From: D. J. Bernstein <djb@cr.yp.to> via pqc-forum@list.nist.gov
To: pqc-comments@nist.gov
CC: pqc-forum@list.nist.gov
Subject: [pqc-forum] ROUND 3 OFFICIAL COMMENT: CRYSTALS-KYBER
Date: Wednesday, July 27, 2022 02:33:12 AM ET
Attachments: smime.p7m

NIST's round-3 report claims a better FO proof picture for Kyber than
it does for NTRU. This is wrong——exactly the opposite of what the
literature says on this topic.

Here's what the report says about FO proofs for Kyber:

> The security proofs hold tightly in the ROM [169, 170] and
> non-tightly in the QROM. Yet under various other natural assumptions,
> KYBER may also achieve a tight security reduction in the QROM [184].

Here's what the report says about FO proofs for NTRU:

> The NTRU KEMs have tight CCA-security reductions to the underlying
> PKEs in the ROM, and non-tight security reductions in the QROM.
> Making some additional non-standard assumptions, one of the QROM
> security proofs can be made tight.

This portrays NTRU as having a worse security-proof picture than Kyber:
NTRU needs "non-standard assumptions" for a tight QROM proof, whereas
Kyber "may" have a tight QROM proof under "natural" assumptions.

In fact, the situation for years has been that the literature has better
FO proofs——better tradeoffs between tightness and the strength of the
PKE assumption——for deterministic PKEs than for randomized PKEs:

* ROM, deterministic PKE: https://eprint.iacr.org/2018/526 Theorem
  14.3 obtains IND-CCA2 very tightly from the standard minimal PKE
  security assumption, OW-CPA. (The techniques are older, but this
  paper is designed to support proof verification and identifies
  errors in some previously claimed theorems.)

* ROM, randomized PKE: https://eprint.iacr.org/2017/604 has a proof that's (almost as) tight, but assumes that the PKE is IND-CPA. IND-CPA is still standard, but it's stronger and more complicated than OW-CPA, and has received less attention from cryptanalysts. (See generally Section 6 of https://eprint.iacr.org/2019/691.)

* QROM, deterministic PKE: https://eprint.iacr.org/2019/590 obtains IND-CCA2 tightly from OW-CPA. I'm assuming here that sqrt(epsilon) is allowed as tight, unlike a number-of-queries loss factor.

* QROM, randomized PKE: All tight proofs in the literature make non-standard assumptions, such as the "disjoint simulatability" assumption from the paper [184] that NIST cites. This is stronger than standard assumptions; declaring that it's "natural" doesn't make cryptanalysis magically appear, and doesn't tell us whether the security levels are as high as desired.

The Kyber PKE (like other GAM/LPR variants) is randomized, so it definitely can't use the better proofs. The NTRU PKE is deterministic (since round 2), so presumably the better proofs apply. Someone should check the details of this application, but the risk of an error here doesn't justify NIST making claims that are out of whack with the applicable literature.

NIST's report thus needs an erratum to say that, oops, the report said that NTRU needs a "non-standard assumption" for a tight QROM proof and didn't say this about Kyber, whereas in fact the literature indicates that Kyber needs a non-standard assumption for a tight QROM proof while NTRU doesn't.

If NIST _isn't_ allowing sqrt(epsilon) as "tight", then the report needs to clarify the "tight" dividing line. An erratum is still required for the misinformation that Kyber has a better FO proof picture than NTRU: in fact, Kyber has a worse FO proof picture than NTRU.

This is important because this Kyber proof gap could be hiding a big

security loss. See https://eprint.iacr.org/2021/912 for examples where FO IND-CCA2 security is far below OW-CPA security of the underlying PKE.

This is exactly the "derandomization" risk described in Sections 3.8 and 5.8 of https://ntruprime.cr.yp.to/latticerisks-20211031.pdf, which was filed before NIST's deadline for round-3 input and which, unfortunately, NIST doesn't seem to have read. But simply reading through the previous FO proofs is sufficient to see that NIST's report gets this security comparison backwards.

——D. J. Bernstein

P.S. This is unrelated to the objections that have been raised to the handling of hashing in Kyber's FO security proofs. Qualitatively, those objections are identifying an error in the proofs, which of course is worrisome in a security analysis that NIST's report calls "thorough". However, the idea that someone is going to find a collision in these hash functions is very far down any reasonable list of post-quantum security risks; and plugging in known indifferentiability results closes the proof gap at the expense of a quantitatively minor loss of tightness. Derandomization is a much bigger issue.

P.P.S. Kyber has had a new version in every round, and presumably one should wait to see the next version before filing comments on it, so I'm filing this is a round-3 comment. However, unless there's a radical change in Kyber, I would expect the same comment to continue to apply.

P.P.P.S. This comment is of course also regarding NTRU, which NIST's report says NIST could still select. The underlying issues are also applicable to the split between deterministic PKEs and randomized PKEs in other submissions, although unfortunately NIST's report is structured in a way that obfuscates such comparisons.