidence didn't magically make the numbers correct. But the declaration made me guess that NIST's underlying mistake wasn't as simple as conflating the round-2 and round-4 speeds.

So I looked more closely at the web site that NIST had cited. I ended up tracking down exactly how **NIST's numbers actually came from a slower third-party algorithm**, an algorithm that NIST was and is falsely claiming to be the round-4 Classic McEliece submission.

Here are the details. Let me again emphasize that none of this detective work should have been necessary.

10.1. The source cited by NIST has reported many different "Classic McEliece" numbers from different "OQS" versions and different "Raw Data" versions. Clicking on the "KEM algorithms runtime" link from https://web.archive.org/web/20240425050637/https://openquantumsafe.org/benchmarking/ finds https://web.archive.org/web/20240318011921/https://openquantumsafe.org/benchmarking/visualization/speed_kem.html, which includes a line of cycle counts for "Classic-McEliece-348864f (x86_64)". Clicking to earlier and later archived versions of the same page finds different cycle counts for "Classic-McEliece-348864f (x86_64)"—often radically different; for example, pages saying "OQS version 0.7.3-dev" or "OQS version 0.8.0-dev" report roughly 280000 kcycles for "Classic-McEliece-348864f (x86_64)", while pages saying "OQS version 0.9.0-rc1" report roughly 140000 kcycles for "Classic-McEliece-348864f (x86_64)". There are also smaller variations in pages showing different "Raw Data" dates.

Almost every archived page has keygen cycle counts that, when rounded to kcycles, *don't* exactly match the 114189 claimed in NIST IR 8545. However, the last data before changes to the web page in July 2025 (and therefore the last data before NIST IR 8545 was issued), namely "Raw Data 2024-04-02" with "OQS version 0.9.0-rc1", says 114188667 keygen cycles (see https://web.archive.org/web/20250120141536/https://openquantumsafe.org/benchmarking/visualization/2024-04-02/speed_kem.json), which matches NIST's claimed 114189 kcycles. It also says 44812 enc kcycles, which matches NIST's claimed 45 enc kcycles, and 120425 dec cycles, which matches NIST's claimed 120 dec kcycles.

(The same data set shows "BIKE-L1" as 636580 keygen cycles, matching 637 kcycles in Figure 2; 110537 enc cycles, matching 111 kcycles; and 1750669 dec cycles, *not* matching 1428 kcycles. The data set also translates the 1750669 dec cycles to 1428 dec per second. Evidently NIST copied this dec-per-second number to a kcycles number. Another mistake that was pointed out before and that NIST *did* promptly fix is that the data set shows 446071 cycles for HQC-256 keygen while Figure 2 claimed 4246 kcycles.)

This pinpoints when NIST obtained its claimed numbers for NIST IR 8545, namely 2024-04-02 or later. This also shows that NIST was obtaining numbers from "OQS version 0.9.0-rc1" rather than "OQS version 0.8.0-dev" or earlier. The obvious next question is how these "OQS" versions were related to different versions of the Classic McEliece submission, and in particular the round-4 Classic McEliece submission. Remember that NIST wasn't presenting these as numbers for "OQS version 0.9.0-rc1"; it was presenting these as numbers for the "fourth-round candidates".

10.2. Every version of the "OQS" software differs from the round-4 Classic McEliece software. The web page seemed to indicate that "OQS" referred to measurements of a library called "liboqs". https://cr.yp.to/2025/20250807-liboqs.tar.gz is a copy of the "liboqs" repository at the time of this writing (officially hosted on GitHub, but sometimes repositories disappear from GitHub, so I'm more comfortable linking to a copy saved on my own https://cr.yp.to site). This repository includes the current version of "liboqs" and previous versions such as 0.9.0-rc1.

To check that a measurement of "Classic McEliece" in some "OQS" version is measuring the round-4 Classic McEliece software, NIST would have had to check that this software in "liboqs" is the round-4 Classic McEliece software. So, focusing again on mceliece348864f, I compared the "`pqclean_mceliece348864f_avx2`" software in "liboqs" to the authentic software in the round-4 Classic McEliece submission, including a scan for different versions of this software in different versions of "liboqs". My comparison immediately showed that, no, this software *isn't* the round-4 Classic McEliece software.

There is no reason *a priori* to believe that the "liboqs" speeds should match the speed of the round-4 Classic McEliece submission. The speeds are in fact radically different; see Section 10.3 below. See also Section 10.4 pinpointing the main reason that "liboqs" is so slow, and explaining how to verify this.

I also see nowhere that "liboqs" makes any claim of a match. As far as I can tell, that claim is a fabrication by NIST. (See Section 10.6 below.) But, even if "liboqs" did make such a claim somewhere, NIST had an independent responsibility to *check* that claim, and failing to check was reckless. Either way, NIST has a responsibility for correcting the false information that NIST has distributed regarding round-4 Classic McEliece.

10.3. The "OQS" software is much slower than the round-4 Classic McEliece software. I checked the cost of "liboqs" 0.9.0-rc1 on an Intel Xeon E3-1220 v5 (Skylake) under Ubuntu 22.04 using gcc 11.4. Here are commands to rerun this experiment:

```
wget https://cr.yp.to/2025/20250807-liboqs.tar.gz
tar -xf 20250807-liboqs.tar.gz
cd liboqs
git checkout 0.9.0-rc1
mkdir build
cd build
cmake -GNinja ..
ninja
tests/speed_kem Classic-McEliece-348864f
```

The cycle counts reported by "liboqs" were mean 126297630 (deviation 848857) for keygen, mean 69420 (deviation 21258) for enc, mean 147988 (deviation 1115) for dec. (A similar experiment with newer 0.14.0-rc1 produced essentially the same numbers.)

This is the same computer that produced the measurements in https://bench.cr.yp.to/results-kem/amd64-samba.html mentioned above: 31314 kcycles for keygen, 30 kcycles for enc, and 118 kcycles for dec.

One could also try further computers to see the "liboqs" slowness on each computer. I instead looked at how the "liboqs" software differs from the round-4 Classic McEliece software, and spotted a difference that was *obviously* a big slowdown in "liboqs", as explained below.

10.4. The "OQS" software uses a much slower integer-sorting algorithm than the round-4 Classic McEliece software. The round-4 Classic McEliece software includes an `int32_sort` subroutine that quickly sorts an array of 32-bit integers, whereas "liboqs" substitutes a much slower (in particular, non-vectorized) sorting algorithm.

I copied the authentic `int32_sort` subroutine from the round-4 Classic McEliece submission into "liboqs". Here are commands to rerun this experiment:

```
wget https://cr.yp.to/2025/20250807-liboqs.tar.gz
tar -xf 20250807-liboqs.tar.gz
cd liboqs
git checkout 0.9.0-rc1
wget https://cr.yp.to/2025/20250805-oqs-348864f-patch.txt
patch -p1 < 20250805-oqs-348864f-patch.txt
mkdir build
cd build
cmake -GNinja ..
ninja
tests/speed_kem Classic-McEliece-348864f
```

This modification of "liboqs" reported around 40 Skylake Mcycles for keygen. That's still behind the authentic software (sorting isn't the only thing that "liboqs" slowed down), but it's not in the same ballpark as the insanely wrong numbers from NIST IR 8545.

10.5. NIST falsified Table 5. To be clear, I'm not saying "liboqs" did something wrong here. The "liboqs" documentation states many goals other than speed. Presumably those goals explain the differences between the round-4 Classic McEliece software and "liboqs", including the "liboqs" substitution of a much slower sorting algorithm.

Meanwhile NIST's behavior was *definitely* wrong. NIST IR 8545 claims to report measurements of the round-4 Classic McEliece submission; in fact, NIST took measurements of a (much slower) third-party algorithm, and misrepresented those as measurements of the round-4 Classic McEliece submission.

The file `README.md` in "liboqs" says "WE DO NOT CURRENTLY RECOMMEND RELYING ON THIS LIBRARY IN A PRODUCTION ENVIRONMENT OR TO PROTECT ANY SENSITIVE DATA" (uppercase in original). Basically, "liboqs" is a teaching tool, whereas the round-4 Classic McEliece submission provides real software that is designed for deployment and that is in fact deployed. NIST took measurements of the teaching tool and misrepresented those as measurements of the round-4 submission.

Presumably the culprits at NIST will say that they didn't *know* that they were taking slower algorithms. That's merely the difference between recklessness and dishonesty. Either way, NIST IR 8545 has to be fixed.

NIST could have very easily caught its mistake here by checking what the submission said about speed, or by checking what `https://bench.cr.yp.to` said about speed, or by trying the submitted software. NIST's failure to check its numbers is inexcusable.

It's also inexcusable that NIST failed to provide the transparency that it had promised. Transparency would have allowed the public to catch NIST's mistake *before* that mistake fed into NIST's decision-making process, and presumably the mistake would then have been rapidly corrected without a fuss.

Instead NIST published this misinformation as part of a prominent table in a report that was announcing various final decisions. I promptly pointed out the error; NIST then stonewalled, evidently judging that denial would be better for NIST than admitting an embarrassing error.

10.6. Words matter. There is a "`docs/algorithms`" directory in "liboqs", evidently documentation of the algorithms inside "liboqs". Inside that directory, "`kem/classic_mceliece.md`" has title "Classic McEliece" and various bullet items such as "Algorithm type: Key encapsulation mechanism", "Specification version: SUPERCOP-20221025", "Source: `https://github.com/PQClean/PQClean/commit/1eacfdafc15ddc5d5759d0b85b4cef26627df181`", and "Ancestors of primary source: SUPERCOP-20221025 'clean' and 'avx2' implementations".

The Haswell-optimized round-4 Classic McEliece software for, e.g., mceliece348864f is practically identical to the software in the `crypto_kem/mceliece348864f/avx` directory in SUPERCOP version 20221025; in particular, I've checked that the speeds are the same. But are these bullet items claiming that "liboqs" is using that software? No, they aren't. They are labeling that software as an *ancestor* of the "liboqs" source. In other words, the "liboqs" software is *derived from* that software. This doesn't mean that the "liboqs" software *is* that software. As we've seen, it *isn't* that software; it's much slower, most importantly because it substitutes a much slower sorting algorithm.

Is it possible that NIST somehow thought that "ancestor" meant "identical"? This is the most charitable explanation that I can see at this point for NIST falsely claiming that the measurements of "liboqs" were measurements of the round-4 Classic McEliece submission.

To be clear, this was still astonishingly reckless. The correct information was readily available in the submission (and from SUPERCOP), and was easy to verify. NIST had promised to carry out "a thorough analysis of the submitted algorithms", as one would expect for such an important project. Evidently NIST failed.

11 The law

Beyond NIST's *ethical* obligation to promptly correct the misinformation in NIST IR 8545, NIST has a *legal* obligation to do so under the following statutes and regulations.

11.1. OMB's authority over NIST publications. 44 U.S.C. §3504(d)(1) requires the Director of OMB to "develop and oversee the implementation of policies, principles, standards, and guidelines" that "apply to

Federal agency dissemination of public information, regardless of the form or format in which such information is disseminated". 44 U.S.C. §3516 require the Director to "promulgate rules, regulations, or procedures necessary to exercise the authority provided by this subchapter".

11.2. The obligation to ensure and maximize information quality. Section 515(a) of Public Law 106-554 states the following: "The Director of the Office of Management and Budget shall, by not later than September 30, 2001, and with public and Federal agency involvement, issue guidelines under sections 3504(d)(1) and 3516 of title 44, United States Code, that provide policy and procedural guidance to Federal agencies for ensuring and maximizing the quality, objectivity, utility, and integrity of information (including statistical information) disseminated by Federal agencies in fulfillment of the purposes and provisions of chapter 35 of title 44, United States Code, commonly referred to as the Paperwork Reduction Act."

Section 515(b) requires OMB's guidelines to, in turn, "require that each Federal agency to which the guidelines apply" issue "guidelines ensuring and maximizing the quality, objectivity, utility, and integrity of information (including statistical information) disseminated by the agency, by not later than 1 year after the date of issuance of the guidelines under subsection (a)", and to "establish administrative mechanisms allowing affected persons to seek and obtain correction of information maintained and disseminated by the agency that does not comply with the guidelines issued under subsection (a)".

11.3. OMB's rules. Pursuant to the above statutes, OMB issued rules (https://obamawhitehouse.archives.gov/sites/default/files/omb/fedreg/reproducible2.pdf later supplemented by https://web.archive.org/web/20250422071633/https://trumpwhitehouse.archives.gov/wp-content/uploads/2019/04/M-19-15.pdf) imposing more specific requirements upon each agency, such as requiring each agency to be able to "substantiate the quality of the information it has disseminated through documentation or other means appropriate to the information"; requiring time limits on the handling of correction requests; requiring "a high degree of transparency about data and methods to facilitate the reproducibility of such information by qualified third parties" when "an agency is responsible for disseminating influential scientific, financial, or statistical information"; and, for such information, requiring "sufficient transparency about data and methods that an independent reanalysis could be undertaken by a qualified member of the public".

11.4. DOC's rules. The Department of Commerce, NIST's parent agency, delegated compliance to individual units as NIST (see https://www.federalregister.gov/documents/2002/05/03/02-10991/guidelines-for-ensuring-and-maximizing-the-quality-objectivity-utility-and-integrity-of-disseminated and https://www.federalregister.gov/documents/2002/10/08/02-25340/guidelines-for-ensuring-and-maximizing-the-quality-objectivity-utility-and-integrity-of-disseminated), although it imposed some procedural constraints such as the following: "The operating unit will respond to all initial requests within 60 calendar days of receipt. If the request requires more than 60 calendar days to resolve, the operating unit will inform the complainant that more time is required and indicate the reason why and an estimated decision date."

11.5. NIST's rules. As required, NIST issued its own information-quality standards (https://web.archive.org/web/20250507112244/https://www.nist.gov/director/nist-information-quality-standards). The specifics of these information-quality standards—and of NIST's violations of those information-quality standards—are reviewed in detail below and provide the main legal force for this appeal.

It is, however, important to note that this appeal *also* has legal force provided by the statutes, the OMB rules, and the DOC rules. This is important because NIST's information-quality standards impermissibly deviate in various ways from those statutes and rules. For example, NIST's information-quality standards say that appeals will be handled "usually within 60 calendar days"; because of the word "usually", this does not meet OMB's requirement of a time limit.

As another example, NIST's information-quality standards claim that "Because NIST is a non-regulatory federal agency, it is not anticipated that NIST will disseminate influential scientific information". But the OMB definition of "influential" asks whether "the agency can reasonably determine that dissemination of the information will have or does have a clear and substantial impact on important public policies or important

private sector decisions". This isn't limited to regulations issued by the same agency. NIST isn't allowed to ignore the influence of its post-quantum reports on industry decisions, on White House orders, etc.

12 NIST IR 8545's failure to comply with NIST's information-quality standards

NIST's correction procedures, beyond simply asking "why the requester believes that the information is not correct", also ask "how the information at issue fails to comply with applicable information quality guidelines and standards".

NIST IR 8545—in particular, the table at issue—does not comply with the following NIST information-quality standards. I am not saying that this list is complete, but this list is clearly sufficient for invoking Section 515 of Public Law 106-554.

12.1. Quality assurance. NIST's information-quality standards include the following: "Quality will be ensured and established at levels appropriate to the nature and timeliness of the information to be disseminated." The word "quality" is defined to include "objectivity", which includes "a focus on ensuring accurate, reliable, and unbiased information", along with a procedural requirement of "sound statistical and research methods".

What NIST did in Table 5 of NIST IR 8545 wasn't sound research. It wasn't research at all. Despite NIST's promise to "perform a thorough analysis of the submitted algorithms", NIST built its table by copying cost numbers from another source.

Critically, that source was presenting measurements of "OQS version 0.9.0-rc1", but NIST falsely presented the numbers as measurements of the round-4 Classic McEliece submission. This is the opposite of ensuring accuracy: it is NIST actively falsifying data.

Regarding bias, recall from Section 9 that—after I challenged this table—NIST wrote "we did not cite every source that we reviewed". Then what sources of performance data *did* NIST review? How did NIST decide to present just these particular (wrong) numbers, rather than the much smaller (correct) numbers from the submission?

This is an agency that promised its analysis would be "open and transparent to the public", and that subsequently claimed "We operate transparently. We've shown all our work" (https://web.archive.org/web/20211115191840/https://www.nist.gov/blogs/taking-measure/post-quantum-encryption-qa-nists-matt-scholl). Where are the public records of NIST's work that led to Table 5?

Ultimately NIST manipulated the available numbers in a way that disfavors this submission. Whether or not this was intentional, it is the opposite of ensuring unbiased information.

As for timeliness, recall from Section 2 that various cryptographic standards are failing *right now* to protect the long-term confidentiality of user data. Once encrypted data is sent through the network and recorded by an attacker, there is nothing we can do to retroactively upgrade that data to better encryption.

12.2. Taking responsibility. NIST's information-quality standards include the following: "Although third-party sources may not be directly subject to Section 515, information from such sources, when used by NIST to develop information products or to form the basis of a decision or policy must be of known quality and consistent with the applicable information quality guidelines and standards. When such information is used, any limitations, assumptions, collection methods, or uncertainties concerning it are taken into account and disclosed."

NIST IR 8545 violated this "disclosed" rule by repeating numbers from a third-party source while suppressing mention of various limitations stated in the source—such as the critical fact that the source was talking about

speeds of "OQS version 0.9.0-rc1". Of course, properly disclosing this would have made it more difficult for NIST to pretend that the source was talking about speeds of the round-4 Classic McEliece submission.

NIST IR 8545's failure to state an access date for reference "[1]" was another violation of the requirement to disclose limitations (and a violation of basic citation rules). Web pages continually change. This particular web page provided many different numbers before NIST IR 8545 was issued (see Section 10), and went through another radical change in July 2025.

The "known quality" rule means that NIST cannot rely on ignorance as an excuse for disseminating bad information. For example, imagine that the culprits at NIST suddenly switch to claiming that they *assumed* that these measurements of "liboqs" were actually measurements of the round-4 Classic McEliece submission—big mistake, oops, but mistakes happen! Any such assumption would have violated this NIST information-quality standard. NIST wasn't allowed to disseminate information unless the quality of the information was *known*.

12.3. Interlude: Did NIST miss the "now-defunct" warning? NIST claimed in August 2025 that the prominent disclaimer in https://web.archive.org/web/20240425050637/https://openquantumsafe.org/benchmarking/ appeared only "after NIST collected the data". Amazingly, this claim *still* wasn't accompanied by a statement of the supposed date on which NIST accessed that web page.

One wonders why NIST didn't make this claim back at the beginning of May, in response to my question "Why was NIST taking numbers from this source in a March 2025 report and suppressing any mention of the disclaimer?" (https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/CSYRn8A-0zE/m/B_4ng8ShAAAJ) on NIST's pqc-forum mailing list.

Could NIST's claim be correct? Recall from Section 10 that the numbers that NIST copied (and misrepresented) appeared in the "Raw Data" file dated 2 April 2024. This archive of the disclaimer was on 25 April 2024. This leaves a few weeks during which NIST could have obtained those numbers without seeing the disclaimer, if NIST never rechecked the web page before issuing its report in March 2025.

But this can't excuse NIST suppressing *other* limitations that *were* already mentioned on the web page—such as, again, the fact that the source was measuring "OQS version 0.9.0-rc1". It also can't excuse NIST actively fabricating data, pretending that the source was instead measuring the round-4 Classic McEliece submission.

12.4. Limiting the degree of error. NIST's information-quality standards include the following: "Scientific and financial information disseminated by NIST is reliable and accurate to an acceptable degree of error as determined by factors such as the importance of the information, its intended use, time sensitivity, expected degree of permanence, relation to the primary mission(s) of the disseminating office, and the context of the dissemination, balanced against the resources required and the time available." I'll now go through these criteria.

12.4.1. Scientific information. The table at issue is certainly "scientific information", which NIST's information-quality standards define as "factual inputs, data, models, analyses, technical information, or scientific assessments based on the behavioral and social sciences, public health and medical sciences, life and earth sciences, engineering, or physical sciences", specifically including "any communication or representation of knowledge such as facts or data, in any medium or form, including textual, numerical, graphic, cartographic, narrative or audiovisual forms that involves a field identified in the preceding sentence". Measurements of a computer running software fit under both "engineering" and "physical sciences".

NIST excludes "hyperlinks to scientific information that others disseminate" and statements labeled as "opinion rather than fact". Neither of these exclusions applies to Table 5 of NIST IR 8545.

12.4.2. Importance of the information. NIST failed to quantify the weights placed upon individual comparison factors in this NIST project, so NIST can try weaseling out of *any* error by claiming that fixing the error wouldn't have changed NIST's decisions. NIST also has an incentive to avoid admitting that any of its decisions were botched.

It would, however, not be credible for NIST to retroactively claim that Table 5 is an unimportant part of NIST IR 8545. NIST IR 8545 is a 34-page report; NIST chose to place the three tables shown in Figure 2 prominently on page 6 of that report. These tables are the central speed measurements in the report. The text of the report comments again and again on speed (e.g., "Key generation in Classic McEliece is an outlier, being three orders of magnitude more costly than HQC"), and the reader understands the tables to be the report's justification for those comments.

Beyond the question of how NIST used the information, other decision-makers comparing application resources to NIST's table of key-generation costs will sometimes incorrectly conclude that they cannot afford Classic McEliece, and will select riskier options instead. A delay in correcting NIST's misinformation then means a delay in Classic McEliece deployment.

12.4.3. Time sensitivity and intended use. See Section 2, Section 4, and Section 5. A few further comments also seem warranted here.

NIST appears to have now terminated activity under the name "Post-Quantum Standardization Process" (perhaps because of my ongoing FOIA lawsuits regarding that process), but NIST's stated goal to "select a number of acceptable candidate cryptosystems for standardization" has not disappeared. NIST and other standardization agencies are continuing to work on post-quantum standardization, as illustrated by https://web.archive.org/web/20250415051610/https://csrc.nist.gov/Projects/pqc-dig-sig/. Recall also the following statement from NIST IR 8545: "After the ISO standardization process has been completed, NIST may consider developing a standard for Classic McEliece based on the ISO standard."

Furthermore, standardization is only one part of the overall decision-making process regarding which cryptographic systems to use. There is ample evidence of cost information playing a major role in this process: see my survey https://cr.yp.to/papers.html#competitions.

The primary purpose of NIST collecting and disseminating information about the speed of various post-quantum algorithms is as input to decisions of which algorithms to use. This is a broad, ongoing process.

Given how NIST has been claiming that the errors in Table 5 don't matter (see Section 16 below), presumably NIST will try arguing that the intended use of its post-quantum reports is purely as historical information about NIST's past standardization decisions. This (1) doesn't pass the laugh test, (2) isn't consistent with NIST claiming that these reports are "fundamental research" (see Section 15 below), and (3) isn't consistent with the continued usage of the report contents *by NIST*. For example, https://web.archive.org/web/20250606142005/https://www.nccoe.nist.gov/sites/default/files/2023-12/pqc-migration-nist-sp-1800-38b-preliminary-draft.pdf is a NIST report in 2025 saying "symmetric cryptographic primitives, such as block ciphers and hash functions, are not as drastically impacted by the advent of a (cryptographically relevant) quantum computer" and attributing this to NIST IR 8413, NIST's round-3 report.

12.4.4. Dissemination context and the expected degree of permanence. NIST IR 8545 is not a momentary excerpt from a continually updated series of data releases: it is a permanent report advertised by NIST very close to the top of NIST's "Post-quantum cryptography" page (see https://web.archive.org/web/20250726044753/https://csrc.nist.gov/projects/post-quantum-cryptography).

Formally, "permanent" isn't exactly right—I should say "permanent by default", since NIST has the power to issue corrections to its reports. But this is not helpful when NIST stonewalls in response to correction requests. In the context of a correction request, a question about permanence doesn't make sense if it's considering the case of NIST accepting the correction request.

12.4.5. Relation to the primary mission(s) of the disseminating office. NIST IR 8545 is a prominent part of a several-year several-person project by the disseminating office at NIST, a project that has been repeatedly highlighted in activity reports by that office and that plays an indispensable role in White House actions. See, e.g., https://nist.pqcrypto.org/foia/20250211/RE_%20PQC%20announcement%20is%20ready.pdf as an example of NIST discussing a White House request for NIST to make an announcement within this project, and https://bidenwhitehouse.archives.gov/briefing-room/statements-release

s/2022/05/04/national-security-memorandum-on-promoting-united-states-leadership-in-quantum-computing-while-mitigating-risks-to-vulnerable-cryptographic-systems/ ordering agencies to use the NIST standards resulting from this project.

12.4.6. The resources required and the time available. It would have taken under a minute for NIST to copy the correct speed information from the submission. NIST should have also *checked* the information by re-running the software—which is slightly more effort, but still a straightforward process that NIST reported carrying out for many round-1 submissions within a month of the initial submission deadline in 2017 (https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/zeHJEzWdv2Y/m/yTg4wG7cBwA). This particular employee subsequently left NIST, but NIST has other employees capable of following the same straightforward process.

12.4.7. The degree of error. NIST's 114189 claim is an astounding 4× above reality (see Section 8.1). This gap comes from NIST taking numbers from a third party talking about the cost of "liboqs" and falsely presenting those numbers as the cost of the round-4 Classic McEliece submission.

Misrepresenting what a source wrote wouldn't even be acceptable from a student carrying out a homework assignment, never mind an agency carrying out a several-year project to evaluate technology choices for a multi-billion-dollar industry.

12.5. Reporting uncertainty. NIST's information-quality standards include the following: "NIST Policy 1800.00 requires that NIST measurement results be accompanied by quantitative statements of their uncertainty, and that a uniform approach to expressing measurement uncertainty be followed."

As noted in Section 6, the tables in NIST IR 8545 fail to include this basic information.

For comparison, Figure 1 includes quartiles—which are far apart for mceliece348864 keygen. This is not because of measurement error but because the cost of mceliece348864 keygen has a significant dependence upon random choices made in each algorithm run. Robust statistics for this scenario are presented in https://cr.yp.to/papers/rsrst-20250727.pdf but readers in any case need to be warned regarding the variability.

I've taken mceliece348864f, where keygen quartiles are much closer together, as my running example rather than mceliece348864. The lower variability of mceliece348864f makes comparisons easier. But NIST's numbers are also wrong for mceliece348864, for mceliece460896, etc.

12.6. Scrutiny. NIST's information-quality standards include the following: "NIST treats information quality as integral to every step in its process of developing the information, including creation, collection, maintenance, and dissemination. All scientific information disseminated by NIST receives a level of scrutiny commensurate with the critical nature of the information and its intended use."

Obviously this scrutiny did not take place here. NIST copied and pasted third-party numbers regarding "liboqs" and falsely presented those as measurements of the round-4 Classic McEliece submission.

13 Procedural burdens in NIST's information-quality standards

NIST's information-quality standards prominently state that "NIST is committed to maintaining a high level of quality in the information it disseminates". Despite this supposed commitment, the subsequent details of the standards add various exemptions, placing extra burdens upon people requesting corrections.

13.1. Affected people. NIST's correction procedures require "an explanation of how the requester is affected", and define "affected person" as "an individual or entity that uses, benefits from, or is harmed by the disseminated information at issue".

To answer this, I'll start by quoting the Policy Statement on Ethical Guidelines from the American Mathematical Society (https://www.ams.org/about-us/governance/policy-statements/sec-ethics): "The correct attribution of mathematical results is essential, both because it encourages creativity, by benefiting the creator whose career may depend on the recognition of the work and because it informs the community of when, where, and sometimes how original ideas entered into the chain of mathematical thought." Specific responsibilities include giving "appropriate credit" and using "no language that suppresses or improperly detracts from the work of others".

As noted above, I am one of the members of the team that developed Classic McEliece, both initially and through subsequent improvements that led to round-4 Classic McEliece. What NIST is doing in Table 5 of NIST IR 8545 is exaggerating the cost, specifically the CPU time, of round-4 Classic McEliece. NIST is improperly detracting from the work of all of the scientists who contributed to round-4 Classic McEliece.

Furthermore, NIST's citation and usage of the "liboqs" results fails to give appropriate credit to https://bench.cr.yp.to, a project that had the relevant benchmarks before anyone else, along with careful attention to accuracy, verifiability, comprehensiveness, etc.

Given that failure to provide appropriate credit is recognized by professional societies as an ethics violation, obviously the failure is harmful, meeting the "harm" criterion in NIST's procedures. But I think it's also important to emphasize how inappropriate this procedural constraint is. The real damage done by NIST's misinformation is via people using the information *without realizing that the information is wrong*. Corrections should not have to wait until those people *recognize* the damage.

The same AMS ethics statement has a requirement "to correct in a timely way or to withdraw work that is erroneous"—no exceptions. NIST's failure to correct Table 5 of NIST IR 8545 is unethical.

13.2. Useful corrections. NIST's procedures also say that NIST will not consider corrections that "would serve no useful purpose", such as corrections to information that is "not valid, used, or useful after a stated short period of time".

NIST IR 8545 does not state a time limit on the usage of its information regarding Classic McEliece. On the contrary, as noted above, there is ongoing activity by NIST and others regarding this. Consider again the following statement in NIST IR 8545: "After the ISO standardization process has been completed, NIST may consider developing a standard for Classic McEliece based on the ISO standard." And consider again https://mceliece.org showing that there are already many users of Classic McEliece. Cryptographic decisions are happening all the time; there is no excuse for NIST misrepresenting what is known about the features of the options that are available.

By writing that it did not rely "solely" on the specific information at issue, NIST implicitly admits that it used this information. The correct information is indisputably useful.

14 NIST's unjustified delays in handling this correction request

Recall the following DOC rule quoted in Section 11.4: "The operating unit will respond to all initial requests within 60 calendar days of receipt. If the request requires more than 60 calendar days to resolve, the operating unit will inform the complainant that more time is required and indicate the reason why and an estimated decision date." DOC also wrote in https://www.federalregister.gov/documents/2002/10/08/02-25340/guidelines-for-ensuring-and-maximizing-the-quality-objectivity-utility-and-integrity-of-disseminated that responses will be "usually within 60 calendar days".

I filed a formal request for correction on 7 May 2025. Two months later, NIST wrote to me (email dated 7 Jul 2025 15:09:20 +0000) as follows: "I am writing to provide you with an update on the response to your request for correction under NIST's Information Quality Standard Guidelines. The response is currently undergoing review, and we will send it to you as soon as the review process is concluded."

In violation of DOC's rules, NIST failed to state an "estimated decision date". Furthermore, the generic

reference to "undergoing review" does not meet DOC's rule requiring an indication of "the reason why" there is a delay. The rule applies *if* a specific request takes more than 60 days, so "the reason why" also has to be specific. This is not a free pass allowing NIST to generically delay *all* responses, or to delay responses on issues where it has internally decided to stonewall.

I wrote back (email dated 7 Jul 2025 17:37:26 +0200) as follows: "Thank you for the update. Do you have an estimated decision date?"

NIST wrote back (email dated 7 Jul 2025 17:15:31 +0000) as follows: "We are required to forward all responses to requests for correction to OMB for review, those reviews are coordinated by the Department of Commerce, and as such, we are unable to provide an expected completion date."

But the DOC rules (see quote and link above) *require* NIST to provide an estimated decision date. I shouldn't even have had to ask. The generic commentary about slowdowns for "all responses to requests for correction" again fails to meet DOC's requirement to explain "the reason why" a specific request is not being handled within 60 days.

I wrote back (email dated 7 Jul 2025 19:45:33 +0200) quoting the following NIST rule: "The Director, NIST Management and Organization Office will attempt to communicate either a decision on the request, or a statement of the status of the request and an estimated decision date, within 60 calendar days after receipt of the request." I continued as follows:

> I understand that ‘‘attempt’’ isn't a guarantee, but failing to reach the 60-day target and failing to provide an estimated decision date would seem to require justification specific to the circumstances at hand, rather than a generic statement about OMB involvement in all responses.
>
> I find it particularly puzzling to not see a prompt correction for an amply documented case of NIST fabricating data. Please say concretely what communication has occurred regarding this specific matter and why this is taking so long.

NIST didn't respond. When NIST finally wrote to me in August answering my formal correction request, it still didn't explain why this had taken so long.

15 NIST's fantasy exemption

Back in May, my formal correction request quoted the information-quality standards that NIST was violating, and traced how NIST was violating those standards. Recall that Public Law 106-554 requires "administrative mechanisms allowing affected persons to seek and obtain correction of information maintained and disseminated by the agency that does not comply with the guidelines issued under subsection (a)"; this *requires* NIST to correct information not complying with NIST's information-quality standards, whenever an affected person invokes these mechanisms to request a correction.

Procedurally, NIST's information-quality standards require "a point-by-point response to any relevant data quality arguments contained in the request". But NIST's response in early August avoided quoting my request, and avoided quoting NIST's information-quality standards.

NIST started by making a bizarre *claim* about those information-quality standards: namely, that the entire NIST IR series is categorically exempt from NIST's information-quality standards. That's thousands of public NIST reports on a wide range of topics. Basically, NIST is claiming that it can publish whatever information it wants without complying with its own information-quality standards, simply by labeling the information as a NIST IR.

Rather than relying on this claim, NIST continued by attempting to argue that the content of NIST IR 8545 is just fine. I'll address that in Section 16. This section focuses on debunking NIST's claim of an exemption.

15.1. NIST's statement of the fantasy. Here's the paragraph where, without quoting its information-quality standards, NIST attempts to claim a giant exemption from those standards:

> NIST IRs are part of the NIST Technical Series Publication and defined as interagency or internal reports of research findings, including background information for Federal Information Processing Standards (FIPS) and NIST Special Publications (NIST SP)[3]. NIST treats NIST IRs as fundamental research communication. As stated in the NIST Guidelines, while fundamental research communications are expected to adhere to information aquality standards and do receive technical, policy, and editorial review by the NIST Editorial Review Board in addition to the pre-dissemination review methods listed in the NIST Guidelines, they are not included in the scope of the OMB and NIST Guidelines.

(Typo "aquality" in original.) Footnote 3 links to https://csrc.nist.gov/publications#:~:text=NIST%20Interagency%20or%20Internal%20Reports%20(NIST%20IR),NIST%20Cybersecurity%20and%20Privacy%20Program.

This is claiming that NIST IRs are categorically "fundamental research communication", and that "fundamental research communications" are "not included in the scope of the OMB and NIST Guidelines". Every aspect of this claim is wrong.

15.2. Public Law 106-554. Let's start with what the statute actually says, keeping in mind that the statute is binding upon OMB, DOC, and NIST. Public Law 106-554 requires "guidelines ensuring and maximizing the quality, objectivity, utility, and integrity of information (including statistical information) disseminated by the agency". It makes no exceptions for "fundamental research communications".

15.3. The OMB rules. Next, let's look at what the OMB rules (https://obamawhitehouse.archives.gov/sites/default/files/omb/fedreg/reproducible2.pdf and https://web.archive.org/web/20250422071633/https://trumpwhitehouse.archives.gov/wp-content/uploads/2019/04/M-19-15.pdf) actually say. The rules define "dissemination" as "agency initiated or sponsored distribution of information to the public".

OMB reports that concerns were expressed "particularly in the context of scientific research conducted by Federally employed scientists or Federal grantees who publish and communicate their research findings in the same manner as their academic colleagues". OMB responded that "information generated and disseminated in these contexts is not covered by these guidelines unless the agency represents the information as, or uses the information in support of, an official position of the agency".

In more detail, OMB explained that an agency-initiated distribution "refers to information that the agency disseminates, e.g., a risk assessment prepared by the agency to inform the agency's formulation of possible regulatory or other action", along with cases where "an agency, as an institution, disseminates information prepared by an outside party in a manner that reasonably suggests that the agency agrees with the information".

OMB contrasted this with the situation where "a Federally employed scientist or Federal grantee or contractor publishes and communicates his or her research findings in the same manner as his or her academic colleagues". OMB concluded as follows: "To avoid confusion regarding whether the agency agrees with the information (and is therefore disseminating it through the employee or grantee), the researcher should include an appropriate disclaimer in the publication or speech to the effect that the 'views are mine, and do not necessarily reflect the view' of the agency."

Does NIST IR 8545 say that it is merely a publication of views of individual NIST employees, not necessarily representing the view of NIST? No. NIST IR 8545 is an official report of NIST's selection of what to standardize. It bears the NIST logo, the name of the Secretary of Commerce, and the name of the acting NIST director. Its contents issue a variety of new statements in the name of NIST, such as "NIST may consider

developing a standard for Classic McEliece based on the ISO standard" and "NIST determined that HQC would provide a good complement to ML-KEM".

Is NIST IR 8545 published in the same way as academic research, for example as a paper in a journal? No. Is NIST IR 8545 research in the first place? No. It's a report that merely *cites* some research. It also doesn't claim to be research; it claims to be a report describing "the evaluation and selection process of these fourth-round candidates based on public feedback and internal review".

Is the specific table at issue, Table 5, an example where "a Federally employed scientist or Federal grantee or contractor publishes and communicates his or her research findings" (never mind "in the same manner as his or her academic colleagues")? No. The numbers come from another source. Misrepresenting what that source measured (and miscopying some of the numbers) is not research.

NIST is grossly mischaracterizing the OMB rules when it claims that the scope of those rules categorically excludes "fundamental research communications". The actual OMB distinction is as follows:

- OMB's rules cover *actions by the agency*: in particular, "agency initiated or sponsored distribution of information to the public", such as information "prepared by the agency to inform the agency's formulation of possible regulatory or other action".
- OMB's comments explain that this is different from *actions by individual agency employees*: for example, "a Federally employed scientist or Federal grantee or contractor publishes and communicates his or her research findings in the same manner as his or her academic colleagues", with "an appropriate disclaimer" to avoid confusion.

NIST's distribution of Table 5 of NIST IR 8545 to the public was both initiated and sponsored by NIST, so it constitutes NIST dissemination of this (mis)information, making NIST responsible for issuing a correction.

15.4. The NIST rules. NIST also claims that NIST's information-quality standards categorically exclude "fundamental research communications"; that this is "stated" in those standards; and that "NIST treats NIST IRs as fundamental research communication".

The actual text of NIST's information-quality standards (https://web.archive.org/web/20250507112244/https://www.nist.gov/director/nist-information-quality-standards) does not back up any of these claims.

The phrase "fundamental research communications" appears twice in the actual text. First, there is a sentence "It does not include scientific information disseminated in the form of fundamental research communications". But the word "It" here does not refer to *the information-quality standards*: it refers to NIST's definition of "Scientific, financial, and statistical information". Here is the complete paragraph:

> **Scientific, financial, and statistical information** includes all scientific, financial, and statistical information disseminated to the public by NIST in formats including but not limited to:
>
> * Standard Reference Data
>
> * Standard Reference Materials
>
> * Economic impact studies
>
> * Exhibits
>
> * Web sites.
>
> It does not include scientific information disseminated in the form of fundamental research communications.

(Boldface in original.)

In Section 12, I listed five examples of how the NIST IR 8545 table at issue fails to comply with NIST's information-quality standards. Two of the standards that I quoted involve the phrase "scientific information". But that phrase has a separate definition in NIST's information-quality standards, with its own inclusions and exclusions. Even if that definition were somehow overridden by the definition of "scientific, financial, and statistical information" in a way relevant to the table at issue, there would still be three other failures of the table to comply with NIST's information-quality standards—and *any* failure is enough to force a correction under Public Law 106-554.

The second appearance of the phrase "fundamental research communications" is as follows:

> Although they are not included in the scope of the OMB Guidelines, NIST also expects fundamental research communications to adhere to information quality standards.

This is not an exclusion of "fundamental research communications" from NIST's information-quality standards: it is saying that "fundamental research communications" *are* covered by NIST's information-quality standards, despite (supposedly) not being covered by OMB's rules.

NIST's information-quality standards even state specific requirements for "fundamental research communications" to be reviewed "by the NIST Editorial Review Board". How can the information-quality standards impose rules on "fundamental research communications" if, as NIST now claims, "fundamental research communications" are outside the scope of the information-quality standards? This makes no sense.

Furthermore, NIST's information-quality standards do not say that NIST IRs are categorically "fundamental research communication". The particular report at issue is not even research, let alone "fundamental research".

15.5. Relative powers of authorities. Imagine NIST issuing replacement information-quality standards that categorically exempt "fundamental research communication" and that categorically exempt NIST IRs.

This hypothetical exemption would be in direct conflict with Public Law 106-554, which covers all "information (including statistical information) disseminated by the agency". Public Law 106-554 is a statute; NIST is powerless to make rules contrary to a statute. Public Law 106-554 does not exempt "fundamental research communication", and does not exempt NIST IRs.

This hypothetical exemption would also be in direct conflict with OMB's rules, which define "disseminated by the agency" to include all "agency initiated or sponsored distribution of information to the public". OMB has statutory authority over publications by all agencies (see Section 11.1), and NIST is powerless to issue publications that violate OMB's rules. OMB's rules do not exempt "fundamental research communication", and do not exempt NIST IRs.

In any event, NIST IR 8545 has to comply with the information-quality standards that were in place when NIST IR 8545 was prepared and issued. Those standards do not have NIST's fantasy exemption.

16 NIST's August 2025 refusal to admit error

I'll now go point by point (focusing on content, not trying to preserve line breaks) through NIST's early-August refusal to correct NIST IR 8545.

> This letter is in response to your May 7, 2025 request for correction (Request) of certain information contained in the NIST IR 8545, Status Report on the Fourth Rounds of the NIST Post-Quantum Cryptography Standardization Process[1].

(Footnote 1 links to https://nvlpubs.nist.gov/nistpubs/ir/2025/NIST.IR.8545.pdf.)

That's not exactly the title of NIST IR 8545, but, yes, NIST IR 8545 is the report for which I'm requesting a correction.

> Your Request is submitted under the Information Quality Act, Section 515 of Public Law 106-554; the Office of Management Budget's Guidelines for Ensuring and Maximizing the Quality, Objectivity, Utility, and Integrity of Information Disseminated by Federal Agencies (67 FR 8451, Feb. 22, 2002) (''OMB Guidelines''); the Office of Management and Budget's Memorandum on Improving Implementation of the Information Quality Act, M-19-15 (''OMB Memo''); and NIST's Guidelines, Information Quality Standards, and Administrative Mechanism (''NIST Guidelines'').

This is not accurately describing my request. My request was labeled as "a request under Section 515 of Public Law 106-554" and cited NIST's information-quality standards; it didn't cite the OMB rules.

However, it's true that NIST is bound by OMB's rules, via OMB's authority under statute to control NIST's publications (see Section 11.1).

> The Request asserts that the following information in NIST IR 8545 requires correction: 1. That the number ''114 189'' in Table 5 is incorrect, overstating the number of clock cycles by approximately a factor of 4.

This is sloppily worded. Readers expect the items enumerated after the colon to be examples of the "information in NIST IR 8545", but item 1 (namely, that 114189 is incorrect) is not information in NIST IR 8545; it is information from my correction request.

Anyway, yes, the 4× gap between 114189 and reality is the issue in a nutshell. Of course, since all of the claims in Table 5 arise from the same NIST falsification, all of the claims should be replaced with the correct numbers from Figure 1 (scaled by 0.001 to kcycles to avoid confusion given the context).

> 2. That this number incorrectly reflects performance of the second-round software for Classic McEliece, not the fourth-round implementation.

This has an overlap with the actual problem here—namely, Table 5 of NIST IR 8545 falsely purports to be measurements of the round-4 Classic McEliece submission—but this sentence by NIST misstates my correction request in two basic ways.

First: No, my correction request did not use the word "implementation", which has precise meanings in some contexts (e.g., in SUPERCOP) but is ambiguous and confusing for a general audience. I don't know what NIST thinks it means by "the fourth-round implementation" of Classic McEliece.

Second: No, I didn't state round 2 vs. round 4 as a separate request for correction. I stated it as a diagnosis of *how* NIST arrived at the grossly inaccurate numbers in Table 5. See Section 8 of this document for how NIST's procedural violations ended up naturally leading to this diagnosis, and see Section 10 for bulletproof evidence for the correct diagnosis.

> 3. That NIST cited a defunct and unmaintained source without disclosing the limitations and omitted access dates for the citation.

No, I didn't state this as a separate request for correction. NIST's procedures demanded an explanation not just of how NIST's data was *wrong* but also of how NIST was *violating its own information-quality standards*. My explanation of the details was explicitly in this context.

NIST is also not accurately summarizing my explanation. For example, within the section reviewing NIST's violations, my correction request had a subsection regarding NIST's violations of "NIST Policy 1800.00 requires that NIST measurement results be accompanied by quantitative statements of their uncertainty". NIST simply ignores this point, in violation of the procedural requirement of "a point-by-point response to any relevant data quality arguments contained in the request".

Of course, the discipline of simply *quoting* the request would have made it much more difficult for NIST to avoid a point-by-point response.

> 4. That NIST disregarded better benchmarking sources, such as the SUPERCOP/bench.cr.yp.to platform, which is more accurate, current, and comprehensive.

Yes, that's a better benchmarking source; but, no, I didn't state this as a separate request for correction. Also, "disregarded" is a strawman from NIST, not something my correction request had said.

I mentioned https://bench.cr.yp.to in three contexts. First, I mentioned it as an example of how easy it is to find the correct numbers; this was explicitly in the context of addressing NIST's "balanced against the resources required and the time available" criterion. Second, I mentioned it as an example of presenting quantitative indications of variability; this was explicitly by comparison to NIST IR 8545's violation of NIST's requirement on that topic. Third, I mentioned it as an example of how NIST was failing to give appropriate credit; this was explicitly in the context of addressing NIST's demand for "an explanation of how the requester is affected".

> You cite previous correspondence regarding the Table 5 number on a ``pqc-forum'' mailing list[2].

(Footnote 2 links to https://groups.google.com/u/1/a/list.nist.gov/g/pqc-forum/c/CSYRn8A-0zE/m/6rhWGzEgDgAJ.)

It's true that the request cites correspondence on pqc-forum, but it is strange that NIST describes this as "a" mailing list rather than describing it as *the* public mailing list that NIST operates for discussions of NIST's post-quantum standardization efforts.

NIST also does not admit why I was citing this particular correspondence, namely to quote and challenge NIST's claim to have "verified that the performance figures given in the tables in our report accurately reflect the benchmarks we cited". Gee, what a surprise to see NIST repeating essentially the same claim three months later.

> According to the Request, the supposed error misinformed the cryptographic community and delayed deployment of Classic McEliece,

No, this is another fabrication by NIST. Here's what my request actually said: "I don't expect NIST's misinformation about Classic McEliece to be able to stop Classic McEliece deployment, but I do expect it to slow down deployment if it is not promptly corrected." Also: "You can see from https://mceliece.org that there are already many users of Classic McEliece, and further users who have expressed interest in Classic McEliece. Cryptographic decisions are happening all the time. There is no excuse for NIST misrepresenting what is known about the features of the options that are available."

> thereby impacting efficiency and public security.

What is "public security" supposed to mean?

This is a mangling by NIST of the phrase "damaging security, damaging efficiency, and harming the general public" from my request.

There is no legitimate reason for NIST to avoid quoting what the request actually said.

> NIST IRs are part of the NIST Technical Series Publication and defined as interagency or internal reports of research findings, including background information for Federal Information Processing Standards (FIPS) and NIST Special Publications (NIST SP)[3]. NIST treats NIST IRs as fundamental research communication. As stated in the NIST Guidelines, while fundamental research communications are expected to adhere to information aquality standards and do receive technical, policy, and editorial review by the NIST Editorial Review Board in addition to the pre-dissemination review methods listed in the NIST Guidelines, they are not included in the scope of the OMB and NIST Guidelines.

See Section 15 for a detailed response to this.

> Even if, assuming arguendo, NIST IR 8545 were to be included within the scope of the NIST Guidelines, your requests for correction would not be warranted, as discussed below.

NIST IR 8545 is within the scope of the NIST guidelines, and my requests for correction are amply justified.

> *1. **Request for Correction:** The number ''114 189'' in Table 5 is incorrect, overstating the number of clock cycles by approximately a factor of 4.*

(Italics and boldface in original.) Yes, that's again the issue in a nutshell.

> NIST recognizes that different benchmarking platforms have reported varying results using different implementations of the Classic McEliece algorithm,

Most readers will understand this as follows: (1) "benchmarking platforms" disagree regarding the performance of the round-4 Classic McEliece submission; (2) NIST picked the numbers from one "benchmarking platform" for its report; (3) the report "recognizes" that there are variations.

#3 is not true. The report *suppresses* variations rather than *recognizing* them. The ambiguous phrasing of "NIST recognizes" allows NIST to claim that it wasn't referring to what the *report* said—for example, it's referring to NIST's quiet awareness of variations, or to NIST now admitting variations—but NIST can't defend *the report's* suppression of variations by saying that the correct information is available elsewhere.

But there's a much bigger problem with NIST's "different benchmarking platforms" narrative: namely, #1 is wildly misrepresenting what's actually going on here.

I don't know what NIST thinks it means by a "benchmarking platform", but let's just assume that the measurement tool in "liboqs" somehow qualifies. Yes, that's different from SUPERCOP. No, this difference is *not* the reason for the immense gap between the "liboqs" numbers and the speeds of the round-4 Classic McEliece submission.

The actual reason for the gap is that the "liboqs" numbers *are not measurements of the performance of the round-4 Classic McEliece submission.* They are instead measurements of slower software (specifically the third-party software identified in Section 10). NIST misrepresented the results as measurements of the round-4 Classic McEliece submission.

Given past events, I presume that the culprits at NIST won't admit how they're deceiving readers with this "different benchmarking platforms" narrative. Instead they'll exploit the ambiguity of their own wording

"implementations of the Classic McEliece algorithm" to argue that "liboqs" qualifies as such an "implementation". But this is irrelevant to my correction request regarding NIST IR 8545, whereas what NIST leads readers to believe—namely, that the gaps come from "different benchmarking platforms" measuring the submission differently—would be highly relevant if it weren't divorced from reality.

This ambiguous "implementations" concept is neither necessary nor helpful for understanding the central facts:

- NIST solicited submissions "from any interested party for candidate algorithms to be considered for public-key post-quantum standards".
- NIST's regulations for the project pointed to "criteria that will be used to appraise the candidate algorithms".
- The fourth-round project report, NIST IR 8545, claims to describe "the evaluation and selection process of these fourth-round candidates".
- The report specifically says that the speed tables are for these submissions: "Tables 3 through 5 show representative benchmarks for key generations, encapsulations, and decapsulations of BIKE, HQC, and Classic McEliece, respectively. Each row is a specific parameter set from the corresponding submission."
- But Table 5 of NIST IR 8545 actually describes a slower algorithm assembled by a third party, *not* the round-4 Classic McEliece submission.

NIST misrepresented that third-party algorithm as the fourth-round candidate algorithm—and, even worse, is refusing to correct this.

> `with some reporting lower cycle counts than those listed in Table 5.`

The "some" wording here, along with the earlier comment about "varying results", makes the reader imagine a large array of benchmarking platforms, with the majority producing numbers in line with NIST's Table 5, while "some" produced smaller numbers.

There's a reason that NIST doesn't—and can't—provide URLs supporting the image that NIST is creating in the reader's mind: namely, the image is pure fiction. Did NIST spend three months trying to figure out wording that would provide the best tradeoffs between (1) leading readers to the wrong idea and (2) evading NIST responsibility for this deception?

NIST took the numbers in Table 5 of NIST IR 8545 from *one* source. The source said it was measuring "OQS version 0.9.0-rc1". NIST suppressed the "OQS version 0.9.0-rc1" information and instead misrepresented the numbers as the speeds of the round-4 Classic McEliece submission. See Section 10.

> `However, this figure in Table 5 was accurately cited from a publicly available benchmark source--- specifically, Reference [1] in NIST IR 8545, which corresponds to the Open Quantum Safe (OQS) project[4].`

(Footnote 4 repeats bibliography entry 1 from NIST IR 8545.)

What is "accurately cited" supposed to mean?

I'm not disputing that NIST provided the right URL. I'm not disputing that a cycle count matching 114189 kcycles appeared at that URL (starting in April 2024). I'm saying that NIST is (1) suppressing the source's statement of what the source was measuring and (2) pretending that the source was measuring something else.

> `NIST IR 8545 accurately reproduces the benchmark results published there.`

No, it does not. See Section 10.

> Importantly, only one out of 48 data points drawn from Reference [1] across Tables 3, 4, and 5 has been challenged,

"Importantly"? Why exactly is this supposed to be important for my correction request? Surely NIST isn't claiming that falsifying data is acceptable if it doesn't happen too often.

Anyway, no, what NIST is saying here (and claiming to be "important") simply isn't true. I challenged the entirety of Table 5. I gave the number 114189 as an early example to explain "the problem in a nutshell", but my very next paragraph explicitly broadened this to challenging "the numbers in that table", and said "NIST is falsely portraying these numbers as evaluations of the fourth-round Classic McEliece software."

> and even this challenged data point was correctly cited.

Covered above.

> The allegation that the report fabricated or misstated data is unfounded.

If a data point says "The observed temperature in D.C. at 23:59 on 8 August 2025 was 67 degrees", manipulating it to say "The observed temperature in D.C. at 23:59 on 8 August 2025 was 73 degrees" is data fabrication. Manipulating it to say "The observed temperature in Anchorage at 23:59 on 8 August 2025 was 67 degrees" is also data fabrication. The data point includes the number *and* the statement of what the number was observing.

NIST took information from a source that said it was measuring "OQS version 0.9.0-rc1". NIST removed the original labeling of these numbers as measurements of "OQS version 0.9.0-rc1". NIST pretended that these were measurements of the round-4 Classic McEliece submission. This is NIST fabricating data, and more specifically NIST falsifying data. The misconduct here is similar to, e.g., the case reported in https://retractionwatch.com/2025/03/21/osaka-dental-university-fabricated-data-investigation/ ("Because images in the other two articles were identical to images in the 2014 paper, the university determined the data in the later papers were fabricated").

> As stated in NIST IR 8545, the performance figures are intended to be representative,

It's true that NIST IR 8545 says "Tables 3 through 5 show representative benchmarks for key generations, encapsulations, and decapsulations of BIKE, HQC, and Classic McEliece, respectively. Each row is a specific parameter set from the corresponding submission." But the numbers in Table 5 *aren't* benchmarks of the Classic McEliece submission.

> not definitive,

No, NIST IR 8545 doesn't say that the performance figures aren't definitive, nor is this relevant to my request for NIST to correct its falsified numbers.

> and reasonable differences in benchmark platforms or methodologies do not constitute misinformation.

This is another smokescreen. The reason Figure 1 and Figure 2 are so different is not because of NIST's silent switch away from its designated comparison platform (see Section 8.1), and is not because of different measurement methodologies. It is because NIST took measurements of a much slower third-party algorithm and misrepresented those as measurements of the round-4 Classic McEliece submission.

> *2. **Request for Correction:** This number reflects the performance of the second-round software for Classic McEliece, not the fourth-round software implementation.*

(Italics and boldface in original.) No, this was not a separate request for correction. See above. NIST should be quoting and responding to the actual request for correction, not to NIST's mischaracterizations of the request.

> NIST disagrees with this assertion.

I agree that—because NIST's procedural violations obscured the facts—I was originally misdiagnosing *how* NIST was obtaining its fake numbers. See Sections 8 and 10. But this doesn't rescue NIST IR 8545: the numbers are still fake, and still ludicrously inaccurate.

> NIST is confident that the benchmarks cited from Reference [1] accurately reflect the fourth-round version of Classic McEliece.

NIST provides no evidence backing up this claim of accuracy, and in any event the claim is false. See above.

> NIST IR 8545 focuses specifically on fourth-round submissions, and the benchmark data used aligns with that scope. Accordingly, the performance measurement in question is consistent with the version under evaluation.

No. NIST took measurements of a slower third-party algorithm and misrepresented those as measurements of the round-4 Classic McEliece submission.

> *3. **Request for Correction:** NIST cited a defunct and unmaintained source without disclosing the limitations and omitted access dates for the cited source.*

(Italics and boldface in original.) No, this was not a separate request for correction. See above.

> The cited source--- Open Quantum Safe algorithm performance visualizations --- was an active and widely used project at the time NIST gathered benchmark data.

Active in the sense of having a functioning web page, yes. Widely used, no. As of 8 August 2025, Google Scholar reports just 18 publications mentioning https://openquantumsafe.org/benchmarking (see https://scholar.google.com/scholar?hl=en&as_sdt=0%2C14&q=%22openquantumsafe.org%2Fbenchmarking%22&btnG=), two of which are NIST IR 8545 and an earlier NIST report.

> It played a visible and constructive role in the NIST post-quantum cryptography process.

This has no evident content beyond saying that NIST learned about this source and chose to cite it.

> The disclaimer currently shown on the OQS website was added after NIST collected the data.

This might be correct (see Section 12.3), but does not rescue NIST IR 8545.

> NIST accurately cited the relevant information as it existed at the time. The addition of a later disclaimer does not retroactively invalidate the benchmark or its citation.

NIST suppressed critical limitations of the information from that source and added fake information instead, pretending that the source was measuring the round-4 Classic McEliece submission.

> *4. **Request for Correction:** NIST disregarded better benchmarking sources, such as the SUPERCOP/bench.cr.yp.to platform, which is more accurate, current, and comprehensive.*

(Italics and boldface in original.) No, this was not a separate request for correction. See above.

> Benchmarking results naturally vary depending on the platform, implementation, and methodology.

Covered above.

> As stated in the report, the benchmark tables were intended to provide representative data points,

Covered above.

> not authoritative

No, the report had no such disclaimer, nor would such a disclaimer be relevant to the problem at hand, namely NIST fabricating data.

> or exhaustive metrics.

Irrelevant. Table 5 uses the typical speed metrics (keygen cycles, enc cycles, dec cycles), and presents misinformation regarding the performance of the round-4 Classic McEliece submission in those metrics.

> NIST did not disregard SUPERCOP;

Here NIST is responding to its own strawman. My correction request did not say that NIST disregarded SUPERCOP.

> it was among the sources NIST reviewed during the evaluation process.

One is forced to wonder, then, how NIST could possibly have missed the gigantic discrepancies between NIST's claims and the correct numbers from SUPERCOP.

Meanwhile NIST had promised to carry out "a thorough analysis of the submitted algorithms". The submission provided the software and accurately documented its speed.

> Ultimately, NIST selected Reference [1] because it provided a consistently structured and well-documented dataset that served NIST's goal of presenting a fair, representative comparison across submissions.

No, that reference did not cover the round-4 Classic McEliece submission. That reference presented measurements of a much slower third-party algorithm. Having NIST substitute this for the actual speeds would have been glaringly unfair even if the substitution had been stated openly: reporting slowed-down numbers for a submission deceives readers who miss the statement that those are slowed-down numbers, so why do this selectively for only *one* submission? Certainly the reader is at fault, but this doesn't remove the unfairness. Anyway, having this as a *silent* substitution was even worse, and means that the reader *isn't* at fault.

I do agree that "liboqs" seems to adequately document its data. But this raises the question of how NIST could possibly have misunderstood this reference as being about the round-4 Classic McEliece submission. (See Section 10.6.)

As a separate matter, NIST's claim about its specific goal here lacks credibility. Consider the following facts:

- NIST asked submission teams to provide "an AVX2 (Haswell) optimized implementation".
- NIST said the optimized implementation was "to demonstrate the performance of the algorithm".
- NIST suddenly switched from Intel Haswell CPUs to an Intel Cascade Lake CPU for NIST IR 8545.

Doesn't fairness force NIST to stick to its announced comparison target? Shouldn't NIST have immediately rejected the "liboqs" numbers for not covering Haswell? Only a few CPUs are covered by the "liboqs" source that NIST cites.

NIST's selected data set also flunks the basic quality requirement of accompanying measurements by quantitative statements of their uncertainty, such as standard deviations or quartiles. (I saw deviations when I tried "liboqs", but the data sets assembled on the web pages cited by NIST fail to include those.) NIST is required to quantify uncertainty, so NIST is not allowed to use third-party numbers that fail to quantify uncertainty. Remember that information from third-party sources "when used by NIST to develop information products or to form the basis of a decision or policy must be of known quality and consistent with the applicable information quality guidelines and standards". I already quoted this rule in my formal correction request; NIST didn't respond.

Meanwhile the SUPERCOP data is consistently structured, is amply documented, covers many more CPUs (including Haswell CPUs), provides quartiles, and, unlike the reference that NIST used (and misrepresented), covers the round-4 Classic McEliece submission. Also, SUPERCOP's tables report speeds of the *fastest* software for each "primitive", so slower software—software farther from the speeds that the user will end up seeing, if the user cares about performance—doesn't corrupt the tables. What NIST did in Table 5 is the same type of corruption.

> NIST recognizes that experts may have different preferences for benchmarking sources, but choosing one over another for clarity and consistency does not imply negligence or bias.

This sentence sounds reasonable in isolation; it's also irrelevant to the problem at hand.

> Regarding the alleged harm, we emphasize that key generation cycle counts are only one of many considerations for weighing alternatives for standardization. Suitability and competitiveness of a given algorithm for standardization depend on a variety of security, performance, and implementation characteristics across different use cases. Improved key generation benchmarks alone would not change the conclusions.

See Section 12.4 for a detailed response to this.

> Based on the foregoing, your Request is denied. In preparing this response, NIST carefully reviewed the relevant information and sources cited to ensure that its

> conclusions are accurate, well-founded, and consistent with NIST's information quality standards.

The table at issue isn't accurate and isn't well-founded. The table violates several of NIST's information-quality standards. If NIST were applying an appropriate level of care then it would have fixed this table a long time ago.

> You may appeal this decision within 30 calendar days from the date of this decision. Such an appeal must be in writing and addressed to: NIST Associate Director for Laboratory Programs National Institute of Standards and Technology 100 Bureau Drive, Mail Stop 1000 Gaithersburg, MD 20899-1000 An appeal of an initial denial must include: 1. The requestor's name, current home or business address, and telephone number or e-mail address (to ensure timely communication); 2. A copy of the original request and any correspondence regarding the initial denial; and 3. A statement of the reasons why the requester believes the initial denial was in error. Additional details regarding the appeal process can be found in Section D. of the NIST Guidelines.

I'm now appealing.

17 Formalities regarding the appeal

NIST's information-quality standards impose the following requirements upon an appeal.

An appeal "must be made within 30 calendar days of the date of the initial decision". This appeal is within this time limit.

An appeal must include "the requester's name, current home or business address, and telephone number or e-mail address (in order to ensure timely communication)". My name is Daniel J. Bernstein. My address is Department of Computer Science, 851 S. Morgan St, MC 152, Chicago, IL, 60607. My email address is djb@math.uic.edu. Given the slowness and unreliability of postal services (not to mention a move of the department this year to a new building), I request the usage of email for all communication regarding this matter.

An appeal "must be in writing and addressed to" the "NIST Associate Director for Laboratory Programs". At this point there is an ambiguity: is email allowed? An argument for "no" would be that no email address is provided. An argument for "yes" would be that excluding email is contrary to the "ensure timely communication" quote. I contacted the NIST Associate Director for Laboratory Programs by email and was happy to receive a prompt response saying "An email is fine to send the appeal", resolving this ambiguity.

An appeal must include "a copy of the original request and any correspondence regarding the initial denial". This correspondence appears in chronological order in the remaining sections of this document.

Finally, an appeal must include "a statement of the reasons why the requester believes the initial denial was in error".

Procedurally, NIST's denial violated various regulatory requirements such as the requirement to "contain a point-by-point response to any relevant data quality arguments contained in the request". NIST's unexplained delays also violated the regulations and exacerbated the damage. More broadly, NIST's obstructionism is the opposite of providing due process.

Content-wise, NIST is obliged by Public Law 106-554 to provide "correction of information maintained and disseminated by the agency that does not comply with the guidelines issued under subsection (a)". The guidelines that NIST issued under this law are NIST's information-quality standards. Table 5 of NIST IR

8545 is information maintained and disseminated by the agency. Earlier sections of this document state in detail why that table is not just embarrassingly inaccurate but also violates quite a few of NIST's information-quality standards. Ergo, NIST is obliged by law to correct the table.

18 Correction email dated 7 May 2025 17:39:21 +0200

Subject: Section 515 of Public Law 106-554
From: "D. J. Bernstein" <djb@math.uic.edu>
Date: Wed, 7 May 2025 17:39:21 +0200
To: info.quality@nist.gov
Message-ID: <20250507153921.1826142.qmail@cr.yp.to>

Dear Director, NIST Management and Organization Office:

1. I am writing to you to submit a request under Section 515 of Public Law 106-554 for correction of information, as per the following procedures:

> https://web.archive.org/web/20250507112244/https://www.nist.gov/director/nist-information-quality-standards

My name is Daniel J. Bernstein. My address is Department of Computer Science, 851 S. Morgan St, MC 152, Chicago, IL, 60607. My email address is djb@math.uic.edu. Given the slowness and unreliability of postal services, please use email for all communication regarding this matter.

2. The particular information at issue is Table 5 in this document:

> https://web.archive.org/web/20250313083030/https://nvlpubs.nist.gov/nistpubs/ir/2025/NIST.IR.8545.pdf

This document, NIST IR 8545, was issued by NIST in March 2025 and is also currently available on the NIST web site. It was issued within NIST's Post-Quantum Cryptography Standardization Process, and lists pqc-comments@nist.gov as a contact address.

The aforementioned NIST project solicited four rounds of submissions, including software, with due dates in 2017, 2019, 2020, and 2022. There were various changes to the submissions, described in the submission documents and sometimes also as intermediate announcements on NIST's "pqc-forum" mailing list. For example,

> https://web.archive.org/web/20250325060224/https://csrc.nist.gov/CSRC/media/Projects/post-quantum-cryptography/documents/round-2/official-comments/Classic-McEliece-round2-official-comment.pdf

shows that NIST received a public "OFFICIAL COMMENT" from Tung Chou in 2020 (during the second round) announcing a big software speedup for Classic McEliece, one of the submissions.

NIST IR 8545 says it evaluates the fourth-round submissions. In particular, Table 5 of NIST IR 8545 claims to be evaluating a cost metric ("cycles") for the Classic McEliece submission.

Now here's the problem in a nutshell: the number "114 189" in Table 5 of NIST IR 8545 is easily verified to be wrong, overstating cost in this metric by approximately a factor of 4.

More broadly, the numbers in that table are actually evaluations of the second-round Classic McEliece software. NIST is falsely portraying these numbers as evaluations of the fourth-round Classic McEliece software.

The source cited by NIST for the table already had a prominent warning in April 2024 that it was giving out-of-date numbers from a defunct measurement project:

> https://web.archive.org/web/20240425050637/https://openquantumsafe.org/benchmarking/

NIST's bibliography entry violates the standard requirement of stating an access date. When I publicly challenged Table 5 (on a "pqc-forum" mailing list operated by NIST), did NIST claim that, oops, it had visited the URL before the warning appeared, and hadn't rechecked for an entire year before issuing NIST IR 8545, and had somehow misled itself into believing that the data was for the round-4 software?

No. Instead NIST claimed "We have verified that the performance figures given in the tables in our report accurately reflect the benchmarks we cited":

> https://web.archive.org/web/20250507151442/https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/CSYRn8A-0zE/m/1PPIazqcAAAJ

This is absurd. The cited source doesn't claim to be, and isn't, measuring the round-4 submission. The claim that the numbers are measurements of the round-4 submission is a fabrication by NIST.

Along with claiming accuracy, NIST threw up a smokescreen, writing that "performance can vary significantly depending on factors such as platform, application, computing environment, and network conditions". Those variations have nothing to do with the issue at hand.

NIST IR 8545 claims "114 189" kilocycles for mceliece348864f key generation on "x86_64". A closer look at "[1]" shows that this is from an Intel Xeon Platinum 8259CL CPU. The submitted fourth-round software is 4x faster than that on that CPU.

The network conditions and application are irrelevant. The problem is NIST claiming measurements of the fourth-round candidate when in fact NIST is talking about code from 2019, ignoring subsequent speedups.

3. Your procedures, beyond simply asking "why the requester believes that the information is not correct", also ask "how the information at issue fails to comply with applicable information quality guidelines and

standards".

NIST IR 8545 does not comply with the following NIST information-quality standards. I am not saying that this list is complete, but this list is clearly sufficient for invoking Section 515 of Public Law 106-554.

3(a). "Quality will be ensured and established at levels appropriate to the nature and timeliness of the information to be disseminated."

As background, attackers are recording data encrypted today to decrypt with future quantum computers. Post-quantum cryptography is a multi-billion-dollar industry reacting to this security problem.

Various existing NIST standards have the same security problem. In 2016, NIST issued 81 FR 92787, establishing a multi-year process to evaluate and standardize replacements, and citing criteria that "will be used" for that process. Those criteria include a statement that the goal of the process is to "select a number of acceptable candidate cryptosystems for standardization":

> https://web.archive.org/web/20220119113311/https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf

Evaluation of the acceptability of post-quantum cryptosystems is thus an important and urgent task.

NIST IR 8545 says that its purpose is to summarize "the fourth round of the NIST PQC Standardization Process". Unfortunately, in preparing the report, NIST recklessly collected round-2 information and falsely presented it as round-4 information, while ignoring and suppressing warnings about the information being out of date. This was neither ensuring nor establishing an appropriate level of quality.

Note that "quality" is defined to include "objectivity", which includes "a focus on ensuring accurate, reliable, and unbiased information", along with a procedural requirement of "sound statistical and research methods". What NIST did here was not sound research, and did not take reasonable steps to ensure accuracy.

3(b). "Although third-party sources may not be directly subject to Section 515, information from such sources, when used by NIST to develop information products or to form the basis of a decision or policy must be of known quality and consistent with the applicable information quality guidelines and standards. When such information is used, any limitations, assumptions, collection methods, or uncertainties concerning it are taken into account and disclosed."

NIST IR 8545 violated this "disclosed" rule by repeating numbers from a non-governmental source while suppressing the warning prominently provided by that source.

NIST IR 8545's failure to state an access date for the reference was

another violation of this rule.

The "known quality" rule means that NIST cannot rely on ignorance as an excuse for repeating bad information. For example, if NIST suddenly decides to switch to claiming that it _assumed_ that these third-party numbers were measurements of the round-4 software: "known quality" means that NIST wasn't allowed to make any such assumption.

3(c). "Scientific and financial information disseminated by NIST is reliable and accurate to an acceptable degree of error as determined by factors such as the importance of the information, its intended use, time sensitivity, expected degree of permanence, relation to the primary mission(s) of the disseminating office, and the context of the dissemination, balanced against the resources required and the time available."

See above regarding importance and time sensitivity. NIST IR 8545 is a permanent report, a prominent part of a multi-year multi-person project by the disseminating office at NIST, a project that has been repeatedly highlighted in activity reports by that office and that plays an indispensable role in White House orders regarding the topic.

NIST IR 8545 does not quantify the weights placed upon individual comparison factors, such as the specific table at issue; but it would not be credible for NIST to retroactively claim that Table 5 is an unimportant part of NIST IR 8545. The tables are prominently placed (this table is on page 6 of a 34-page document), appear to be the source of various comments in the text, and are the central cost measurements in the report. Cost is the second main evaluation criterion stated in NIST's evaluation criteria for the project:

> https://web.archive.org/web/20220119113311/https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf

NIST is obliged to follow those criteria by 81 FR 92787 unless and until that is overridden. More to the point, there is ample evidence of cost measurements playing a major role in cryptographic decisions.

Regarding "intended use", NIST appears to have now terminated activity under the name "Post-Quantum Standardization Process" (perhaps because of my ongoing FOIA lawsuits regarding that process). However, the stated goal to "select a number of acceptable candidate cryptosystems for standardization" has not disappeared. On the contrary, NIST and other standardization agencies are continuing to work on post-quantum standardization, as illustrated by

> https://web.archive.org/web/20250415051610/https://csrc.nist.gov/Projects/pqc-dig-sig/

and by the following statement from NIST IR 8545: "After the ISO standardization process has been completed, NIST may consider developing a standard for Classic McEliece based on the ISO standard."

So it would not be credible for NIST to claim that the "intended use" of

its evaluations of candidates is merely to document a past decision. Furthermore, such a claim cannot retroactively justify NIST's recklessness in handling this data when it was making that decision. NIST was violating the rules at that point, not taking reasonable steps to ensure accuracy.

Regarding the resources required to obtain the correct information, the correct numbers were stated in a table in

    https://classic.mceliece.org/mceliece-impl-20221023.pdf#page.10

which was part of the fourth-round submission package filed with NIST. Those numbers are for Intel Haswell CPUs; NIST had designated Haswell as its comparison platform in

    https://web.archive.org/web/20220120140641/https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/BjLtcwXALbA/m/Bjj_77pzCAAJ

in 2019 and never announced a revised designation, although the NIST IR 8545 table at issue actually switches to newer Intel CPUs with slightly better numbers. Continually updated measurements for many CPUs are available from

    https://bench.cr.yp.to

which is a long-running measurement project that I co-initiated and continue to manage. Verifying the accuracy of the numbers is a simple matter of re-running the software.

3(d). "NIST Policy 1800.00 requires that NIST measurement results be accompanied by quantitative statements of their uncertainty, and that a uniform approach to expressing measurement uncertainty be followed."

The tables in NIST IR 8545 fail to include this information. For comparison, the numbers from

    https://classic.mceliece.org/mceliece-impl-20221023.pdf#page.10

include medians and quartiles, as do the https://bench.cr.yp.to tables.

3(e). "NIST treats information quality as integral to every step in its process of developing the information, including creation, collection, maintenance, and dissemination. All scientific information disseminated by NIST receives a level of scrutiny commensurate with the critical nature of the information and its intended use."

Obviously this scrutiny did not take place here. NIST copied and pasted third-party round-2 numbers and falsely presented those as round-4 numbers.

4. Your procedures also require "an explanation of how the requester is

affected". By definition, "affected person" means "an individual or entity that uses, benefits from, or is harmed by the disseminated information at issue".

To answer this, I'll start by quoting the Policy Statement on Ethical Guidelines from the American Mathematical Society: "The correct attribution of mathematical results is essential, both because it encourages creativity, by benefiting the creator whose career may depend on the recognition of the work and because it informs the community of when, where, and sometimes how original ideas entered into the chain of mathematical thought." Specific responsibilities include giving "appropriate credit" and using "no language that suppresses or improperly detracts from the work of others":

https://www.ams.org/about-us/governance/policy-statements/sec-ethics

I am speaking for myself in this message, but I am also one of the members of the team that submitted Classic McEliece to NIST. What NIST is doing in the table at issue is exaggerating the cost, specifically the CPU time, of Classic McEliece. NIST is improperly detracting from the work of all of the scientists who contributed to Classic McEliece.

Furthermore, NIST's citation and usage of (obsolete) results from https://openquantumsafe.org/benchmarking/ fails to give appropriate credit to https://bench.cr.yp.to, a project that had the relevant benchmarks before anyone else, along with careful attention to accuracy, verifiability, comprehensiveness, etc.

Beyond my role as a scientist working in this area, I am like billions of other people in relying on cryptography to protect data in transit through the Internet. Classic McEliece has the strongest track record of all available options for post-quantum public-key encryption, and it is also the most efficient option for a large class of applications (what are called "static keys" in, e.g. NIST SP 800-56 and 800-57), contrary to what NIST IR 8545 indicates. I don't expect NIST's misinformation about Classic McEliece to be able to stop Classic McEliece deployment, but I do expect it to slow down deployment if it is not promptly corrected. Ultimately NIST is damaging security, damaging efficiency, and harming the general public.

5. Your procedures also say that you will not consider corrections that "would serve no useful purpose", such as corrections to information that is "not valid, used, or useful after a stated short period of time".

NIST IR 8545 does not state a time limit on the usage of its information regarding Classic McEliece. On the contrary, as noted above, there is ongoing activity by NIST and others regarding this. Consider again the following statement in NIST IR 8545: "After the ISO standardization process has been completed, NIST may consider developing a standard for Classic McEliece based on the ISO standard."

You can see from https://mceliece.org that there are already many users of Classic McEliece, and further users who have expressed interest in

Classic McEliece. Cryptographic decisions are happening all the time.
There is no excuse for NIST misrepresenting what is known about the
features of the options that are available.

6. I believe that a suitable form of correction would be an update to
NIST IR 8545 to correct Table 5, accompanied by an apology for the
indefensible previous misinformation from NIST.

Please let me know if you need any further information.

---D. J. Bernstein

19 NIST email dated 8 May 2025 18:43:23 +0000

Subject: RE: Section 515 of Public Law 106-554
From: "Egan, Amy S. (Fed)" <amy.egan@nist.gov>
Date: Thu, 8 May 2025 18:43:23 +0000
To: "djb@math.uic.edu" <djb@math.uic.edu>
Message-ID: <SJ0PR09MB10422042C7A04A5C74A3088ECFC8BA@SJ0PR09MB10422.namprd09.prod.outlook
.com>

This email acknowledges receipt of your May 7, 2025 Information Quality Request to the
National Institute of Standards and Technology.

We are reviewing your request and the information provided. We will let you know if we
have any questions or need additional information.

Thank you,

Amy Egan
Acting Director, Management and Organization Office
National Institute of Standards and Technology
301-975-2819

20 Correction email dated 6 Jun 2025 13:24:04 +0200

Subject: Re: Section 515 of Public Law 106-554
From: "D. J. Bernstein" <djb@math.uic.edu>
Date: Fri, 6 Jun 2025 13:24:04 +0200
To: "Egan, Amy S. (Fed)" <amy.egan@nist.gov>
Message-ID: <20250606112404.3975134.qmail@cr.yp.to>

Dear Acting Director, NIST Management and Organization Office:

I am writing to check on the status of NIST's correction procedure
regarding this matter. Thanks in advance for any information you can
provide.

---D. J. Bernstein

21 NIST email dated 6 Jun 2025 12:11:48 +0000

Subject: RE: Section 515 of Public Law 106-554
From: "Egan, Amy S. (Fed)" <amy.egan@nist.gov>
Date: Fri, 6 Jun 2025 12:11:48 +0000
To: "D. J. Bernstein" <djb@math.uic.edu>
Message-ID: <SJ0PR09MB10422A27CC90098085FA8ED8CFC6EA@SJ0PR09MB10422.namprd09.prod.outlook.com>

Thank you for writing regarding the status of your information quality request.

NIST is reviewing your request and working on the determination. As stated in the information quality procedures (available at https://www.nist.gov/director/nist-information-quality-standards), we will issue a determination usually within 60 calendar days of receipt of the request.

We will contact you if we have any questions or need additional information regarding your request.

Sincerely,

Amy Egan
Acting Director, Management and Organization Office
National Institute of Standards and Technology
301-975-2819

22 NIST email dated 7 Jul 2025 15:09:20 +0000

Subject: RE: Section 515 of Public Law 106-554
From: "Egan, Amy S. (Fed)" <amy.egan@nist.gov>
Date: Mon, 7 Jul 2025 15:09:20 +0000
To: "D. J. Bernstein" <djb@math.uic.edu>
Message-ID: <SJ0PR09MB10422BFC48D93F4439A7687E7FC4FA@SJ0PR09MB10422.namprd09.prod.outlook.com>

I am writing to provide you with an update on the response to your request for correction under NIST's Information Quality Standard Guidelines. The response is currently undergoing review, and we will send it to you as soon as the review process is concluded.

Thank you,

Amy Egan
Acting Director, Management and Organization Office
National Institute of Standards and Technology
301-975-2819

23 Correction email dated 7 Jul 2025 17:37:26 +0200

Subject: Re: Section 515 of Public Law 106-554
From: "D. J. Bernstein" <djb@math.uic.edu>

Date: Mon, 7 Jul 2025 17:37:26 +0200
To: "Egan, Amy S. (Fed)" <amy.egan@nist.gov>
Message-ID: <20250707153726.1205345.qmail@cr.yp.to>

Dear Acting Director, NIST Management and Organization Office:

Thank you for the update. Do you have an estimated decision date?

---D. J. Bernstein

24 NIST email dated 7 Jul 2025 17:15:31 +0000

Subject: RE: [EXTERNAL] Re: Section 515 of Public Law 106-554
From: "Egan, Amy S. (Fed)" <amy.egan@nist.gov>
Date: Mon, 7 Jul 2025 17:15:31 +0000
To: "D. J. Bernstein" <djb@math.uic.edu>
Message-ID: <SJ0PR09MB10422DABB8BBB71AD319DF50AFC4FA@SJ0PR09MB10422.namprd09.prod.outlook.com>

We are required to forward all responses to requests for correction to OMB for review, those reviews are coordinated by the Department of Commerce, and as such, we are unable to provide an expected completion date.

Thank you,

Amy Egan
Acting Director, Management and Organization Office
National Institute of Standards and Technology
301-975-2819

25 Correction email dated 7 Jul 2025 19:45:33 +0200

Subject: Re: [EXTERNAL] Re: Section 515 of Public Law 106-554
From: "D. J. Bernstein" <djb@math.uic.edu>
Date: Mon, 7 Jul 2025 19:45:33 +0200
To: "Egan, Amy S. (Fed)" <amy.egan@nist.gov>
Message-ID: <20250707174534.1211653.qmail@cr.yp.to>

Dear Acting Director, NIST Management and Organization Office:

Thanks for your message. One reason I'm asking is that

> https://web.archive.org/web/20250507112244/https://www.nist.gov/director/nist-information-quality-standards

includes the following policy: "The Director, NIST Management and Organization Office will attempt to communicate either a decision on the request, or a statement of the status of the request and an estimated decision date, within 60 calendar days after receipt of the request."

I understand that "attempt" isn't a guarantee, but failing to reach the 60-day target and failing to provide an estimated decision date would

seem to require justification specific to the circumstances at hand,
rather than a generic statement about OMB involvement in all responses.

I find it particularly puzzling to not see a prompt correction for an
amply documented case of NIST fabricating data. Please say concretely
what communication has occurred regarding this specific matter and why
this is taking so long.

---D. J. Bernstein

26 NIST email dated 4 Aug 2025 20:01:45 +0000

Subject: RE: Section 515 of Public Law 106-554
From: "Egan, Amy S. (Fed)" <amy.egan@nist.gov>
Date: Mon, 4 Aug 2025 20:01:45 +0000
To: "djb@math.uic.edu" <djb@math.uic.edu>
Message-ID: <SJ0PR09MB104228509F8247746A7647632FC23A@SJ0PR09MB10422.namprd09.prod.outlook
.com>

Please see attached response to your May 7, 2025, Information Quality Request to NIST.

Sincerely,

Amy Egan
Acting Director, Management and Organization Office
National Institute of Standards and Technology
301-975-2819

27 NIST attachment

See subsequent pages.

UNITED STATES DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
Gaithersburg, Maryland 20899-0001

August 4, 2025

Daniel J. Bernstein
Department of Computer Science
University of Illinois Chicago
851 S. Morgan St., MC 152
Chicago, IL 60607
Via email: djb@math.uic.edu

Dear Dr. Bernstein:

This letter is in response to your May 7, 2025 request for correction (Request) of certain information contained in the NIST IR 8545, Status Report on the Fourth Rounds of the NIST Post-Quantum Cryptography Standardization Process[1]. Your Request is submitted under the Information Quality Act, Section 515 of Public Law 106-554; the Office of Management Budget's Guidelines for Ensuring and Maximizing the Quality, Objectivity, Utility, and Integrity of Information Disseminated by Federal Agencies (67 FR 8451, Feb. 22, 2002) ("OMB Guidelines"); the Office of Management and Budget's Memorandum on Improving Implementation of the Information Quality Act, M-19-15 ("OMB Memo"); and NIST's Guidelines, Information Quality Standards, and Administrative Mechanism ("NIST Guidelines").

The Request asserts that the following information in NIST IR 8545 requires correction:

1. That the number "114 189" in Table 5 is incorrect, overstating the number of clock cycles by approximately a factor of 4.
2. That this number incorrectly reflects performance of the second-round software for Classic McEliece, not the fourth-round implementation.
3. That NIST cited a defunct and unmaintained source without disclosing the limitations and omitted access dates for the citation.
4. That NIST disregarded better benchmarking sources, such as the SUPERCOP/bench.cr.yp.to platform, which is more accurate, current, and comprehensive.

You cite previous correspondence regarding the Table 5 number on a "pqc-forum" mailing list[2].

[1] Available at: https://nvlpubs.nist.gov/nistpubs/ir/2025/NIST.IR.8545.pdf
[2] https://groups.google.com/u/1/a/list.nist.gov/g/pqc-forum/c/CSYRn8A-0zE/m/6rhWGzEgDgAJ

According to the Request, the supposed error misinformed the cryptographic community and delayed deployment of Classic McEliece, thereby impacting efficiency and public security.

NIST IRs are part of the NIST Technical Series Publication and defined as interagency or internal reports of research findings, including background information for Federal Information Processing Standards (FIPS) and NIST Special Publications (NIST SP)[3]. NIST treats NIST IRs as fundamental research communication. As stated in the NIST Guidelines, while fundamental research communications are expected to adhere to information aquality standards and do receive technical, policy, and editorial review by the NIST Editorial Review Board in addition to the pre-dissemination review methods listed in the NIST Guidelines, they are not included in the scope of the OMB and NIST Guidelines. NIST IR 8545 was approved by the NIST Editorial Review Board on March 5, 2025.

Even if, assuming *arguendo*, NIST IR 8545 were to be included within the scope of the NIST Guidelines, your requests for correction would not be warranted, as discussed below.

1. ***Request for Correction:*** *The number "114 189" in Table 5 is incorrect, overstating the number of clock cycles by approximately a factor of 4.*

   NIST recognizes that different benchmarking platforms have reported varying results using different implementations of the Classic McEliece algorithm, with some reporting lower cycle counts than those listed in Table 5. However, this figure in Table 5 was accurately cited from a publicly available benchmark source—specifically, Reference [1] in NIST IR 8545, which corresponds to the Open Quantum Safe (OQS) project[4]. NIST IR 8545 accurately reproduces the benchmark results published there. Importantly, only one out of 48 data points drawn from Reference [1] across Tables 3, 4, and 5 has been challenged, and even this challenged data point was correctly cited. The allegation that the report fabricated or misstated data is unfounded. As stated in NIST IR 8545, the performance figures are intended to be representative, not definitive, and reasonable differences in benchmark platforms or methodologies do not constitute misinformation.

2. ***Request for Correction:*** *This number reflects the performance of the second-round software for Classic McEliece, not the fourth-round software implementation.*

   NIST disagrees with this assertion. NIST is confident that the benchmarks cited from Reference [1] accurately reflect the fourth-round version of Classic McEliece. NIST IR 8545 focuses specifically on fourth-round submissions, and the benchmark data

[3]https://csrc.nist.gov/publications#:~:text=NIST%20Interagency%20or%20Internal%20Reports%20(NIST%20IR),NIST%20Cybersecurity%20and%20Privacy%20Program.
[4] Open Quantum Safe (OQS) algorithm performance visualizations. Available at https://openquantumsafe.org/benchmarking

used aligns with that scope. Accordingly, the performance measurement in question is consistent with the version under evaluation.

3. ***Request for Correction:*** *NIST cited a defunct and unmaintained source without disclosing the limitations and omitted access dates for the cited source.*

   The cited source— Open Quantum Safe algorithm performance visualizations — was an active and widely used project at the time NIST gathered benchmark data. It played a visible and constructive role in the NIST post-quantum cryptography process. The disclaimer currently shown on the OQS website was added after NIST collected the data. NIST accurately cited the relevant information as it existed at the time. The addition of a later disclaimer does not retroactively invalidate the benchmark or its citation.

4. ***Request for Correction:*** *NIST disregarded better benchmarking sources, such as the SUPERCOP/bench.cr.yp.to platform, which is more accurate, current, and comprehensive.*

   Benchmarking results naturally vary depending on the platform, implementation, and methodology. As stated in the report, the benchmark tables were intended to provide representative data points, not authoritative or exhaustive metrics. NIST did not disregard SUPERCOP; it was among the sources NIST reviewed during the evaluation process. Ultimately, NIST selected Reference [1] because it provided a consistently structured and well-documented dataset that served NIST's goal of presenting a fair, representative comparison across submissions. NIST recognizes that experts may have different preferences for benchmarking sources, but choosing one over another for clarity and consistency does not imply negligence or bias.

Regarding the alleged harm, we emphasize that key generation cycle counts are only one of many considerations for weighing alternatives for standardization. Suitability and competitiveness of a given algorithm for standardization depend on a variety of security, performance, and implementation characteristics across different use cases. Improved key generation benchmarks alone would not change the conclusions.

Based on the foregoing, your Request is denied. In preparing this response, NIST carefully reviewed the relevant information and sources cited to ensure that its conclusions are accurate, well-founded, and consistent with NIST's information quality standards.

You may appeal this decision within 30 calendar days from the date of this decision. Such an appeal must be in writing and addressed to:

NIST Associate Director for Laboratory Programs
National Institute of Standards and Technology

100 Bureau Drive, Mail Stop 1000
Gaithersburg, MD 20899-1000

An appeal of an initial denial must include:

1. The requestor's name, current home or business address, and telephone number or e-mail address (to ensure timely communication);
2. A copy of the original request and any correspondence regarding the initial denial; and
3. A statement of the reasons why the requester believes the initial denial was in error.

Additional details regarding the appeal process can be found in Section D. of the NIST Guidelines.

Sincerely,

Amy Egan
Acting Director, Management and Organization Office