

---

**From:** D. J. Bernstein <djb@cr.yp.to>  
**Sent:** Monday, September 21, 2020 5:50 AM  
**To:** pqc-comments  
**Cc:** pqc-forum  
**Subject:** ROUND 2 OFFICIAL COMMENT: NTRU Prime  
**Attachments:** signature.asc

I'm writing on behalf of the NTRU Prime team, as a public response to NIST's request for a summary of expected changes in round 3.

NTRU Prime is a small lattice system. Subject to this constraint, our primary objective is to eliminate unnecessary complications in security review. We correctly predicted that such complications would lead to security failures in NISTPQC lattice submissions. We evaluated a variety of trapdoor functions from this perspective before submission, again during round 1, and again during round 2.

On this basis we have once again decided against decryption failures; modules; errors; and all other changes that we have considered to our family of trapdoor functions. We therefore plan to submit the same family of trapdoor functions in round 3. NTRU Prime will therefore have an unchanged family of trapdoor functions throughout round 1, round 2, and round 3.

Our CCA conversion includes various hashing safeguards, some already in round 1 and some added in round 2. These safeguards cost 32 bytes in ciphertext size and a considerable fraction of our CPU time. However, even with these safeguards, NTRU Prime often outperforms other small lattice KEMs, as the following references show:

<https://cr.yp.to/papers.html#paretoviz>  
<https://cr.yp.to/talks/2019.08.23-1/slides-djb-20190823-1-paretoviz-4x3.pdf>  
<https://bench.cr.yp.to/results-kem.html#amd64-hiphop>  
<https://github.com/mupq/pqm4/blob/master/benchmarks.md>

More importantly, the costs of our hashing safeguards are negligible in applications. We plan to submit the same CCA conversion in round 3.

NTRU Prime will thus be fully compatible between round 2 and round 3, when users choose the same parameters.

Regarding parameter selection, we are concerned that pre-quantum Core-SVP levels  $2^{100}$ ,  $2^{106}$ , and  $2^{111}$ , proposed for category 1 for Dilithium, NTRU, and Kyber respectively, will turn out to be inadequate against generic lattice attacks. We will not add dimensions below our 653 (pre-quantum Core-SVP  $2^{129}$ ). We recommend our original dimension 761 (pre-quantum Core-SVP  $2^{153}$ ) for an extra security margin.

We have seen various requests for larger dimensions, even larger than our dimension 857 (pre-quantum Core-SVP  $2^{175}$ ). To accommodate these requests and prevent any accusations of a lack of flexibility, we plan to add some larger dimensions as a supplement to our current dimensions.

We have also considered adding intermediate parameter sets to further illustrate our flexibility, showing that NTRU Prime offers even larger advantages in Core-SVP under various size limits compared to, e.g., Kyber. The call for proposals explicitly allowed multiple parameter sets per category. However, NIST has now made an announcement that seems to discourage "too many parameter sets".

The rest of this message is regarding the problems caused by NIST's unstable definitions of security categories. These are problems for the NISTPQC process broadly, not just for NTRU Prime.

The call for proposals specified AES-128 key search as a "floor" for category 1 in "all metrics that NIST deems to be potentially relevant to practical security". The call for proposals similarly specified floors for other categories. However, this is not a clear and stable category definition unless the metrics are clear and stable.

As an illustration of how much impact metrics have, there is a 40-year literature studying metrics for realistic large-scale two-dimensional models of computation. Standard theorems---see, e.g.,

<https://www.eecs.harvard.edu/~htk/publication/1981-jacm-brent-kung.pdf>

---imply that these metrics assign 50% higher asymptotic exponents to large-integer multiplication, large-array sorting, etc. than a "gates"

metric does. (Analogous three-dimensional metrics studied in, e.g.,

<https://link.springer.com/article/10.1007/BF01744565>

are for machines that appear far more difficult to build than quantum computers, and still have 33% higher exponents than a "gates" metric.) Any attack that has large-scale sorting as a bottleneck is affected by this, whereas AES-128 key search is not.

The call for proposals highlighted "classical gates" and "quantum gates" (with limited depth) as metrics. However, NIST is not requiring lattice submissions to meet the "classical gates" floor. (See examples below.)

NIST also has not defined a replacement metric for submissions to use.

All lattice submissions have Core-SVP evaluations, but AES-128 does not.

Core-SVP is not a metric for the cost of computation: it is a mechanism for claiming security levels (in an undefined metric) specifically for lattices. Ray Perlner's message dated 9 Jun 2020 15:39:09 +0000 stated "we feel that the CoreSVP metric does indicate which lattice schemes are being more and less aggressive in setting their parameters", but the mapping from Core-SVP evaluations to categories remains undefined.

NIST IR 8309's handling of categories is not consistent across lattice submissions. Consider the following examples from three different submissions:

(P80) Category 3 for pre-quantum Core-SVP  $2^{153}$ ; 153 is 80% of 192.

(P78) Category 1 for pre-quantum Core-SVP  $2^{100}$ ; 100 is 78% of 128.

(P71) Category 3 for pre-quantum Core-SVP  $2^{136}$ ; 136 is 71% of 192.

P78 is the lowest of these three examples in Core-SVP, and P71 is the lowest in Core-SVP as a percentage of the AES key size for the category.

However, NIST's wording was strikingly more negative for P80 than for P78 and P71:

(P80) "quite aggressive compared to most of the other submissions targeting the same security categories"; need to study "whether they actually meet their claimed security categories";

(P78) "lowest CoreSVP security strength parameter set of any of the lattice schemes still in the process"; need more study on "understanding the concrete security";

(P71) "lower CoreSVP complexity than many of the other schemes targeting the same security strength categories"; need to "understand exactly ... bit security strengths".

Notice, e.g., that NIST asks whether P80 "actually" meets its "claimed" security category, while NIST does not ask the same question regarding P78 or P71.

If NIST were applying the "classical gates" metric then none of P80, P78, and P71 would be able to confidently claim these categories. For example, the uncertainties in Core-SVP seem very unlikely to turn Core-SVP  $2^{100}$  into  $2^{143}$  "classical gates". Most of the remaining lattice submissions (at least Dilithium, NTRU, Kyber, and NTRU Prime; perhaps also SABER after the announcement that SABER's security levels were miscalculated) would have to adjust their category assignments.

Even worse, some of these submissions (at least Dilithium, NTRU, and Kyber) would have to remove some previously proposed parameters, which seems contrary to the idea of being ready for standardization.

All of these submissions argue, with varying levels of detail and references, that the "classical gates" metric underestimates the actual cost of known attacks. NIST seems receptive to the idea of using a more realistic metric, but has taken four years to post its "preliminary thoughts" on the realism of several different metrics. It is not clear what metrics NIST will end up defining, and it is not clear how long NIST will take to settle on the definitions. What is clear is that NIST has not applied the categories consistently, as illustrated by NIST IR 8309 assigning P80 more negative wording than P78 and P71.

The different wording regarding P80, P78, and P71 appears to have translated into different action, and this seems particularly important for NIST's handling of NTRU Prime. As context, NIST IR 8309 describes finalists in general as

the most promising to fit the majority of use cases and most likely to be ready for standardization soon after the end of the third round.

We have shown that NTRU Prime fits practically all use cases. As far as we can tell, beyond general concerns about the safety of lattice-based cryptography and about the safety of all small lattice proposals, NTRU Prime is ready for standardization now with our existing parameter sets.

The only negative comments that NIST IR 8309 made regarding NTRU Prime were regarding parameter sets. Specifically, NIST seemed to criticize

- \* NTRU Prime's assignment of pre-quantum Core-SVP  $2^{153}$  to Category 3 (this is exactly P80 above),

- \* NTRU Prime's assignment of \_post-quantum\_ Core-SVP  $2^{159}$  (pre-quantum Core-SVP  $2^{175}$ ) to Category 4 (this has a larger security margin than P80),

- \* NTRU Prime's assignment of \_post-quantum\_ Core-SVP  $2^{117}$  (pre-quantum Core-SVP  $2^{129}$ ) to Category 2 (this has a larger security margin than P80), and

\* NTRU Prime's "narrower range of CoreSVP values" (our understanding now is that this wasn't a negative comment but merely a request for larger parameters going forward).

Meanwhile various lattice submissions with objectively more dangerous parameter selections were given less critical wording by NIST and were selected as finalists. We see no explanation for why NIST treated P78 and P71 in those submissions more gently than P80 in NTRU Prime.

An application limited to 1024 bytes for keys and plaintexts reaches Core-SVP  $2^{129}$  with NTRU Prime's proposed parameters and nothing better than  $2^{111}$  with Kyber's proposed parameters.  $2^{129}$  is higher security relative to category 2 than  $2^{111}$  relative to category 1, and obviously higher security on an absolute scale. NIST's report did not acknowledge this security advantage of NTRU Prime.

We are concerned that the lack of clear, stable, consistently applied category definitions will be used in the continuation of NISTPQC to again make NTRU Prime's parameter choices artificially sound worse than more dangerous parameter choices in other submissions. If we try to reduce this risk by downgrading (e.g.)  $2^{129}$  to category 1, while Kyber is allowed to remain in category 1 with just  $2^{111}$ , then NTRU Prime will be unfairly punished in performance comparisons.

We request that NIST issue clear and stable definitions of the metrics used to define NIST's security categories. At this point in the NISTPQC process, clarity and stability are more important than the exact level of realism. Beyond the floor for the categories, one can reasonably argue that users should take higher Core-SVP levels for all lattice submissions in light of continued advances in lattice attacks; but NIST should handle this in a way that is fair to all submissions. As soon as the evaluation criteria are made clear, we will be happy to adjust our category assignments accordingly.

---Dan

---

**From:** D. J. Bernstein <djb@cr.yp.to>  
**Sent:** Sunday, September 27, 2020 4:43 AM  
**To:** pqc-forum  
**Cc:** pqc-comments  
**Subject:** ROUND 2 OFFICIAL COMMENT: NTRU Prime  
**Attachments:** signature.asc

This message has three questions for NIST.

In <https://www.youtube.com/watch?v=CBGX1OMzN1o&t=37m55s> a few days ago, NIST stated "We're we're still uh have some some questions about NTRU Prime" but didn't elaborate. What are NIST's questions about NTRU Prime?

The late notification and lack of information are problematic. NIST has asked for round-3 tweaks by 1 October, which is just a few days from now. Did I miss some NIST publication listing NIST's questions?

I see only one part of NIST IR 8309 that can be understood as a question about NTRU Prime: namely, "whether they actually meet their claimed security categories will need to be determined" regarding the parameter choices. Meanwhile NIST did not ask the same question regarding other submissions that have objectively more dangerous parameter selections.

The NTRU Prime team email dated 21 Sep 2020 11:49:53 +0200 gave examples of this phenomenon. Procedurally, seeing the issue raised by NIST was a prerequisite for responding to it (and the difficulty of responding was exacerbated by the lack of clarity regarding NIST's "categories").

This section of the talk appeared to be presenting NIST's rationale for selecting lattice finalists and lattice alternates. It would thus seem that the existence of NIST's "questions" regarding NTRU Prime played a role in this. Why didn't NIST IR 8309 say this and provide the list of questions?

---Dan (speaking for myself)

---

**From:** Moody, Dustin (Fed)  
**Sent:** Monday, September 28, 2020 5:15 PM  
**To:** D. J. Bernstein; pqc-forum  
**Cc:** pqc-comments  
**Subject:** Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

Dear Dan,

Thanks for your last two official comments on NTRUprime. We'll try to address your questions.

As noted, in my talk at PQCrypto I mentioned that NIST had some questions regarding NTRUprime. In the chat which followed the session, we were asked the same question you asked, i.e. what questions do we have about NTRUprime? We answered that they mainly deal with algebraic cryptanalysis of lattice KEMS which exploit the structure of the ring that they choose, as well as what we wrote in our write-up of NTRUprime in NISTIR 8309. In the chat we also had a few more comments. You were active on the chat, so we assumed you saw our answers. I wasn't trying to bring up anything last-minute. Indeed, some of this has been discussed via email between NIST and NTRUprime back in August. We will expand a bit more below.

While your comment (from Sept 21) focuses on the parametrization of NTRUprime, this was not a major reason for assigning NTRUprime as an alternate rather than a finalist. In particular, we felt that given known facts regarding lattice cryptanalysis, the cyclotomic structures used by the lattice finalists are more researched, and should by default, be considered better understood. In our view, the submitters of NTRUprime have voiced concerns about cyclotomic lattices, but have not proposed any concrete attack against them or any strong argument that NTRUprime's ring choice would rule out a similar attack. We are open to the possibility that there is such an attack, as we are open to the possibility that NTRUprime's ring choice is susceptible to an attack from which cyclotomics are immune. But, without a strong argument that such an attack is ruled out by NTRUprime's ring choice (this seems to us difficult to provide without seeing the hypothetical attack first,) any attack on cyclotomics would undermine our confidence in structured lattices overall. As such, NTRUprime's most probable path to standardization involves two major research results.

- 1) An attack undermining NIST's confidence in the choice of algebraic structures used by Kyber, NTRU, and Saber
- 2) The establishment of a strong theoretical barrier for attacks to be extended to the NTRUprime ring.

As the timeline for this path to standardization is by necessity longer than likely paths to standardization for the other finalists, after much deliberation and debate we classified NTRUprime as an alternate. On rereading the report, I realize we did not make this reasoning as spelled-out as it could have been, and for that we apologize.

It should also be noted that the sheer number of structured lattice KEMs in the process always meant we would have to make hard cuts in choosing finalists and alternates. In contrast, as an alternate, NTRUprime is still very much in the process. The same can not be said for other very strong submissions in the structured lattice KEM category that were eliminated from consideration at the end of the 1st and 2nd round of our process.

The fact that there is a lot of discussion in the report about parameter choices and security levels is not because it has factored heavily into our selections thus far, but rather with the intent that, going forward, all submissions in the 3rd round will have the opportunity, through tweaks, to minimize the chances that we choose not to standardize a scheme that we would have, had they chosen other parameters. The issue that the range of parameters offered by the various lattice submissions is highly variable when assigned to NIST security categories was in fact raised on the pqc-forum late in our selection process by you in late June. This was after we had already settled on our list of finalists and alternates,

but before our 2nd round report was published. We agreed that it could cause problems going forward, so we highlighted it in our report. The wordings we used varied for substantive reasons (e.g. NTRU's choice to highlight two different computational cost models, and the fact that failing to meet category 1 would result in us discarding a parameter set, while failing to meet a higher level could simply mean changing how it is labeled in our standard.) But the wording also varied due to the personal idiosyncrasies of the writing styles the 13 different authors of the NIST report. If you think subtle differences in wording are any indication of how the NIST team as a whole will evaluate the candidates at the end of the 3rd round, 12 to 18 months from now, independent of any tweaks made by the submitters, and any subsequent analysis that may come to light, you are reading too much into the tone here.

Finally, you request further clarification than we have already given regarding how we intend to assign security levels. We feel we have given about as much as we can give without prejudging scientific questions to which we do not know the answer. We have already stated that gate count in the RAM model exceeding that of brute force AES key search or brute force SHA collision search as appropriate will be taken as strong a priori evidence of meeting any of the NIST categories. If you think more aggressive parameters than are justified via the RAM model are appropriate for meeting a given security target, you will need to make a strong argument that convinces us. We've given reasons why previous arguments have not completely convinced us, so you should take that into account. We think there is a real possibility such arguments could convince us -- we just haven't been completely convinced by any yet. Likewise, while variations in wording about how we raise questions about a particular parameter set are likely to be uninformative, we think we were fairly consistent in which lattice parameter sets we did and did not raise concerns about when claiming each security strength category. Parameter sets we didn't raise concerns about are probably ok for their claimed security levels, barring developments in cryptanalysis. If you think we've missed something in this regard, please let us know, because it likely indicates a relevant research result we were not aware of at the time of writing our report.

In any event, we think the public statements we've made on the forum and in our report are sufficient for any submission team working in good faith to determine what parameter sets will be uncontroversial, controversial and unacceptable for the claimed security levels given the current state of knowledge. Keep in mind that extra care is warranted for the lowest (and perhaps highest) security levels offered by a submission. For example, while a controversial assignment of category 1 runs the risk of the parameter set in question not being standardized, a controversial assignment of category 3 probably just runs the risk of the parameter set in question being downgraded to category 2, or at worst, category 1. Stakeholders who care about category 3 can then take our assignment into account when deciding which NIST standardized parameter set to use. A controversial assessment of category 5 runs the risk that the submission will not meet the needs of any users who actually want category 5. We reiterate what we've been saying since the beginning of the PQC process that, barring future cryptanalysis, category 1 parameters are probably enough to thwart any purely computational attack in the near to medium term future, and category 3 is almost certainly enough.

We do try to respond promptly to questions brought to us, but it does take a little bit of time for our team to discuss and draft a response. The updated specifications and implementations are due on October 1st. We've tried to be flexible throughout the process. Teams can always contact us to ask for a bit more time if they feel they need it.

Dustin and Ray  
The NIST PQC team

---

**From:** D. J. Bernstein <djb@cr.yp.to>  
**Sent:** Friday, September 3, 2021 4:53 PM  
**To:** pqc-comments  
**Cc:** pqc-forum  
**Subject:** ROUND 3 OFFICIAL COMMENT: NTRU Prime  
**Attachments:** signature.asc

Two years ago an NTRU Prime update talk announced the new "factored" NTRU Prime software, a wrapper around "modules with separate tests and optimizations". I gave a talk today announcing, among other things, computer verification for most of these modules that the existing "avx" implementation produces the same output as the existing "ref" implementation for `_all_` possible inputs:

<https://cr.yp.to/talks.html#2021.09.03>

This is a big step towards full verification of the optimized NTRU Prime software. Next steps are matching "avx" to "ref" for the other modules (notably multiplication, where another tool has gotten through some important parts of the code but not yet everything) and matching the C code to the Sage reference code.

The "saferewrite" tool used for this verification has a broad range of applicability beyond NTRU Prime. The first example in the talk is how saferewrite catches both of the array-comparison problems that were announced in the official Frodo software. However, to enable this analysis, I had to define a Frodo array-comparison module and write reference code for that module, so that saferewrite could compare the official code to my reference code. This wasn't a big deal since this particular module is so simple, but an analogous analysis for larger components of Frodo (short of taking the entire KEM as a monolith!) would require additional work to write reference code for those modules. For NTRU Prime, this work is already done.

For each module, saferewrite compiles each implementation with clang -O1 and with gcc -O3, uses the angr symbolic-execution toolkit to convert each binary into unrolled code in a much simpler language, and uses the Z3 theorem prover (via angr's claripy) to verify equivalence or find a counterexample. The automatic equivalence chains look like this (although this pattern isn't optimal in general):

```
opt clang -O1 = ref clang -O1 = avx clang -O1
  ||
opt gcc -O3 = ref gcc -O3 = avx gcc -O3
```

There could be compiler bugs affecting outputs, but to evade detection these bugs would have to have the same effect on every node in the diagram simultaneously. (It would also be possible to hook a direct Python-to-the-simpler-language conversion into the picture.) There could be unrolling bugs, but saferewrite also runs the binaries on some random inputs (plus all-0 and all-1 to make sure every bit is touched) and checks that these match the unrolled code; also, angr has been heavily exercised in a variety of reverse-engineering applications. There could be bugs in saferewrite itself, but reviewing saferewrite is a much smaller task than reviewing a ton of optimized post-quantum code.

The examples supplied in the saferewrite package include deliberately buggy code to exercise saferewrite's tests, in particular producing 16 analyses printing "differentfrom" counterexamples (which I've checked by hand), providing some evidence that saferewrite is working as desired.

Some of these bugs are also found by random tests, but some aren't. More advanced fuzzing can do better than random tests but has no hope of finding typical cryptographic overflow bugs.

Seeing C code working with two compilers doesn't mean that the same code will work with further compilers, but if analyses are fast enough then it's realistic to re-apply the analyses whenever the compiler changes. I tried the saferewrite analysis of 107 implementations of 27 functions, times 2 compilers, on a dual EPYC 7742; it finished in 8 minutes of wall-clock time, using 20 cores on average, using under 200GB of RAM.

I'm filing this as an OFFICIAL COMMENT regarding NTRU Prime because it's directly relevant to the official NISTPQC evaluation criteria, notably the following:

The algorithms can be implemented securely and efficiently on a wide variety of platforms, including constrained environments, such as smart cards.

Various NISTPQC submissions have provided fast AVX2 software, fast M4 software, etc., but the primary evidence for "securely" is that the software is constant-time. (I'll skip discussion of the broken masked implementations.) The problem is that the same optimizations add massive complexity to the software, and this complexity is a security threat:

\* <https://arxiv.org/abs/2107.04940> studied the vulnerabilities announced between 2010 and 2020 in eight well-known cryptographic libraries, and found 73 vulnerabilities in the cryptographic computations, including 11 known to be exploitable ("severe"), along with "evidence of a strong correlation between the complexity of these libraries and their (in)security". (There were also hundreds of further bugs, such as buffer overflows.)

\* Post-quantum software is newer, more complicated, and much harder to thoroughly review. Superficial reviews of post-quantum software have caught one devastating bug after another; the only reasonable prediction is that more serious reviews will find many more bugs.

Is it reasonable to say that an algorithm "can be implemented securely and efficiently" if fast implementations are so complex that the experts are getting them wrong? If the answer is pointing to an implementation and saying "No bugs are known in this implementation", then why should we think that the code is correct, rather than thinking that security reviewers are overloaded and that this answer is pure selection bias?

There's stronger evidence for "securely and efficiently" when optimized modules are verified to match much simpler reference implementations.

Covering more modules will further strengthen this evidence.

Another relevant NISTPQC evaluation criterion is the following:

Factors that might hinder or promote widespread adoption of an algorithm or implementation will be considered in the evaluation process, including, but not limited to, ...

The availability of modularized implementations, and the availability of verification tools applicable to some of those modules, certainly help promote widespread adoption of those implementations and the algorithm.

---Dan (speaking for myself)

---

**From:** pqc-forum@list.nist.gov on behalf of Wrenna Robson <wren.robson@gmail.com>  
**Sent:** Saturday, September 4, 2021 5:51 AM  
**To:** pqc-forum  
**Subject:** Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: NTRU Prime

Thank you for this, Dan - really interesting. I am hopeful of having something to share in the next couple of months on my own verification and validation efforts with Classic McEliece using SAW/Cryptol, and certainly much of your thinking here seems to align with conclusions and thoughts I've reached.

(Increasingly I think a hybrid approach, using multiple toolchains that complement each other's weaknesses, with a clear analysis and synthesis of the truth-claims that each makes. is the way forward - though this mean, of course, you have more tools that you need to validate and verify...)

- Wrenna

On Fri, 3 Sept 2021 at 21:53, D. J. Bernstein <[djb@cr.yp.to](mailto:djb@cr.yp.to)> wrote:

Two years ago an NTRU Prime update talk announced the new "factored" NTRU Prime software, a wrapper around "modules with separate tests and optimizations". I gave a talk today announcing, among other things, computer verification for most of these modules that the existing "avx" implementation produces the same output as the existing "ref" implementation for `_all_` possible inputs:

<https://cr.yp.to/talks.html#2021.09.03>

This is a big step towards full verification of the optimized NTRU Prime software. Next steps are matching "avx" to "ref" for the other modules (notably multiplication, where another tool has gotten through some important parts of the code but not yet everything) and matching the C code to the Sage reference code.

The "saferewrite" tool used for this verification has a broad range of applicability beyond NTRU Prime. The first example in the talk is how saferewrite catches both of the array-comparison problems that were announced in the official Frodo software. However, to enable this analysis, I had to define a Frodo array-comparison module and write reference code for that module, so that saferewrite could compare the official code to my reference code. This wasn't a big deal since this particular module is so simple, but an analogous analysis for larger components of Frodo (short of taking the entire KEM as a monolith!) would require additional work to write reference code for those modules. For NTRU Prime, this work is already done.

For each module, saferewrite compiles each implementation with clang -O1 and with gcc -O3, uses the angr symbolic-execution toolkit to convert each binary into unrolled code in a much simpler language, and uses the Z3 theorem prover (via angr's claripy) to verify equivalence or find a counterexample. The automatic equivalence chains look like this

---

**From:** pqc-forum@list.nist.gov on behalf of D. J. Bernstein <djb@cr.yp.to>  
**Sent:** Wednesday, September 8, 2021 10:16 PM  
**To:** pqc-forum  
**Cc:** pqc-comments  
**Subject:** [pqc-forum] ROUND 3 OFFICIAL COMMENT: NTRU Prime  
**Attachments:** complaint-re-apon.pdf; signature.asc

Each round-3 submission team was given a 15-minute slot at the NIST conference three months ago to present updates for, and field questions from, an online audience of about 300 people, of course including NIST.

During the NTRU Prime talk, Dr. Apon posted the text quoted below to the Slack channel for the conference, publicly accusing me of professional misconduct. Specifically, he accused me of initiating private contact with NIST so as to provide false information to NIST regarding the timing of an upcoming announcement relevant to NIST's ongoing decisions.

However:

- \* The contact was requested by NIST.
- \* The false information that Dr. Apon attributes to me is a fabrication by Dr. Apon.

Here's Dr. Apon's text:

9:38 PM

Daniel Apon @djb a quick, gentle question:

Regarding cyclotomic-based attacks: 13-14 months ago (just before the end of the 2nd Round), you privately emailed NIST PQC to suggest that you had a new attack paper (in the line) against cyclotomics that was coming out in 2-3 months. Of course, we are very eager to hear any progress along this line, even attacks that provide progress on this line, even if they don't break a cryptosystem outright (but might threaten cryptosystems in the future). However, no paper came out. We invited you to give a public 3rd Round Seminar talk on this issue in the Fall, and you ended up giving a talk at the Seminar Series in January 15 of this year. The talk presented a variety of algebraic and mathematical background that was quite interesting, but didn't suggest a clear attack vector. At the time, I suggested that you finish the paper and submit the attack paper to this conference. We didn't receive such a submission. Given this background of no attack progress against cyclotomics since the beginning of the pandemic (after claiming at least an epsilon of progress would be coming in "a month or months" well over a year ago), how would you characterize your progress in making a single epsilon of progress in attacking cyclotomic structures in lattice-based cryptography?

I replied quickly on the Slack channel to the beginning of this---

please don't misrepresent the history. nist specifically asked to be informed regarding ongoing projects, and i answered.

---but then, reading further in Dr. Apon's text, I found one outright fabrication after another, and responded accordingly:

i'll post an apon fact check to pqc-forum in due time. in the meantime, happy to answer honest questions.

I promptly contacted Dr. Apon:

I'm writing to request a retraction of the "question" that you issued during my talk. I expect the retraction to include a clear and specific acknowledgment that you fabricated the "coming out in 2-3 months" part, and that this part plays a critical role in your "question". I'll give you 1 week before escalating.

I sent another message to Dr. Apon shortly after that:

Appendix: I also expect the retraction to include the same clear and specific acknowledgment regarding your fabrication of the "month or months" part, which plays the same role in your "question".

A week later I filed a complaint with NIST, reviewing the relevant facts in detail and comparing them to Dr. Apon's accusation, and also raising the obvious procedural points.

NIST replied a week after that, admitting that the timing of Dr. Apon's "question" was "not proper", but did not address any of the gaps that I had identified between the facts and what Dr. Apon wrote. I asked two clarification questions, which were never answered.

A few weeks later I further escalated the complaint within NIST. NIST replied after a week, again not addressing the factual dispute. Two weeks later I asked what options were available within NIST for resolving the factual dispute. There was no reply.

Naturally, I cc'ed Dr. Apon on the complaint and every subsequent message. He has had three months to tell me how exactly he believes I'm getting the facts wrong or missing something that justifies his accusation. He has not taken this opportunity.

I am now posting the fact check as promised. See attached for a copy of the complaint that I filed with NIST. This incident appears to be part of a larger problem with continuing impact on NIST's evaluation of NTRU Prime, so I'm filing this as an OFFICIAL COMMENT regarding NTRU Prime.

---Dan (speaking for myself)

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group. To unsubscribe from this group and stop receiving emails from it, send an email to [pqc-forum+unsubscribe@list.nist.gov](mailto:pqc-forum+unsubscribe@list.nist.gov). To view this discussion on the web visit <https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/20210909021617.1608218.qmail%40cr.yip.to>.

# Formal complaint regarding 8 June 2021 incident

2021.06.15

Daniel J. Bernstein

**Executive summary.** A week ago Dr. Daniel Apon from NIST publicly accused me of professional misconduct. Specifically, he accused me of initiating private contact with NIST so as to provide false information to NIST regarding the timing of an upcoming announcement relevant to NIST's ongoing decisions. However:

- The contact was requested by NIST.
- The false information that Dr. Apon attributes to me is a fabrication by Dr. Apon.

All of this contact was official email exchanged with `dustin.moody@nist.gov`. Anyone at NIST with access to the email can see that it disproves Dr. Apon's accusation.

**Complaint overview.** There are three independent prongs in this complaint:

- Content: Dr. Apon issued a false accusation.
- Procedures: Dr. Apon did not follow proper procedures regarding the choice of venue and timing for issuing this accusation.
- Intent: Dr. Apon's choice of content, venue, and timing for this accusation were malicious.

Dr. Apon's false accusation was published for an audience of more than 300 attendees of a NIST conference. It's not clear at this point how much work will be required to address the damage. It is clear, however, that at a minimum Dr. Apon must issue a prompt retraction of his accusation, including clear acknowledgments of

- the specific fabrications described below and
- the critical role that two of these fabrications play in Dr. Apon's accusation.

For obvious reasons, I request an opportunity to review and approve the text and venue for the retraction before the retraction is published.

Dr. Apon has already had an opportunity to see and respond to the core of each prong of this complaint:

- Regarding content, I sent email to Dr. Apon dated 8 Jun 2021 21:53:21 +0200 quoting a specific part of his accusation, identifying this as a fabrication playing a critical role in his accusation, and requesting a retraction of the accusation. This was approximately 15 minutes after his accusation appeared. I sent followup email dated 8 Jun 2021 23:54:44 +0200 identifying a further fabrication.
- Regarding procedures, I sent email dated 9 Jun 2021 00:03:20 +0200 to Dr. Dustin Moody, cc'ing Dr. Apon, to ask whether Dr. Moody knew in advance that Dr. Apon was going to do this. (Dr. Moody said he didn't know in advance.) This was only a procedural question, not a complaint, but it included a summary of why Dr. Apon's choice of venue and timing were problematic.
- Regarding intent, I followed up in the same venue asking Dr. Apon not to misrepresent the history; saying I would "post an apon fact check to pqc - forum in due time"; and saying I was "happy to answer honest questions". Dr. Apon saw this and replied with an "eyes" emoji.

The only response I've seen from Dr. Apon is (1) generically claiming sincerity and (2) bringing his supervisors into the discussion. Dr. Apon has not addressed the specific fabrications that I identified, or the procedural issues. I conclude that it is appropriate for me to escalate to Dr. Lidong Chen at this point.

**The NIST Post-Quantum Cryptography Standardization Project (NISTPQC).** In 2016, NIST called for public submissions to its Post-Quantum Cryptography Standardization Project. NIST created a public mailing list, pqc - forum, for discussions of submissions.

I had coined the term “post-quantum cryptography” in 2003, and I have various papers on the topic. I joined a few teams preparing submissions.

NIST set a submission deadline of 30 November 2017; posted 69 “complete and proper” submissions on 21 December 2017; held a First PQC Standardization Conference on 11–13 April 2018; announced its selection of 26 submissions on 30 January 2019 for further consideration; held a Second PQC Standardization Conference on 22–24 August 2019; announced its selection of 15 submissions on 22 July 2020 for further consideration; and held a Third PQC Standardization Conference online last week, 7–9 June 2021.

After the submission deadline, Dr. Apon asked me for a faculty job (email dated 18 Dec 2017 13:55:44 -0800). I immediately filed his application as “reject without notification”. He appears to have taken this personally, and appears to be abusing his position at NIST. Given his failed job application, he should have recused himself from every action and decision related to my NISTPQC submissions, but he did not do so.

**NIST’s pattern of encouraging private input in NISTPQC.** NIST ran what it called a “survey” at the Second PQC Standardization Conference. This “survey” was not limited to any specific questions: it was an open-ended request for public inputs regarding all aspects of NISTPQC. NIST posted anonymized versions of some of the inputs: e.g., one comment advocated prioritizing “security” and “maybe size” while ignoring “performance”. NIST promised to take all of the inputs into account, not just the ones it had posted.

There were subsequent examples of NIST requesting private input. For example, in email to pqc - forum dated 30 Oct 2019 15:38:10 +0000 (2019), NIST posted technical comments regarding hybrid encryption modes and asked for feedback “either here on the pqc-forum **or** by contacting us at pqc - comments@nist.gov” (emphasis added).

In July 2020, during the second round, I finally realized what was wrong with this picture. I tweeted the following on 22 Jul 2020 13:01 GMT (<https://twitter.com/hashbreaker/status/1285922808392908800>):

After NIST’s Dual EC standard was revealed in 2013 to be an actual (rather than just potential) NSA back door, NIST promised more transparency. Why does NIST keep soliciting private #NISTPQC input? (The submissions I’m involved in seem well positioned; that’s not the point.)

This happened to be about 2 minutes before NSA sent its first message to pqc - forum. NIST announced round 3 some hours later, and eventually admitted coordination with NSA, although the exact extent of the coordination remains unclear to the public.

**Applicable transparency principles.** NSA has a long history of manipulating standards. Before NIST (at the time NBS) issued its first cryptographic standard, NSA established an internal policy goal of making this standard “weak enough to still permit an attack”. See <https://cr.jp.to/papers.html#competitions> for the full quote and references.

News reports in 2013 indicated that NSA had successfully manipulated NIST into issuing a Dual EC standard backdoored by NSA. In response, the NIST VCAT Dual EC report (see <https://www.nist.gov/system/files/documents/2017/05/09/VCAT-Report-on-NIST-Cryptographic-Standards-and-Guidelines-Process>).

pdf) strongly encouraged NIST to run “open competitions” and to follow various transparency principles. As Bart Preneel put it in that report, transparency includes “version control on all documents from an early stage, a full documentation of all decisions, and clear processes for the disposition of each and every comment received”. As Steven B. Lipner put it in the same report, transparency includes being open about “what steps were followed, what authorities were consulted or reviews sought, what comments were received, and what actions or resolutions reached”.

For NISTPQC, the call for submissions stated that NIST would “perform a thorough analysis of the submitted algorithms in a manner that is open and transparent to the public”, and that NIST’s decisions would also take into account “public comments” received in response to the submissions. There was nothing in the call indicating that NIST would also take into account *private* comments, from NSA or from anyone else.

In fact, NIST’s analysis within NISTPQC is not transparent to the public. For example, despite the NIST VCAT Dual EC report, NIST is not open about which comments it is receiving from which sources, and does not have transparent processes for the disposition of comments.<sup>1</sup> Beyond the lack of basic transparency, NIST refuses to designate NISTPQC as a “competition”, despite the VCAT report having strongly encouraged open competitions; it appears that NIST decided that it didn’t want to follow the additional rules that would have been triggered by the word “competition”.

On 28 August 2020, NIST gave a talk explaining “NIST’s decision process for Round 3”, with information that had not appeared in NIST’s report a month earlier. There was no public announcement of the talk. I filed a FOIA request (email dated 1 Sep 2020 17:10:24 +0200), in response to which NIST publicly (email dated 1 Sep 2020 16:11:44 -0700) described the talk as having been “given to the University of Maryland Crypto Reading Group” and posted the slides.

The FOIA request was actually for “the list of people that NIST emailed to give access to the talk”. The response eventually showed that NIST had invited *some* round-3 submitters—not from the University of Maryland—to attend the talk and the subsequent Q&A session. This included, for example, members of submission teams for six lattice submissions: Dilithium (2×), Falcon (2×), Frodo, Kyber (3×), NTRU (2×), and SABER. It did not include any members of the submission team for the only other lattice submission, NTRU Prime.

**Examples of consequences of the lack of transparency in NISTPQC.** One aspect of NIST’s round-3 report that surprised me was a new provable-security argument that “the security of NewHope is never better than that of KYBER”. This was almost half of NIST’s text in that document regarding NewHope, and appeared to play a prominent role in NIST’s decision to eliminate NewHope.

The reason this surprised me is that it was contradicted by the literature on the topic. What NIST wrote has *some* overlap with a proof in the literature, but there are also clear and irreconcilable gaps. NIST downplayed these gaps as “a few minor caveats”, but the simple fact is that the proof never said what NIST claimed.

I filed an official comment (email dated 23 Jul 2020 11:32:30 +0200) saying “I don’t believe the argument. I’m filing this comment to request that NIST spell out the argument in more detail for public review.” Various cryptographers followed up, pointing to some

---

<sup>1</sup>Long before July 2020, I had heard from colleagues about submitters having detailed private discussions with NIST. I had dismissed these as isolated rumors. Later, when I asked what talks had been given to NIST before round 3, Dr. Apon wrote (email dated 11 Nov 2020 17:29:56 +0000) that “we didn’t have any talks that I recall in the prior rounds”.

of the reasons that “there is no known reduction between Kyber and NewHope”. I followed up (email dated 25 Jul 2020 10:36:57 +0200) pointing to even more obstacles to making NIST’s argument work, and then pinpointed how a series of technical errors would lead to exactly what NIST had written. I concluded with a procedural objection:

NIST promised more transparency after Dual EC. I don’t understand why NIST keeps soliciting *private* NISTPQC input rather than asking for the whole evaluation to be done in public. (I also said this on the record before round 3 was announced.) This isn’t just an NSA issue; anyone who has served on program committees sees that some scientists use privacy as a shield for exaggeration. I don’t see NISTPQC procedures to compensate for overconfidence and confirmation bias; I don’t see where NIST asked for public feedback regarding these NTRU-vs.-something and NewHope-vs.-Kyber provable-security claims before the assessments appeared in the latest report; and I don’t see how similar NIST errors in round 3 are going to be corrected before they’re used for decisions.

NIST replied to my objection by reviewing what the theorem says, admitting one specific “caveat”, and then switching to *different* arguments against NewHope, arguments that had not appeared in the report.

Another aspect of NIST’s round-3 report that surprised me was the following. The report said that NIST “strongly encourages” submissions to provide a parameter set “that meets category 5”, and then criticized some submissions that provided at most category 4. The reason this surprised me is that the call for submissions had said something different: NIST “recommends” that submissions provide a parameter set “above category 3”. The report didn’t mention that NIST had changed from “recommends . . . above category 3” to “strongly encourages . . . category 5”, never mind explaining why the change had happened. When I publicly objected, one of NIST’s attempted defenses was the following remarkable admission (email dated 31 Jul 2020 14:42:02 +0000):

Throughout the process we’ve been in dialogue with various teams as they have adjusted parameter sets.

I responded as follows (email dated 2 Aug 2020 11:50:26 +0200):

You’re talking about *private* discussions that NIST has had with various teams? This is supposed to somehow be a replacement for having a change in evaluation criteria proposed and discussed in public? Wow. When did NIST announce that submitters were expected to use this private source of information rather than relying on the public announcements?

There was no reply.

I’ve skimmed the pre-round-3 email that NIST sent to the submission teams that I’m on. There were various administrative messages. There were requests to add certain types of public software documentation; these requests could also have been made public. The only bit that looks to me anything like “dialogue” regarding algorithm desiderata was a clarification question from the SPHINCS+ team:

- NIST’s report at the end of the 1st round said “A possible second-round tweak might involve a more efficient technique for protecting against multi-target attacks.”
- The SPHINCS+ team asked whether NIST was asking for “an LMS variant”.
- NIST wrote “Yes – this is what we were meaning with our comment. It would be great if you could include an LMS variant as one of your tweaks for round 2.”

In retrospect, obviously this clarification should have been public; or, even better, stated more clearly in NIST's public report in the first place. Anyway, none of these teams—including the team for a submission where the largest category proposed at the time was 4—were given any hint that NIST was going to change “recommends ... above category 3” to “strongly encourages ... category 5”.

In September 2020 I posted a paper <https://cr.yyp.to/papers.html#categories> observing how easy it was to manipulate NIST's procedures to favor particular submissions (e.g., weaker submissions), and pointing out a variety of transparency failures in NISTPQC. Many of the questions for NIST in that paper remain unanswered today.

**NIST's requests for information by 15 April 2020.** Let's rewind to the beginning of 2020. By email to pqc - forum dated 28 Jan 2020 16:23:39 +0000, NIST set a deadline of 15 April 2020 for input regarding its end-of-round-2 decisions:

In order to give NIST time to finalize our evaluation and analysis for the candidate algorithms in the 2<sup>nd</sup> round, NIST kindly requests that we be notified of new implementations, benchmarks, research papers, cryptanalysis, etc. by April 15<sup>th</sup>. After that date, factoring any of that information into our decision-making process may overly tax our resources.

In email to pqc - forum dated 26 Mar 2020 19:58:32 +0000, NIST repeated this request. It also asked people to use pqc - forum to “announce results, discuss relevant topics, ask questions, etc.” I sent email to pqc - forum dated 3 Apr 2020 00:15:35 +0200 to recommend a deadline extension:

My impression is that many people who don't know any COVID-19 victims are nevertheless losing large fractions of their previously expected work time as a result of actions taken to slow the spread of COVID-19. It's not realistic to ask for full work days from, e.g., parents suddenly taking care of kids at home all day.

I'm probably close to the low end of involuntary COVID-19 disruption but this low end certainly isn't zero; also, I've been volunteering some of my own research time to COVID-19 modeling (including scripts online and already a paper a few days ago). I still hope to make the 15 April target date for posting updated NIST-PQC benchmarks and the results of some other NISTPQC experiments I've been running, but the big picture makes me think that shifting the timeline would be a good idea.

NIST sent email to pqc - forum dated 3 Apr 2020 16:01:41 +0000 as follows (emphasis added):

If you or anyone else in the community has something important in the works, but don't think it will be done by April 15, **please notify us (by the 15th) with a brief description of the expected results and an estimate of how much longer might be needed.**

Any experienced scientist knows how time-consuming it is to fill out reports describing “expected results” that aren't yet ready to announce. However, given NIST's power over NISTPQC, people involved in the process really don't have the option of ignoring NIST's requests.

In retrospect, it's clear to me (1) that NIST has been continually manipulated by years of private lobbying regarding NISTPQC; (2) that the best response to NIST's requests, already at the 2019 conference, would have been to object to NIST taking private input; and (3) that following NIST's April 2020 request opened me up to ad-hominem attacks

as retribution for my procedural objections.<sup>2</sup> But, again, I didn't realize the problem until July 2020.

So, by email dated 15 Apr 2020 02:35:20 +0200, complying with NIST's request, I sent NIST an "overview of what's cooking from my perspective". Other cryptographers sent me copies of various messages that they had sent to NIST. Presumably there were many further messages to NIST that I never saw. There were also dozens of public team announcements on pqc - forum around that date.

**What the email said regarding timelines.** This email dated 15 Apr 2020 02:35:20 +0200 began with an "easy part" listing nine projects (some that I said I was working on, and some that I said I had heard about). This part of the message had various estimates of when announcements would be ready:

- "momentarily" although "maybe not exactly on the 15th" for updated SUPERCOP benchmark results (**2021 update:** this took slightly more than 24 hours);
- "by the end of the year" for verification of "*some* big subroutines" and "*a few* proofs" (**2021 update:** I posted a paper "Verified fast formulas for control bits for permutation networks" in September 2020);
- "doesn't look too hard to finish but has been on the back burner. Please let me know if this is something that you think is high priority" for "a step-by-step proof survey for Classic McEliece" (**2021 update:** NIST never asked me to prioritize this, and it's still on the back burner);
- "about to release" for Cortex-M4 NTRU Prime code (**2021 update:** this code was released a week later);
- "after the updated SUPERCOP Haswell results are online for everybody" for new Haswell NTRU Prime speeds (**2021 update:** this took under 48 hours);
- "this week" for "what I think is fair to describe as a dramatic performance improvement for sntrup761 keygen" (**2021 update:** that week I announced 166000 cycles where previous work was around 800000 cycles).

Notice the wide range between "this week" and "by the end of the year"; saying that this was the easy part of the email doesn't mean that these were easy projects! I also included three pointers to existing announcements that NIST might have missed.

The email continued with a "hard part", namely that "I've been doing a deep dive into lattice security", and listed five "expected results" within this. Kris Gaj had proposed a two-month extension, and it was clear to me that if NIST did this then by reshuffling research time I could get some of the results from the "hard part" online by then. So I wrote that the "hard part" would "clearly benefit from the timeline that Kris Gaj proposed to you". (**2021 update:** NIST didn't extend the deadline, so each part of the research continued on its natural schedule.)

Within the "hard part", the five "expected results" listed in the email were as follows:

- "there's a hybrid attack that reduces the 'Core-SVP security level' of, e.g., Kyber512" (**2021 update:** I spelled out the critical points in email to pqc - forum dated 7 Jul 2020 19:06:42 +0200);
- there are "indisputably reasonable 'Ring-LWE' parameters for which switching to 'NTRU' would increase the 'Core-SVP security level' while decreasing the probability of decryption failures" (**2021 update:** the point of this is the looseness of a partic-

---

<sup>2</sup>For example, in reply to a technical message that I had posted on another topic, Dr. Apon sent email to pqc - forum dated 22 Aug 2020 13:27:43 -0700 going far out of his way to mention my "private input" to NIST. NIST has not similarly mentioned the private input it received from other submitters.

- ular proof; after checking examples of NIST ignoring previous examples of proof looseness, I decided to put this project on the back burner);
- the “lattices that the NTRU Prime submission presents for attacking ‘Ring-LWE’ generalize the standard Kannan–Bai–Galbraith lattices to allow more short vectors, and (new result) also outperform the standard lattices” (**2021 update:** I have enough computer experiments to be confident in what I wrote, but analysis and optimization are continuing so as to understand what happens for larger sizes);
  - “the huge enumeration speedup announced as work in progress” by another team “also makes a big difference in the enumeration-vs.-sieving cutoffs, especially in realistic models of large-scale computation that account for the costs of memory” (**2021 update:** this analysis ended up in the round-3 NTRU Prime submission in October 2020); and
  - there’s a “new way to exploit the structure of cyclotomics”—I said this wasn’t “big enough to reach the lattices in NISTPQC submissions, but it breaks solidly through barriers claimed in previous work”<sup>3</sup> (**2021 update:** I gave a talk on “Valuations and S-units” in January 2021 explaining the most important number-theoretic background; I have a talk on “S-unit attacks” scheduled for August 2021, and expect that my coauthors will authorize revealing the most important new idea during that talk).

The email did not provide time estimates for any of these five specific items.

I hadn’t realized at this point that NIST should have been categorically forbidding private input from the outset. However, given the long history of errors regarding lattice security in particular, I knew that public review was essential in this area. The email closed with a clear warning on this topic:

I can imagine selection of lattice systems being the hardest decision that NIST is facing, and I can imagine some of these lattice-security issues showing up as part of the decision. **In a non-COVID-19 world, I would have had some of this online already as papers, and then there would be time for the community to comment, which I think is especially important in an error-prone area.**

On the other hand, I also find it easy to imagine that you’re planning to avoid tricky issues in this round and make decisions based primarily on undisputed attacks, performance features, simplicity, stability, diversity, etc. This reminds me: has NIST decided how high a weight it’s putting on patents in this round?

(Emphasis added.) This was before NIST issued its round-3 report with, e.g., a decision based on a “security of NewHope is never better than that of KYBER” provable-security claim that, as mentioned above, had not been publicly reviewed and did not match what the literature said.

NIST engaged me in followup discussion over the next two weeks. None of my followup email had any timeline predictions.

**Content of Dr. Apon’s accusation.** In a nutshell, Dr. Apon accuses me of having deceived NIST regarding the timing of an announcement regarding cyclotomics. This accusation relies critically on the following two fabrications by Dr. Apon:

---

<sup>3</sup>Since cyclotomics show up later in this complaint, here’s the full quote: “Deepest result, joint work with various people: There’s a new way to exploit the structure of cyclotomics, quickly finding a short secret vector from any lattice in a much larger class than handled by previous algorithms. This class isn’t big enough to reach the lattices in NISTPQC submissions, but it breaks solidly through barriers claimed in previous work. Top consequence from my perspective: New reason to be scared of cyclotomics.”

- Dr. Apon quotes me as telling NIST 13–14 months ago that a cyclotomic announcement would be in “a month or months”. I said no such thing.
- Dr. Apon paraphrases me as saying that a “paper” was “coming out in 2-3 months”. I said no such thing.

Dr. Apon contrasts “coming out in 2-3 months” and “month or months” to the lack of any publication for more than a year. Without Dr. Apon’s fabrications of a much shorter claimed timeline, there would be nothing remarkable about a serious scientific project taking more than a year.

The following paragraphs go step by step through Dr. Apon’s accusation, pinpointing the above timeline fabrications and further fabrications by Dr. Apon. The “9:38 PM” time shown here is in the CEST time zone on 8 June 2021.

Daniel Apon 9:38 PM @djb a quick, gentle question:

This introduction was followed by more than 1000 characters (216 words) of text (all posted at once as part of the same message), so the reader instantly sees that “quick” is facetious, and presumes that “gentle” is also facetious. This doesn’t make the reader imagine that Dr. Apon’s subsequent claims regarding the history are outright fabrications.<sup>4</sup>

[Next words in Dr. Apon’s message:] Regarding cyclotomic-based attacks:

The reader interprets further text within this scope.

[Next words in Dr. Apon’s message:] 13-14 months ago (just before the end of the 2nd Round), you privately emailed NIST PQC

Let’s compare this to the facts:

- NIST set a deadline of 15 April 2020 for input, saying that input after this “may overly tax our resources”.
- Beyond public announcements, NIST specifically asked to be given a “brief description” of any “expected results” and an “estimate of how much longer might be needed”.
- Unlike many other government agencies, NIST refused requests for deadline extensions in light of COVID-19.
- Many people provided input on or around 15 April 2020, the deadline that NIST had set. I was one of these people.
- Three months later, NIST ended the second round with an announcement on 22 July 2020. NIST could have taken longer if it wanted to—there was never any public commitment to this particular date until the announcement happened.

Does Dr. Apon inform the reader that NIST had asked for input by 15 April 2020? No. That many other people also provided input on or around this date? No. That NIST also received private input from many other submitters? No. That NIST received private input from other submitters long before this? No. That my email cautioned NIST regarding the importance of public review especially in the area of lattice security? No.

Yes, 15 April 2020 was 13–14 months ago. But is it correct that this was “just before the end of the 2nd Round”? No. First, “just before” can’t reasonably be understood to include a stretch of three months. Second, like other members of the general public, I didn’t know when round 2 would end—as far as I know, NIST was free to take as much time as it wanted—so saying “just before the end of the 2nd Round” is nonsensical as

---

<sup>4</sup>If I were to change the title of this document to “A quick, gentle complaint regarding Dr. Apon”, the reader would still take the rest of the document at face value.

a description of what I did in April 2020.

Some readers will know that NIST never publicly scheduled the end of the round until issuing the 22 July 2020 announcement; so, if they think about it for a moment, they'll realize that what Dr. Apon said here is nonsense. However, for a reader who doesn't know this and who takes Dr. Apon's text at face value, the text implies that NIST *had* publicly scheduled the end of the round, and says that I targeted NIST with a surprise just before this. That isn't true.

[Next words in Dr. Apon's message:] **to suggest that you had a new attack paper (in the line) against cyclotomics that was coming out in 2-3 months.**

Does Dr. Apon inform the reader that the message actually covered 14 different projects, with many different timelines? No.

The core of what the reader has been told at this point is that I wrote to NIST 13–14 months ago to suggest that a new attack paper on cyclotomics was coming out in 2–3 months. The reader assumes that no such paper has appeared—why else would Dr. Apon start by talking about the timeline and using words like “**suggest**”?—and concludes that I provided grossly inaccurate information to NIST.

But the timeline information is a fabrication by Dr. Apon. I never gave NIST a timeline for the cyclotomic attacks.

Perhaps Dr. Apon will point to the comment “In a non-COVID-19 world, I would have had some of this online already as papers, and then there would be time for the community to comment, which I think is especially important in an error-prone area” and say that these papers obviously can't have lost more than a few months given that this was a message in April 2020. However:

- “Some” is not “all”. Can someone who reads “I would have had some of this online already as papers” claim that this says “I would have had all of this online already as papers”—or select one part X of this and claim that it says “I would have had X online already as a paper”? No.
- Suppose a message is saying, in April 2020, that something was within a few months of done in January 2020, but because of COVID-19 isn't done yet. Can one claim that this message is saying that it will be done in July 2020? No.

The bottom line is that the email never said what Dr. Apon claims it said.

[Next words in Dr. Apon's message:] **Of course, we are very eager to hear any progress along this line, even attacks that provide progress on this line, even if they don't break a cryptosystem outright (but might threaten cryptosystems in the future).**

Does Dr. Apon inform the reader that the email had already put this cyclotomic attack into this category, saying that the attack wasn't “big enough to reach the lattices in NISTPQC submissions, but it breaks solidly through barriers claimed in previous work”? No. Dr. Apon suppresses this information, making it sound as if NIST had been given no information regarding the nature of the advance and was merely guessing on its own that the advance wasn't a full break of NISTPQC submissions yet.

Omitting this information is exaggerating what the message said. When the results *are* released, is Dr. Apon going to rely on the same exaggeration to claim that he thought the results were already supposed to be a full break of a NISTPQC submission? Will he express disappointment that the results are “only” breaking some underlying problems? I haven't seen NIST trying to preemptively sabotage announcements that it hears

about in advance from other people.

[Next words in Dr. Apon's message:] **However, no paper came out.**

It's correct that this project has not released a paper yet.

The "however" is drawing a contrast between (1) this fact and (2) claims that are entirely fabricated by Dr. Apon. When the underlying fabrication disappears, the contrast also disappears.

[Next words in Dr. Apon's message:] **We invited you to give a public 3rd Round Seminar talk on this issue in the Fall,**

The reader will understand "this issue" to be the "new attack paper" against "cyclotomics". The reader is thus being told that NIST was asking me to give a talk specifically on the cyclotomic results; and that NIST was asking me to make the results available to the public. But both parts of this are fabrications by Dr. Apon.

NIST's actual invitation, by email dated 17 Aug 2020 16:36:05 +0000, was to give a talk "about the topic(s) of your choice in the security of lattice cryptography". The only specific lattice-security topic mentioned was "the state of the art of lattice reduction algorithms". Experts understand "lattice reduction algorithms" to include, e.g., LLL and BKZ, and to exclude, e.g., the algorithms featured in previous cyclotomic breaks. (I do have some work on reduction algorithms.)

In followup discussion, NIST said that "one area in particular that we would like to be as informed about as possible is the concrete security of lattice-based KEMs" against "lattice reduction", and that "another lattice topic we would like to hear about is an update on the status of any work towards algebraically cryptanalyzing structured lattices, e.g. closely approximating the shortest vector in (power-of-2) cyclotomic lattices".

So, no, it's not true that NIST invited me specifically to give a talk about cyclotomics. This was eventually on NIST's list of topics, but wasn't the only topic, and wasn't the first topic in the list, and wasn't the topic mentioned in NIST's first invitation message.

It's also not correct that the invitation was specifically to give a "public" talk. Having the talk be public was only an option mentioned "if you would prefer". I did state a preference for this option (email dated 3 Sep 2020 09:49:57 +0200):

Procedurally, I now have some inkling of how much communication to and from NIST has been outside public view, and I am extremely uncomfortable with the lack of transparency, so I will certainly insist on having a public announcement and recording of any talk that I give to NIST.

Dr. Apon's fabrications regarding NIST's talk invitation are less important than Dr. Apon's central fabrications of timeline information, but they still tilt the overall picture.

[Next words in Dr. Apon's message:] **and you ended up giving a talk at the Seminar Series in January 15 of this year.**

This is correct if "a talk" is taken out of context. However, the context makes the reader understand "a talk" as a "talk on this issue", i.e., a talk about the "new attack paper" against "cyclotomics"—and that's not true. Let's look at what actually happened.

In its initial invitation (email dated 17 Aug 2020 16:36:05 +0000), NIST had said that its talk series was about "important technical areas" and that "For many of these areas, we have 2 or 3 team members who understand the areas in greater depth, but the broader team does not have a good understanding of all these issues yet".

I wrote (email dated 3 Sep 2020 09:49:57 +0200): "So that I can better understand the

context and objectives for this talk series, and what the audience is likely to know, can you please send me a list of the previous talks and other planned upcoming talks?" One of the examples NIST mentioned (email dated 3 Sep 2020 14:53:03 +0000) was an upcoming "introductory style talk on the basics of LLL/BKZ followed by their new results".

My assessment was that NIST was asking for remedial education. As noted above, given NIST's power over NISTPQC, people involved in NISTPQC really don't have the option of ignoring NIST's requests. But I was also extremely busy with other work items—for example, round-3 submissions filed in October 2020—so it took me some time to respond. I then commented (email dated 10 Nov 2020 13:51:36 +0100) that "it's informative, and disturbing, to learn that the audience doesn't already know the basics of LLL/BKZ".

NIST claimed (email dated 11 Nov 2020 17:29:56 +0000) to be "confident that the NIST PQC team members are familiar with e.g. LLL/BKZ (and at least reasonably familiar with all of the technical minutiae regarding more modern tweaks to lattice reduction algorithms and algebraic cryptanalysis of lattices over various number fields)".

I wrote (email dated 28 Nov 2020 13:28:34 +0100) "Can you please clarify what level of understanding you mean by 'reasonably familiar with all of the technical minutiae'? I'd expect this to mean that audience members would be able to fully define (e.g., implement) these algorithms without referring to notes, but I'm having trouble reconciling this with the September comment indicating NIST's interest in 'basics of LLL/BKZ' as part of someone's introductory talk." NIST never answered the question.

Based on the limited information available, I proposed (also in the 28 Nov 2020 13:28:34 +0100 email) a talk on "Valuations and S-units", with the following abstract:

This talk reviews a standard infinite-dimensional number-theoretic lattice that simultaneously shows how large numbers are and how they factor. The ability to decode this lattice in some surprisingly large cases plays a critical role in a new wave of attacks against ideal-lattice problems. This talk will focus on defining the lattice, with many examples to illustrate.

This is an introductory talk aimed at a broad audience. Prerequisites: mathematics education through a course in undergraduate abstract algebra (commutative rings and fields).

NIST said "this is perfect", and that's what the talk ended up being about.

[Next words in Dr. Apon's message:] **The talk presented a variety of algebraic and mathematical background that was quite interesting, but didn't suggest a clear attack vector.**

The reader has been led to believe that NIST had invited me to give a talk about the "**new attack paper**" against "**cyclotomics**", and that I accepted—"but" then the talk didn't suggest an attack against cyclotomics.

It's not true that this is what NIST invited me to talk about. It's not true that this is what I agreed to talk about. The talk abstract that I proposed, and that NIST accepted, said that there's a standard number-theoretic lattice used in a new wave of attacks, and that this talk would focus on *defining* the lattice.

Does Dr. Apon inform the reader that attacks against cyclotomics were never within the agreed talk scope? No. Instead he deceives the reader into believing that I had been invited to give a talk on cyclotomic attacks, and had agreed to give a talk on cyclotomic attacks, "**but**" then failed to deliver.

[Next words in Dr. Apon's message:] **At the time, I suggested that you finish the paper and submit the attack paper to this conference.**

The recording shows that, in the Q&A session, I mentioned work in progress; Dr. Apon said "Do you have an ETA for this unreleased paper?"; and I said "Well, I'm one of six coauthors, so—it's one of these fun things where the results somehow keep getting better and better. The starting point was, okay, definitely smashed through this barrier, now let's see how much further we can go." Publications are, needless to say, driven by what makes sense scientifically.

Dr. Apon then sent email dated 15 Jan 2021 18:38:50 +0000 saying "I'm definitely looking forward to your unreleased paper making progress on the topic. As a reminder, the 3rd NIST PQC Standardization Conference is coming up; perhaps that is a good place to submit (and then also to a separate, academic conference with published proceedings as well)." So it's correct that he suggested this as a place to submit. So what? This wouldn't be remarkable without his timeline fabrications.

[Next words in Dr. Apon's message:] **We didn't receive such a submission.**

Dr. Apon has already told the reader that I had written to NIST 13–14 months ago to suggest that a new attack paper on cyclotomics was coming out in 2–3 months. Now Dr. Apon contrasts this with my not submitting any such paper to a NIST deadline a year after that.

It's correct that I didn't submit anything like this to the NIST conference.<sup>5</sup> But the entire contrast is against a "2–3 months" fabrication by Dr. Apon, not against something I ever said.

[Next words in Dr. Apon's message:] **Given this background of no attack progress against cyclotomics since the beginning of the pandemic**

What does the reader understand "progress" to mean here?

Often the word "progress" is comparing publications to earlier publications: this paper is making progress in the following way compared to that paper. Given the word "background", I'd expect the reader to understand Dr. Apon to be saying that, since the beginning of 2020, nobody has published any papers making progress on cyclotomic attacks. But then how does Dr. Apon explain <https://eprint.iacr.org/2021/600>, which was posted a month before Dr. Apon's text, and says in the abstract that it shows "that the decomposition group of a cyclotomic ring of arbitrary conductor may be utilised in order to significantly decrease the dimension of the ideal (or module) lattice required to solve a given instance of SVP"? That paper has certain limitations, but I don't see how Dr. Apon can claim that the paper isn't "progress".

It's possible that the reader instead understands Dr. Apon to be referring specifically to the project I'm involved in. The word "progress" in reference to one project generally refers to the steps inside that project, rather than comparing the project to previous work, so this type of reader understands Dr. Apon to be saying that this project hasn't made progress since the beginning of 2020. The case that Dr. Apon lays out for this relies entirely upon his fabricated timeline claims.

[Next words in Dr. Apon's message:] **(after claiming at least an epsilon of progress would be coming in "a month or months" well over a year ago),**

Given the context, the reader understands the quotation marks around "a month or

---

<sup>5</sup>Has NIST issued any statements regarding other people's supposed failure to submit something to NIST's conference? Not as far as I know.

months” to indicate that these words appeared in my email to NIST regarding this cyclotomic project. Dr. Apon is telling the reader that my email (1) claimed that something would already happen for this project in “a month or months” and (2) suggested that a paper was coming out for this project in 2–3 months.

But this “a month or months” quote is another fabrication by Dr. Apon. This is even worse than the earlier “coming out in 2-3 months” fabrication: not only does Dr. Apon attribute to the email something that the email simply did not say, but Dr. Apon inserts quotation marks so as to bolster the credibility of this fabrication.

[Next words in Dr. Apon’s message:] **how would you characterize your progress in making a single epsilon of progress in attacking cyclotomic structures in lattice-based cryptography?**

Starting from his fabricated timeline claims, Dr. Apon accuses me of deceiving NIST regarding the timeline, and then closes by publicly disputing that any progress has happened. He couches this dispute as a question. Since this isn’t an honest question, it doesn’t deserve an answer.

**Venue and timing of Dr. Apon’s accusation: disrupting a scheduled talk.** NIST’s Third PQC Standardization Conference took place online last week, 7–9 June 2021, as mentioned above. The conference charged a \$25 registration fee but was open to the general public. NIST gave each round-3 submission team a 15-minute talk slot at the conference to present updates and field questions, and scheduled further talks regarding, e.g., performance comparisons. NIST also set up a chat system for the public to “submit questions and engage in group discussions” for the conference. The chat system and the video system each listed more than 300 people.

I gave the talk for one of the round-3 submissions, NTRU Prime, on behalf of a team of 10 people. Please note that I’m speaking for myself in this complaint.

NIST has full power over what happens in NISTPQC, and can send messages to submitters or to pqc - forum at any moment. Dr. Apon has sent 88 messages to pqc - forum. **But Dr. Apon chose to issue his accusation during the NTRU Prime talk.**

The talk was scheduled to run from 15:25 to 15:40 D.C. time. However, this was the fourth talk in the session, and the session was running late. My best reconstruction based on the data available at this point is that the talk ran from 15:34 to 15:49 D.C. time. Dr. Apon published his accusation on the chat system at 15:38 D.C. time.

As noted above, Dr. Apon introduced his accusation as “a quick, gentle question”, followed by lengthy text. The typical reader would be captured by the clickbait; read further; see the accusation of misbehavior in Dr. Apon’s first sentence after the introduction; find this interesting; continue reading, finding further text amplifying the accusation of misbehavior; and, eventually, find a “question” at the very end of Dr. Apon’s text. Any such reader would have been

- distracted from listening to the NTRU Prime talk, and
- immediately prejudiced against the NTRU Prime talk, given these accusations against the speaker.

Note that NIST relies heavily on public analysis of submissions—for example, public work to make submissions perform well in various environments. The talk was not just for NIST but for the entire audience, explaining reasons to be interested in NTRU Prime. It’s easy to predict that Dr. Apon’s disruption will end up taking some public analysis away from NTRU Prime.

Presumably Dr. Apon will respond that this talk disruption was justified. Dr. Apon knew that my talk was going to mention the history of attack advances *against cyclotomics* and the possibility of further attack advances *against cyclotomics*; Dr. Apon's accusation is specifically that I had deceived NIST regarding the timing of an announcement *regarding cyclotomics*; Dr. Apon thought it was important for people listening to the talk to know this, and to distrust the talk accordingly.

However, *sometimes* accusations are false (as illustrated by Dr. Apon's accusation), which is why *in all cases* it's important to follow proper dispute-resolution procedures, including normal due-process safeguards. Did Dr. Apon follow these procedures before broadcasting his accusation at a conference attended by more than 300 people? No, he didn't. Did he tell Dr. Moody what he was planning to do? According to Dr. Moody, no, he didn't. Perhaps Dr. Apon will try to claim that his fabrications regarding my email (and regarding NIST's talk invitation and so on) were merely the result of an amazing series of memory failures rather than dishonesty; but did he *check* the email before issuing his accusation? No, he didn't.

Furthermore, anyone who looks at the talk slides or listens to the talk video (I've made the slides and video available at <https://cr.yp.to/talks.html#2021.06.08>) can see that the talk was covering many points other than cyclotomics. It appears that Dr. Apon posted his 216-word "Regarding cyclotomic-based attacks" text at the first moment cyclotomics were mentioned—but, before this, the talk had already covered

- the fact that public security reviewers are overloaded;
- a broad overview of recent advances in lattice attacks;
- the failure of "provable security" to stop small lattice systems from being broken;
- performance requirements sometimes forcing the use of small lattice systems;
- NTRU Prime's goal of, and success in, reducing the attack surface for small lattice systems;
- decryption failures as an example of a NISTPQC attack tool that NTRU Prime had eliminated from the outset; and
- the advantages of proactively reducing attack surface rather than merely reacting to breaks.

Cyclotomics then appeared as another example of a NISTPQC attack tool that NTRU Prime had eliminated from the outset—but, even if Dr. Apon objected to this example, how could this objection possibly have been so important as to justify disrupting a talk that was obviously broader than this? Why not allow the talk to finish, and calmly follow up with objections afterwards?

**Subsequent events.** The conference organizers at NIST had offered the option of sending a video in advance. I had taken this option, because of concerns regarding time zones and regarding difficulties connecting to NIST's video system. Like typical audience members, I was watching the chat system for questions while the video played.

I was, to put it mildly, surprised seeing Dr. Apon's "question" on the chat system. I saw that the message was starting by deceiving readers regarding NIST's request for input, so I replied regarding that:

djb 9:39 PM please don't misrepresent the history. nist specifically asked to be informed regarding ongoing projects, and i answered.

Dr. Apon replied, repeating his "question" to avoid commenting on his misrepresentation of the history—a classic example of an ad-hominem attack:

Daniel Apon 9:39 PM **So, no progress?**

Reading further in Dr. Apon's text, I found one outright fabrication after another, and responded accordingly:

djb 9:44 PM i'll post an apon fact check to pqc-forum in due time. in the meantime, happy to answer honest questions.

[emoji reactions: "eyes" from Daniel Apon; "hushed" from Yuji Suga]

The damage to the talk was already done. The only further questions at the conference regarding the talk were from NIST employees:

angela.robinson 9:51 PM I am curious to know if you are planning a follow-up talk to your talk given at the public NIST PQC seminar series. @djb

[emoji reaction: "+1" from Daniel Apon]

djb 9:54 PM @angela.robinson Thanks for your question! As I mentioned in the talk, there's a second part coming up. This is scheduled for August.

angela.robinson [reply in thread] Great! Sorry if I missed that part. The pace was a bit fast, but I look forward to it.

Daniel Apon [reply in thread] **Looking forward to it!**

djb 9:56 PM To clarify since there are multiple talks at issue: I mentioned in the January talk that a second talk was coming up on S-unit attacks; that talk is now scheduled for August; I didn't mention that talk in my talk today.

[emoji reactions: "+1" from Angela Robinson and Daniel Apon]

Daniel Apon [reply in thread] **Where will the talk be?**

Daniel Smith-Tone 10:04 PM @djb Did the parameters for sntrup4591761 and ntrulpr4591761 change between round 1 and round 2? I am asking because in round 1 the specification document claims NIST security level 5, but then level 3 is claimed in the round 2 spec and afterwards.

djb 10:14 PM @Daniel Smith-Tone I believe the submission document addresses everything, but let me try to answer your question, speaking for myself. There's no change in the parameter set. There have been worrisome advances in attacks against all small-lattice submissions. There are also massive ambiguities in NIST's definitions of security levels, as illustrated by NIST's SHA3-256 security evaluation jumping from  $2^{80}$  to  $2^{146}$  in 2016. Some ambiguities in NIST's definitions of security levels were partially resolved by NIST statements in round 1, so the NTRU Prime assignments were revised accordingly. Many problematic ambiguities remain; see generally <https://cr.yp.to/papers.html#categories>.

Daniel Smith-Tone 10:15 PM I think that the submission document does address this, but I thought I would ask directly since the source is available. I also agree with some of what you just said.

djb 10:15 PM Section 5.4 of that document identifies various specific definitional questions that NIST still hasn't answered.

Daniel Smith-Tone 10:16 PM I'm reading that right now.

djb 10:17 PM Please let me know if anything is unclear.

Meanwhile I wrote to Dr. Apon (email dated 8 Jun 2021 21:53:21 +0200) as follows:

I'm writing to request a retraction of the "question" that you issued during my talk. I expect the retraction to include a clear and specific acknowledgment that you

fabricated the “coming out in 2-3 months” part, and that this part plays a critical role in your “question”. I’ll give you 1 week before escalating.

Later I wrote to Dr. Apon (email dated 8 Jun 2021 23:54:44 +0200) as follows:

Appendix: I also expect the retraction to include the same clear and specific acknowledgment regarding your fabrication of the “month or months” part, which plays the same role in your “question”.

I wrote to Dr. Moody, cc’ing Dr. Apon and the NTRU Prime team, as follows (email dated 9 Jun 2021 00:03:20 +0200):

As you know, each round-3 submission team was given a 15-minute slot at the ongoing NIST conference to present updates for, and field questions from, an online audience of about 300 people, of course including NIST.

For NTRU Prime in particular, I gave the talk. I presume you saw that, early in the talk, Dr. Apon posted a “question” to the Slack channel, a “question” that typical audience members would understand as accusing me of misbehavior. A copy of his accusation appears below.

I will, needless to say, take appropriate action to respond to these accusations. I’ve already contacted Dr. Apon regarding the substance of the accusations, and have offered him a week before I escalate.

However, in the meantime, I believe it’s appropriate to immediately address the following procedural question to you, since Dr. Apon is, if I understand correctly, a NIST PQC team member under your supervision. If I’ve misunderstood the management structure within NIST and should contact Dr. Chen instead, please let me know.

Dr. Apon is a member of the NIST PQC team. The NIST PQC team has full power over what happens in this competition, and can ask questions on pqc-forum at any moment. However, Dr. Apon chose to issue accusations *during the NTRU Prime talk*. Given the length, timing, and nature of his text, it seems reasonably clear that he prepared the text in advance and chose the timing to maximally disrupt the talk. Obviously the video kept playing, but many people

- \* watch the Slack channel for the occasional questions while listening to talks,
- \* would have been distracted from listening because of the obviously interesting nature of Dr. Apon’s text, and
- \* would then have been prejudiced against the talk because of the accusations communicated by the text.

So here’s the question for you: Did you know he was going to do this?

Since what happened here is an issue for the NTRU Prime team as a whole, I’m cc’ing the team.

Dr. Apon replied as follows (email dated 9 Jun 2021 00:10:48 +0000):

My question was asked honestly and in earnest (that is, as a person sincere and serious in behavior and convictions).

I think your question is best addressed to my supervisors. I’ve included Lily Chen and Matt Scholl on the CC list. You’re welcome to discuss with them.

Dr. Moody replied as follows (email dated 9 Jun 2021 12:51:39 +0000):

Thank you for your message and sharing your concern with me. As a quick re-

sponse to your question - no, I did not know in advance that Daniel Apon would ask you a question on slack during the NTRUprime update. Our PQC team members can certainly post questions without needing to ask me. I actually had to skip the last 30 minutes of the day (my son was running his last race of his high school career at a track meet), so I wasn't following live. I didn't check back into work things until this morning.

I am sorry if you felt his question disrupted your talk. I'll talk to Daniel, but from his response to you he says he wasn't intending to disrupt, but wanted to ask you an honest question.

**Concluding remarks.** I've reviewed this complaint, am confident in its accuracy, and see many reasons to have it online as soon as possible. Am I simply posting it?

No. I'm sending it to a limited audience: Dr. Apon (the subject of this complaint), Dr. Moody (in charge of NISTPQC), Dr. Chen (the supervisor that this complaint is addressed to), Dr. Matthew Scholl (Dr. Apon already brought Dr. Scholl into the discussion), and the NTRU Prime team (directly damaged by Dr. Apon's talk disruption, beyond the damage caused by the content of Dr. Apon's accusation).

Why am I not posting the complaint immediately? Because that wouldn't be the proper procedure. I am accusing Dr. Apon of misconduct, and he is entitled to an opportunity to defend himself. I can't imagine how he can justify his ludicrously unprofessional, clearly dishonest actions—but that's not the point. Sometimes people *can* successfully defend themselves, so civilized society sets up procedures that recognize this and that, as noted above, are to be followed in *all* cases.

Part of what I'm complaining about, obviously, is that Dr. Apon didn't follow such procedures: on the contrary, he simply went ahead and *published* his accusation. Given that the accusation is public, I'm certainly entitled to defend myself in public (while, for comparison, Dr. Apon has no such excuse for his procedural violations), so I could easily justify posting this complaint today. Delaying this posting means that Dr. Apon's fabrications will have more time to spread and will cause more damage. But delaying this posting—for a limited time—is still the right thing to do.

One week is ample time for NIST management to compare all relevant email to Dr. Apon's fabrications and to take appropriate action. Given the nature of Dr. Apon's action, I expect this matter to be given high priority. I will, of course, give due consideration to anything that Dr. Apon says now in his defense, and I will give due consideration to any explanation of why more time is requested.

---

**From:** 'Moody, Dustin (Fed)' via pqc-forum <pqc-forum@list.nist.gov>  
**Sent:** Thursday, September 16, 2021 7:49 AM  
**To:** D. J. Bernstein; pqc-forum  
**Subject:** [pqc-forum] Re: ROUND 3 OFFICIAL COMMENT: NTRU Prime

Dan,

That there was an impolite exchange on a Slack thread does not mean the PQC standardization process is unfair or biased in any way. We all have individual perspectives, as researchers and scientists, and it is inevitable that some disagreements and misunderstandings will occur during this process. However, we work as a team to ensure that the final outcomes of this process are rigorous and fair. We take our responsibilities very seriously, and we will continue to be fair and impartial as we evaluate the different algorithms. We make our decisions scientifically based upon the technical results we see published or presented. The community is welcome to provide feedback if they feel this is not the case.

Also, as a reminder, we would like to keep the pqc-forum primarily for technical discussions. We will follow up with you directly on other issues.

Dustin  
The NIST PQC team

---

**From:** D. J. Bernstein  
**Sent:** Wednesday, September 8, 2021 10:16 PM  
**To:** pqc-forum  
**Cc:** pqc-comments  
**Subject:** ROUND 3 OFFICIAL COMMENT: NTRU Prime

Each round-3 submission team was given a 15-minute slot at the NIST conference three months ago to present updates for, and field questions from, an online audience of about 300 people, of course including NIST.

During the NTRU Prime talk, Dr. Apon posted the text quoted below to the Slack channel for the conference, publicly accusing me of professional misconduct. Specifically, he accused me of initiating private contact with NIST so as to provide false information to NIST regarding the timing of an upcoming announcement relevant to NIST's ongoing decisions. However:

- \* The contact was requested by NIST.
- \* The false information that Dr. Apon attributes to me is a fabrication by Dr. Apon.

---

**From:** pqc-forum@list.nist.gov on behalf of Christopher J Peikert <cpeikert@alum.mit.edu>  
**Sent:** Thursday, September 30, 2021 10:30 AM  
**To:** pqc-forum  
**Subject:** [pqc-forum] ROUND 3 OFFICIAL COMMENT: NTRU Prime

Summary of this comment:

1. The NTRU Prime FAQ starts with an objectively false factual claim about competing submissions Kyber and SABER.
2. This false claim is central to the FAQ's misleading attempt to suggest that these systems infringe on a patent.
3. Requests to the NTRU Prime team to remove the false claim and insinuation were refused.
4. I therefore believe that this is not an honest mistake, but a deliberate attempt to smear competing proposals with false disparaging claims and FUD.
5. I request that NIST consider what to do about patterns of behavior like this.

The first answer of the NTRU Prime FAQ, which appears on the official project website [\[link\]](#) and was repeated by Dan Bernstein on the pqc-form on 11 December 2020 [\[link\]](#), says this (emphasis added):

"There are known patent threats against ... Kyber, SABER, and NTRU LPrime (ntrulpr). **These proposals use ... a 2x ciphertext-compression mechanism that appears to be covered** by U.S. patent 9246675 expiring 2033."

(The ellipses replace another claim of threat from a different patent, which is outside the scope of this message, but was also shown to be severely flawed; see, e.g., [\[link\]](#), [\[link\]](#).)

**As a matter of objective fact, the "2x" claim is false.** While Kyber and SABER do perform some mild ciphertext compression, they do not, and could not, come close to 2x with the mechanism they use.

(I pointed out this false "2x" claim on the pqc-forum on 11 December 2020 [\[link\]](#), and again on 21 May 2021 [\[link\]](#), and in private correspondence with the NTRU Prime team, with an explicit request to correct it, but the team refused.)

Why does this matter?

**The false "2x" claim is central to the FAQ's attempt to tie Kyber and SABER to the cited patent.** Specifically, the "appears to be covered" claim implicitly conflates the patented mechanism, which does provide (near-)2x compression, with the unpatented prior-art method that Kyber/SABER use.

Kyber/SABER's compression mechanism is, informally: "drop some low bits of certain integers, keeping the several high bits needed for correct decryption." This method appears in at least four well known works of prior art to the cited patent, some of which are cited in every version of the Kyber submission. For details, see the last part of my pqc-forum message from 21 May 2021 [\[link\]](#).

The patent describes a *different* compression mechanism whose main benefit is that it can provide near-2x in certain contexts, by keeping just a single bit of certain integers. Kyber/SABER do not use this method, and the patent does not claim the above prior-art method that they do use. A detailed explanation of the prior art and the differences between the methods is given in my pqc-forum message from 22 May 2021 [\[link\]](#).

In private correspondence with the NTRU Prime team, based on the above reasoning I stated that **the FAQ's "appears to be covered" claim is highly misleading**, and requested that it be removed. The team refused.

The above summarizes, for the record, the history regarding the facts and analysis. The rest of this message contains my conclusions about the situation, and a discussion of how to proceed from here.

**I believe that the FAQ entry is a deliberate attempt to smear competing proposals with false disparaging claims and FUD.** Of course, the false "2x" claim could originally have been an unintentional error---albeit a sloppy one, showing unfamiliarity with basic properties of the schemes.

However, the team's refusal to fix even this elementary factual error leads me to conclude that the claim has been made intentionally to deceive, i.e., to conflate the unpatented prior art with the patent's near-2x method, and to misleadingly suggest that Kyber/SABER infringe on the patent. Without "2x," there's no link to the patent, and the FAQ entry falls apart (along with subsequent entries that are premised on it).

### **What next?**

I hope the above material sets the record straight. But this example raises the broader issue of NIST PQC participants who exhibit a pattern of the following behavior:

1. Falsely disparage other submissions and/or the process itself.
2. Receive corrections showing these claims to be factually false or otherwise meritless.
3. Make no withdrawal of the false claims. Even worse, give no acknowledgment of the corrections. Even worse than that, persist in spreading the false claims.

(Some other examples of this pattern appear at the end of this message.)

This kind of behavior is outside the bounds of fair play. It sows confusion among non-experts who may only be able to see a "controversy," and it badly wastes the community's time that could be better spent on more productive matters. (Brandolini's law estimates the cost at 10x, but I think that's too low in this context.)

To be absolutely clear: I am not talking about honest mistakes or misunderstandings that are acknowledged and corrected. Indeed, this describes the vast majority of situations in the NIST PQC process, in which submitters and other participants have resolved matters without difficulty.

Procedurally, I think NIST should seriously consider this issue. I can think of a few options for how it could respond, such as:

1. Take no official action. Let people say whatever they want to, and hope that other (unspecified) mechanisms address such behavior. This has the big disadvantage that it does not offer any clarity to non-experts and the broader community.
2. Make an official statement on its findings of the relevant facts, and perhaps its analysis of the consequences. This has the advantage of offering clarity to the community.
3. Do 2, and also penalize submissions/submitters who show a pattern of this kind of behavior, perhaps after a warning and a failure to remedy matters. This has the additional advantage of providing a disincentive to wasting the community's time with FUD and nonsense.

As mentioned above, here are two more examples fitting the pattern of false disparagement, followed by debunking, with no withdrawal or even acknowledgment:

1. The false accusation that round-3 Kyber "switched from Core-SVP to a modified metric," which was conclusively shown [\[link\]](#) to be based on nothing but the accuser's severe misunderstanding (or worse, deliberate mischaracterization) of what Core-SVP means.

2. The striking accusation that "NIST started trying, with considerable success, to delay and deter public analysis of the patent threats" [\[link\]](#). A follow-up message [\[link\]](#) requested evidence to support this accusation---which was never provided---and showed prior statements from NIST *encouraging* comments about patent issues.

Sincerely yours in cryptography,

Chris

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.

To unsubscribe from this group and stop receiving emails from it, send an email to [pqc-forum+unsubscribe@list.nist.gov](mailto:pqc-forum+unsubscribe@list.nist.gov).

To view this discussion on the web visit <https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/CAC0o0QiT6fuopPFB3mUaH%3DArALy52vcDj8sjQSTQguaGXN-oqA%40mail.gmail.com>.