From: Ward Beullens <ward@beullens.com>
Sent: Monday, October 26, 2020 5:34 AM
To: pqc-comments
Cc: pqc-forum
Subject: ROUND 3 OFFICIAL COMMENT: Rainbow

Hi everyone,

I found 2 new attacks on the Rainbow signature scheme. The first attack uses the ideas from the Kipnis-Shamir attack [1] and reduces the security of Rainbow I, III and V by 7 bits, 4 bits and 19 bits respectively. This attack also applies to the UOV scheme.

The second attack is specific to Rainbow and is also more efficient. It reduces a key recovery to a new instance of the MinRank problem, which can then be solved with the methods of Bardet et al.[2] . This attack reduces the security level of Rainbow I, III and V by 20 bits, 40 bits and 55 bits respectively.

The paper is on ePrint: https://eprint.iacr.org/2020/1343

I would like to thank the Rainbow team for reviewing an earlier version of the paper and providing me with helpful feedback.

All the best,
Ward

[1] https://link.springer.com/chapter/10.1007/BFb0055733
[2] https://arxiv.org/abs/2002.08322

From: Bo-Yin Yang <moscito@gmail.com>
Sent: Monday, October 26, 2020 9:46 AM
To: pqc-forum
Cc: Ward Beullens; pqc-forum; pqc-comments
Subject: Re: ROUND 3 OFFICIAL COMMENT: Rainbow

Dear All,

The Rainbow team acknowledges that the new attacks by Ward Beullens are fundamentally correct and does cut down the security of Rainbow somewhat, in the same manner as most other schemes still in the NIST competition have lost a considerable number of security bits. We are actively studying the consequences and will formulate a comprehensive response including a new set of parameters soon.

Bo-Yin, for the Rainbow team members

On Monday, October 26, 2020 at 5:34:41 PM UTC+8 Ward Beullens wrote:

> Hi everyone,
>
> I found 2 new attacks on the Rainbow signature scheme. The first attack uses the ideas from the Kipnis-Shamir attack [1] and reduces the security of Rainbow I, III and V by 7 bits, 4 bits and 19 bits respectively. This attack also applies to the UOV scheme.
>
> The second attack is specific to Rainbow and is also more efficient. It reduces a key recovery to a new instance of the MinRank problem, which can then be solved with the methods of Bardet et al.[2] . This attack reduces the security level of Rainbow I, III and V by 20 bits, 40 bits and 55 bits respectively.
>
> The paper is on ePrint: https://eprint.iacr.org/2020/1343
>
> I would like to thank the Rainbow team for reviewing an earlier version of the paper and providing me with helpful feedback.
>
> All the best,
> Ward
>
> [1] https://link.springer.com/chapter/10.1007/BFb0055733
> [2] https://arxiv.org/abs/2002.08322

From: Ward Beullens <ward@beullens.com>
Sent: Friday, February 25, 2022 2:59 AM
To: pqc-comments
Cc: pqc-forum
Subject: ROUND 3 OFFICIAL COMMENT: Rainbow

Hi everyone,

I found a practical key-recovery attack against Rainbow. With this attack my laptop can recover a private key for the SL I parameter set from the second round in on average only 53 hours, using the implementation of block Wiedemann XL by Niederhagen et al. [1]. The same attack against the SL I parameters for the third-round submission would be more costly by only a factor 2^8.

Details of the attack (and an extension of the attack that performs better at higher security levels) are on ePrint:
https://eprint.iacr.org/2022/214

All the best,
Ward

[1] http://polycephaly.org/projects/xl/

From:	Jintai Ding <jintai.ding@gmail.com>
Sent:	Saturday, February 26, 2022 10:37 PM
To:	Ward Beullens
Cc:	pqc-comments; pqc-forum
Subject:	Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

Dear Ward,

We would like to thank you for sharing your results.

We implemented and tested his attack and it works. We made a mistake in missing this simple attack, and we would like to thank you for the great work.

We proposed to Nist to replace the Rainbow Level 1 parameters with our Level 3 parameters and Level 3 with Level 5 parameters.

Cheers, JT on behalf of Rainbow team.

On Fri, Feb 25, 2022 at 2:59 AM Ward Beullens <ward@beullens.com> wrote:

> Hi everyone,
>
> I found a practical key-recovery attack against Rainbow. With this attack my laptop can recover a private key for the SL I parameter set from the second round in on average only 53 hours, using the implementation of block Wiedemann XL by Niederhagen et al. [1]. The same attack against the SL I parameters for the third-round submission would be more costly by only a factor 2^8.
>
> Details of the attack (and an extension of the attack that performs better at higher security levels) are on ePrint: https://eprint.iacr.org/2022/214
>
> All the best,
>
> Ward
>
> [1] http://polycephaly.org/projects/xl/
>
> --
> 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.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/C5A83F59-86F7-429A-A35B-67B0A03612E0%40hxcore.ol.

From: Quan Thoi Minh Nguyen <msuntmquan@gmail.com>
Sent: Sunday, February 27, 2022 3:16 PM
To: pqc-forum
Cc: jinta...@gmail.com; pqc-comments; pqc-forum; Ward Beullens
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

On Saturday, February 26, 2022 at 7:37:27 PM UTC-8 jinta...@gmail.com wrote:

> Dear Ward,
>
> We would like to thank you for sharing your results.
>
> We implemented and tested his attack and it works. We made a mistake in missing this simple attack, and we would like to thank you for the great work.
>
> We proposed to Nist to replace the Rainbow Level 1 parameters with our Level 3 parameters and Level 3 with Level 5 parameters.

Considering the devastating attack is confirmed by Rainbow's author, I think the following actions are reasonable security-wise:

1/ Rainbow should *not* be considered to advance to the next round.

2/ Equivalently important, NIST should revise NIST PQC finalist documentation and make clear the security risk of Rainbow. This is to avoid the case that people would make a *wrong* assumption that NIST PQC finalists have reasonable security while Rainbow doesn't have enough security confidence.

- Quan

> Cheers, JT on behalf of Rainbow team.
>
> On Fri, Feb 25, 2022 at 2:59 AM Ward Beullens <wa...@beullens.com> wrote:
>
>> Hi everyone,
>>
>> I found a practical key-recovery attack against Rainbow. With this attack my laptop can recover a private key for the SL I parameter set from the second round in on average only 53 hours, using the implementation of block Wiedemann XL by Niederhagen et al. [1]. The same attack against the SL I parameters for the third-round submission would be more costly by only a factor 2^8.
>>
>> Details of the attack (and an extension of the attack that performs better at higher security levels) are on ePrint: https://eprint.iacr.org/2022/214
>>
>> All the best,

From: Ruben Niederhagen <ruben@polycephaly.org>
Sent: Monday, February 28, 2022 3:08 AM
To: pqc-comments
Cc: pqc-forum
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

Dear Quan,

May I ask you to quantify in what way you consider the attack by Ward Beullens to be "devastating"? While the Level 1 parameter sets clearly are broken by the attack, its impact on the larger parameter sets is much more "moderate".

I believe the conclusion by Ward in [1] sums it up quite clearly:

"In principle, it would be possible to move to larger parameters to protect against the attacks presented in this paper, at the cost of larger key sizes and signature sizes. E.g., the SL 3 parameters of the third-round submission seem to provide enough security for SL 1, but those parameters have signatures and public keys that are larger by a factor 2.5 and 4.4 respectively compared to the SL 1 parameters.
However, there seems to be some room for improvement for the attacks in Section 4, so more cryptanalysis would be required before we can have confidence in the security of Rainbow. Moreover, the resulting Rainbow signature scheme would be less efficient than the Oil and Vinegar scheme. So there is seemingly no reason to prefer Rainbow over the Oil and Vinegar scheme, on which Rainbow is based, and which is older, simpler, and has a strictly smaller attack surface in comparison to Rainbow. (E.g., none of the attacks in this paper seem to apply to the Oil and Vinegar scheme)."

So, currently (pending further analysis), the parameters newly proposed by the Rainbow team earlier in this thread provide reasonable security.
Weather or not Rainbow with these parameters is competitive in regard to performance is a different (though valid) question.

Best regards
Ruben

[1] https://eprint.iacr.org/2022/214

From: s zhang <shezhangth@gmail.com>
Sent: Monday, February 28, 2022 3:31 AM
To: pqc-forum
Cc: Ruben Niederhagen; pqc-forum; pqc-comments
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

What will happen to ABCMint in the future and will it become less secure?

Ruben Niederhagen 在 2022 年 2 月 28 日 星期一下午 4:07:57 [UTC+8] 的信中寫道：

> Dear Quan,
>
> May I ask you to quantify in what way you consider the attack by Ward Beullens to be ″devastating″? While the Level 1 parameter sets clearly are broken by the attack, its impact on the larger parameter sets is much more ″moderate″.
>
> I believe the conclusion by Ward in [1] sums it up quite clearly:
>
> ″In principle, it would be possible to move to larger parameters to protect against the attacks presented in this paper, at the cost of larger key sizes and signature sizes. E.g., the SL 3 parameters of the third-round submission seem to provide enough security for SL 1, but those parameters have signatures and public keys that are larger by a factor 2.5 and 4.4 respectively compared to the SL 1 parameters. However, there seems to be some room for improvement for the attacks in Section 4, so more cryptanalysis would be required before we can have confidence in the security of Rainbow. Moreover, the resulting Rainbow signature scheme would be less efficient than the Oil and Vinegar scheme. So there is seemingly no reason to prefer Rainbow over the Oil and Vinegar scheme, on which Rainbow is based, and which is older, simpler, and has a strictly smaller attack surface in comparison to Rainbow. (E.g., none of the attacks in this paper seem to apply to the Oil and Vinegar scheme).″
>
> So, currently (pending further analysis), the parameters newly proposed by the Rainbow team earlier in this thread provide reasonable security. Weather or not Rainbow with these parameters is competitive in regard to performance is a different (though valid) question.
>
> Best regards
> Ruben

From:	D. J. Bernstein <djb@cr.yp.to>
Sent:	Monday, February 28, 2022 7:12 AM
To:	pqc-comments
Cc:	pqc-forum
Subject:	Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow
Attachments:	signature.asc

Quan Thoi Minh Nguyen writes:

> 2/ Equivalently important, NIST should revise NIST PQC finalist
> documentation and make clear the security risk of Rainbow. This is to
> avoid the case that people would make a *wrong* assumption that NIST
> PQC finalists have reasonable security while Rainbow doesn’t have
> enough security confidence.

I agree that there has to be a warning about the drop of security levels for Rainbow. The issue isn't just the demonstrated attack on rainbow1a:
the new paper says 2^131 operations for rainbow3c and 2^164 operations for rainbow5c, where the round-1 submission in 2017 says 2^217.4 and
2^275.4 respectively. (My impression is that the operations being counted are of sufficiently similar types that a serious AT analysis would show a security drop similar to what these numbers suggest.)

But surely NIST also has to catch up on issuing warnings about the drop of security levels for lattice systems. One easy way to see the drop is to rewind to "Better Key Sizes (and Attacks) for LWE-Based Encryption"
from Lindner and Peikert:

https://eprint.iacr.org/2010/613.pdf

The Frodo submission says it's an "instantiation and implementation" of that paper with modifications. The security levels conjectured for Frodo and other lattice systems today are much lower than in that paper:

* In that paper, Lindner and Peikert propose using dimension 256, giving matrix sizes of 400 kilobits, and, for the ring version, public keys of just 2 kilobits, i.e., 256 bytes. They evaluate attacks against these parameters as using "2^120 seconds" on a 2.3GHz core, and conclude that these parameters "appear to be at least as secure as AES-128".

* The current version of Frodo targeting the AES-128 security level uses much larger parameters for this, dimension 640, internally generating matrices on the scale of a megabyte. Similarly, the most aggressive Kyber parameters use dimension 512 and use 800-byte public keys, obviously far above 256 bytes.

If there's supposed to be a dividing line saying that Rainbow needs a warning and lattices don't, let's hear a clear definition of this dividing line and an explanation of why the dividing line is justified.
By default I would think that warnings are required for both.

As an analogy, imagine a report comparing QKD to public-key cryptography and saying that one can't trust the security of public-key cryptography.
Exhibit A is an attack demonstrated against a NISTPQC finalist. The report would also be obliged to point out the terrible security track record of QKD, right?

---D. J. Bernstein

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Tuesday, March 1, 2022 12:39 PM
To: pqc-comments
Cc: pqc-forum
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow
Attachments: signature.asc

I wrote, in pqc-forum email dated 28 Feb 2022 13:12:02 +0100:
> As an analogy, imagine a report comparing QKD to public-key
> cryptography and saying that one can't trust the security of public-key cryptography.
> Exhibit A is an attack demonstrated against a NISTPQC finalist. The
> report would also be obliged to point out the terrible security track
> record of QKD, right?

At "2022-02-28T22:26:54.787Z" the following article appeared:

https://web.archive.org/web/20220228163622/https://medium.com/cambridge-quantum-computing/what-does-the-breaking-of-rainbow-mean-for-cybersecurity-21b125383cab

The article compares QKD to public-key cryptography and says that one can't trust the security of public-key cryptography. Exhibit A is what you'd expect. The report _doesn't_ point out the terrible security track record of QKD. On the contrary, it tries to make the reader believe that the security of QKD is guaranteed by "the fundamental laws of nature", in much the same way that, e.g., a broken dimension-256 cryptosystem was labeled as "theoretically sound" in 2010 Lindner--Peikert.

Given the ethical requirement to "be transparent and provide full disclosure of all pertinent system capabilities, limitations, and potential problems to the appropriate parties", I don't see how it can be acceptable for a report on security of PKC and QKD to highlight risks of PKC while pretending that there are no risks of QKD.

For a report focusing on the NIST finalists, QKD would be out of scope, but I don't see how it can be acceptable for such a report to have a discussion of security degradation that's limited to Rainbow and that conceals the security degradation of lattices. NIST's previous reports on NISTPQC were strikingly selective in their discussion of security degradation; I hope this is corrected in NIST's next report.

On a related note, it's disturbing to see NSA's continued efforts to convince people to _turn off ECC_ in favor of lattices. NIST should endorse ANSSI's statement that "the maturity level of the post-quantum algorithms presented to the NIST process should not be overestimated"
and should join other organizations in recommending that post-quantum algorithms be deployed _only_ as a second layer of encryption (and/or
signatures) together with ECC.

---D. J. Bernstein

From: Christopher J Peikert <cpeikert@alum.mit.edu>
Sent: Tuesday, March 1, 2022 3:26 PM
To: pqc-forum; pqc-comments
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

On Mon, Feb 28, 2022 at 7:12 AM D. J. Bernstein <djb@cr.yp.to> wrote:

> If there's supposed to be a dividing line saying that Rainbow needs a
> warning and lattices don't, let's hear a clear definition of this
> dividing line and an explanation of why the dividing line is justified.

The distinction is obvious and self explanatory: for the Rainbow NISTPQC submission, we now have a conclusive demonstration that its proposed parameters fall far short of their targeted NIST security levels.

For the remaining lattice submissions, all known attacks -- and even optimistic extrapolations thereof -- are quite far from endangering the targeted security levels.

Moreover, even for the cited pre-NISTPQC lattice parameters, it's not clear whether a "serious AT analysis" -- the metric endorsed in the quoted message and its author's research -- of known attacks would actually imply sub-AES-128 security, when accounting for large memory size/accesses and other real costs.

Such an analysis (with realistic modeling choices) could potentially lead to that conclusion, but from an initial look it doesn't seem clear cut. In my opinion it is too close for comfort, but that is a very different matter from affirmatively claiming sub-AES-128 security.

Without such an analysis, it is not justified to call those parameters "broken," as the same commenter just did in a follow-up message. (Insert "where's the erratum?!?" here, if you wish.)

> The security levels conjectured for Frodo
> and other lattice systems today are much lower than in that paper:
>
> * In that paper, Lindner and Peikert propose using dimension 256... and conclude that these parameters "appear to be at
> least as secure as AES-128".
>
> * The current version of Frodo targeting the AES-128 security level
> uses much larger parameters for this, dimension 640... Similarly, the most
> aggressive Kyber parameters use dimension 512...

This is (yet another) quite misleading comparison, due to the very different metrics and security margins these works use to arrive at security bounds.

Compared to LP'11, FrodoKEM and Kyber aim to give *lower bounds* on security that are a good deal lower than the *upper bounds* implied by known attacks. This is due to the use of much more attacker-friendly metrics:

- They ignore the high cost of the huge memory needed for sieving in large dimensions.
- They focus on the "Core-SVP" cost of solving a single SVP in a large enough dimension, ignoring the cost of multiple SVP calls and the surrounding lattice reduction.
- They ignore the bit-operation costs of working with real numbers, etc.

Naturally, this will yield larger parameters to get above a desired security lower bound. But it also gives more room for plausible future improvements in cryptanalysis, like time-memory tradeoffs, amortizing SVP calls, etc.

(Whether it's "good" to use these attacker-friendly metrics is beside the point; the point is the big difference in methodology between LP'11 and FrodoKEM/Kyber, making their security bounds incomparable.)

This does *not* mean that somewhat smaller FrodoKEM/Kyber parameters are less secure than AES-128 -- they likely aren't, but would deserve less confidence. In other words, the proposed parameters probably are not *necessary* for AES-128-level security, but there is good reason to think they are comfortably more than sufficient.

Sincerely yours in cryptography,
Chris

From: Bank BSC <bankofbsc@gmail.com>
Sent: Tuesday, March 1, 2022 3:30 PM
To: pqc-forum
Cc: s zhang; Ruben Niederhagen; pqc-forum; pqc-comments
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

It means anyone can transfer your ABCMint tokens to his/her wallet within a weekend.

On Monday, February 28, 2022 at 3:31:29 AM UTC-5 s zhang wrote:

> What will happen to ABCMint in the future and will it become less secure?
>
> Ruben Niederhagen 在 2022 年 2 月 28 日 星期一下午 4:07:57 [UTC+8] 的信中寫道：
>
>> Dear Quan,
>>
>> May I ask you to quantify in what way you consider the attack by Ward Beullens to be ″devastating″? While the Level 1 parameter sets clearly are broken by the attack, its impact on the larger parameter sets is much more ″moderate″.
>>
>> I believe the conclusion by Ward in [1] sums it up quite clearly:
>>
>> ″In principle, it would be possible to move to larger parameters to protect against the attacks presented in this paper, at the cost of larger key sizes and signature sizes. E.g., the SL 3 parameters of the third-round submission seem to provide enough security for SL 1, but those parameters have signatures and public keys that are larger by a factor 2.5 and 4.4 respectively compared to the SL 1 parameters. However, there seems to be some room for improvement for the attacks in Section 4, so more cryptanalysis would be required before we can have confidence in the security of Rainbow. Moreover, the resulting Rainbow signature scheme would be less efficient than the Oil and Vinegar scheme. So there is seemingly no reason to prefer Rainbow over the Oil and Vinegar scheme, on which Rainbow is based, and which is older, simpler, and has a strictly smaller attack surface in comparison to Rainbow. (E.g., none of the attacks in this paper seem to apply to the Oil and Vinegar scheme).″
>>
>> So, currently (pending further analysis), the parameters newly proposed by the Rainbow team earlier in this thread provide reasonable security. Weather or not Rainbow with these parameters is competitive in regard to performance is a different (though valid) question.
>>
>> Best regards

From: EL HASSANE LAAJI <e.laaji@ump.ac.ma>
Sent: Tuesday, March 1, 2022 5:30 PM
To: pqc-forum; pqc-comments
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

Hi
# I don't see how it can
be acceptable for a report on security of PKC and QKD to highlight risks
of PKC while pretending that there are no risks of QKD.#

QKD uses two channels (BB84): Quantum channel to send quantum information. and classical channel to exchange measurement information. That means we should warranting the security over classical channel.

Then we must combining PKE&QKD.
Best regards
Le mardi 1 mars 2022, D. J. Bernstein <djb@cr.yp.to> a écrit :

> I wrote, in pqc-forum email dated 28 Feb 2022 13:12:02 +0100:
> > As an analogy, imagine a report comparing QKD to public-key cryptography
> > and saying that one can't trust the security of public-key cryptography.
> > Exhibit A is an attack demonstrated against a NISTPQC finalist. The
> > report would also be obliged to point out the terrible security track
> > record of QKD, right?
>
> At "2022-02-28T22:26:54.787Z" the following article appeared:
>
> https://web.archive.org/web/20220228163622/https://medium.com/cambridge-quantum-computing/what-does-the-breaking-of-rainbow-mean-for-cybersecurity-21b125383cab
>
> The article compares QKD to public-key cryptography and says that one can't trust the security of public-key cryptography. Exhibit A is what you'd expect. The report _doesn't_ point out the terrible security track record of QKD. On the contrary, it tries to make the reader believe that the security of QKD is guaranteed by "the fundamental laws of nature", in much the same way that, e.g., a broken dimension-256 cryptosystem was labeled as "theoretically sound" in 2010 Lindner--Peikert.
>
> Given the ethical requirement to "be transparent and provide full disclosure of all pertinent system capabilities, limitations, and potential problems to the appropriate parties", I don't see how it can be acceptable for a report on security of PKC and QKD to highlight risks of PKC while pretending that there are no risks of QKD.
>
> For a report focusing on the NIST finalists, QKD would be out of scope, but I don't see how it can be acceptable for such a report to have a discussion of security degradation that's limited to Rainbow and that conceals the security degradation of lattices. NIST's previous reports on NISTPQC were strikingly selective in their discussion of security degradation; I hope this is corrected in NIST's next report.

From: 'Scott Fluhrer (sfluhrer)' via pqc-forum <pqc-forum@list.nist.gov>
Sent: Tuesday, March 1, 2022 6:00 PM
To: EL HASSANE LAAJI; pqc-forum; pqc-comments
Subject: RE: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

From: pqc-forum@list.nist.gov <pqc-forum@list.nist.gov> **On Behalf Of** EL HASSANE LAAJI
Sent: Tuesday, March 1, 2022 5:30 PM
To: pqc-forum@list.nist.gov; pqc-comments@nist.gov
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow

Hi
# I don't see how it can
be acceptable for a report on security of PKC and QKD to highlight risks
of PKC while pretending that there are no risks of QKD.#

QKD uses two channels (BB84): Quantum channel to send quantum information. and classical channel to exchange measurement information. That means we should warranting the security over classical channel.

Then we must combining PKE&QKD.

While I'm not a fan of QKD, I feel I must come to its defense – QKD systems have their problems, but I don't believe this is one of them.

The classical channel does not need to be private; it does need to be authenticated. I believe what is used in practice is a Message Authentication Code; this assumes that the two sides share a secret key (which could be updated over time using some bits exchanged by the QKD system which are not exposed); this is quantum safe as long as the MAC key is long enough.

Of course, if you do share a secret key, why don't you use that to generate the key bits, and not bother with Quantum hardware (and its expense)? When I ask QKD vendors that, they start talking about the bootstrapping problem – I was never convinced that was sufficient justification for all the downsides of QKD.

--

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Wednesday, March 2, 2022 7:56 AM
To: pqc-comments
Cc: pqc-forum
Subject: Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: Rainbow
Attachments: signature.asc

Christopher J Peikert writes:
> The distinction is obvious and self explanatory: for the Rainbow
> NISTPQC submission, we now have a conclusive demonstration that its
> proposed parameters fall far short of their targeted NIST security levels.

True for rainbow1a, but false for rainbow3c and rainbow5c. Furthermore, asking whether attacks have been _conclusively demonstrated_ is clearly setting the bar too low for issuing warnings.

I already drew this distinction ("The issue isn't just the demonstrated attack on rainbow1a: the new paper says 2^131 operations for rainbow3c and 2^164 operations for rainbow5c, where the round-1 submission in 2017 says 2^217.4 and 2^275.4 respectively") and advocated warnings based on the _drop_ of security levels.

Nobody is disputing the need for warnings regarding the demonstrated break. The question is what to do beyond that: in particular, how to warn people about security degradation.

The official NISTPQC evaluation criteria express concern about "schemes that were designed by repeatedly patching older schemes that were shown vulnerable to cryptanalysis". What we're seeing with Rainbow is a patch, an increase of key sizes, in response to cryptanalysis. But we've also seen lattices increasing key sizes in response to cryptanalysis, as illustrated by the jump from dimension 256 to dimensions 512/640 (in both cases for matching AES-128), a jump big enough to be clearly visible to everybody despite neverending lattice-security obfuscation.

Surely NIST should be consistently issuing warnings regarding all of these security losses. I'm concerned by the selectivity of warnings in NIST's earlier reports, the same way that I'm concerned by QKD-vs.-PKC reports that highlight PKC failures and conceal QKD failures.

> For the remaining lattice submissions, all known attacks -- and even
> optimistic extrapolations thereof -- are quite far from endangering
> the targeted security levels.

False. For example, the round-3 Kyber submission said that known attacks against the newly proposed round-3 Kyber-512 could use just 2^135.5 "gates", i.e., 2^8 times fewer "gates" than AES-128 key search. There have been further attacks since then: e.g., an attack at Asiacrypt 2021 that appears to further reduce the security of round-3 Kyber-512.

The Kyber submission argues that the costs of RAM ("the massive memory requirements that are ignored in the gate-count metric") provide a security margin---but surely such arguments should be applied consistently across submissions!

Furthermore, NIST wrote 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." So the question of whether Kyber-512 reaches "category 1" is important for whether Kyber-512 is kept at all, whereas the question of whether rainbow5c reaches "category 5" could be a mere matter of labeling.

The real reason for issuing a warning about rainbow5c is not the comparison to AES-256, but rather the drop of estimated costs of known attacks, from 2^275.4 operations in 2017 to 2^164 operations in 2022.
It's not that 2^164 operations will happen in the foreseeable future; it's that one has to worry that this isn't the end of the story. But, for the same reason, surely there also have to be warnings about the big advances in lattice attacks that have forced dimension-256 lattice recommendations in 2010 to be replaced by much higher dimensions today.

> Moreover, even for the cited pre-NISTPQC lattice parameters, it's not
> clear whether a "serious AT analysis" -- the metric endorsed in the
> quoted message and its author's research -- of known attacks would
> actually imply
> sub-AES-128 security, when accounting for large memory size/accesses
> and other real costs.

Structurally, are you seriously arguing that the uncertainties in the literature regarding the costs of lattice attacks are a reason to _not_ issue warnings regarding advances in lattice attacks, such as warnings that the 2010 Lindner--Peikert "theoretically sound" dimension-256 lattice parameters are broken and should not be used?

Quantitatively, the best estimates available are that the real costs of sieving (never mind the latest enumeration advances) are at roughly 30% higher security levels than a naive free-RAM analysis would suggest.
This certainly isn't enough to rescue that dimension-256 proposal.

> Compared to LP'11, FrodoKEM and Kyber aim to give *lower bounds* on
> security that are a good deal lower than the *upper bounds* implied by
> known attacks.

This is partially correct (see below), but not meaningful as stated, since it would be satisfied by giving a lower bound of, e.g., 0. What's missing here is a requirement for the "lower bound" to be as secure as AES-128, the minimum allowed security level.

Does Kyber have a lower-bound analysis meeting this requirement? Let's look at the submission.

The original Kyber submission claimed repeatedly that its analysis was "conservative". But it then reported that the result of the analysis for
Kyber-512 was only 2^112 (see Table 3)---which, oops, is much smaller than 2^128.

The submission then (see bottom of page 17) stated various unquantified arguments for the claim that Kyber-512 was as hard to break as AES-128.
This AES-128 conclusion is not the "conservative" analysis and was never claimed to be "conservative"---even though it's easy to see how readers can be misled into thinking otherwise.

Furthermore, a careful look at the supposedly "conservative" analysis (see Section 6 of the round-2 NTRU Prime submission) shows that the analysis combines a series of overestimates, underestimates, potential overestimates, and potential underestimates. Selectively pointing to the underestimates does _not_ justify the claim that this is a "lower bound".

In fact, all available evidence is that the output of the "conservative"
analysis is _above_ the actual attack costs for all sufficiently large parameters. See my pqc-forum email dated 30 May 2020 02:15:31 +0200.
(The issue, in a nutshell, is that "dimensions for free" appears to be a superpolynomial speedup while all of the slowdowns appear to be polynomially bounded.)

The round-3 Kyber submission tried to quantify various issues, leaving a 30-bit range of uncertainty for the number of "gates" in known attacks against round-3 Kyber-512 (while round-2 Kyber-512 has been abandoned).
As noted above, the low end is hundreds of times fewer "gates" than
AES-128 key search.

> - They ignore the high cost of the huge memory needed for sieving in
> large dimensions.

False. The "conservative" analysis ignores this, but the Kyber submission doesn't. For example, as noted above, the current Kyber submission points to "the massive memory requirements that are ignored in the gate-count metric" as a reason to think that it's okay to be breakable by an attack using hundreds of times fewer "gates" than
AES-128 key search.

Page 17 of the original Kyber submission similarly pointed to "the cost of access into exponentially large memory" as part of its unquantified argument that round-1 Kyber-512 was as hard to break as AES-128.

> - They focus on the "Core-SVP" cost of solving a single SVP in a large
> enough dimension, ignoring the cost of multiple SVP calls and
> the surrounding lattice reduction.

False. This is something else that the "conservative" analysis ignores but that the Kyber submission has never ignored.

For example, page 17 of the original Kyber submission pointed to "the
(polynomial) number of calls to the SVP oracle" as another part of its unquantified argument that round-1 Kyber-512 was as hard to break as AES-128. The latest Kyber submission tries to quantify this, and in particular claims (top of page 27) that this adds 11-12 bits of security, but then says (page 29) that it expects to lose 2-8 bits from this because of known speedups.

> - They ignore the bit-operation costs of working with real numbers, etc.

False. For example, page 17 of the original Kyber submission pointed to "the gate count required for one 'operation'" as yet another part of its unquantified argument that round-1 Kyber-512 was as hard to break as AES-128. Again the latest Kyber submission tries to quantify this and leaves considerable uncertainties.

Of course, the 2^128 operations to search all AES-128 keys also involve many bit operations per key. This is why NIST set the minimum at 2^143 "gates". This is also why the original Kyber submission tried to argue that it would gain a factor above 2^30 compared to its 2^112, rather than just a factor 2^16.

> Naturally, this will yield larger parameters to get above a desired
> security lower bound. But it also gives more room for plausible future
> improvements in cryptanalysis, like time-memory tradeoffs, amortizing
> SVP calls, etc.

This is misinformation regarding the Kyber submission.

> (Whether it's "good" to use these attacker-friendly metrics is beside
> the point; the point is the big difference in methodology between
> LP'11 and FrodoKEM/Kyber, making their security bounds incomparable.)

To be clear, are you continuing to propose the "theoretically sound"
dimension-256 lattice systems from 2010 Lindner--Peikert as matching

AES-128 security (even better, requiring "2^120 seconds" to break on a 2.3GHz core), despite all the subsequent advances in lattice attacks?

---D. J. Bernstein